



# Alcatel-Lucent 7450

ETHERNET SERVICE SWITCH | RELEASE 13.0.R1

TRIPLE PLAY SERVICE DELIVERY ARCHITECTURE

Alcatel-Lucent Proprietary  
This document contains proprietary information of Alcatel-Lucent and is not to be disclosed or used except in  
accordance with applicable agreements.  
Copyright 2015 © Alcatel-Lucent. All rights reserved.

All specifications, procedures, and information in this document are subject to change and revision at any time without notice. The information contained herein is believed to be accurate as of the date of publication. Alcatel-Lucent provides no warranty, express or implied, regarding its contents. Users are fully responsible for application or use of the documentation.

Alcatel, Lucent, Alcatel-Lucent and the Alcatel-Lucent logo are trademarks of Alcatel-Lucent. All other trademarks are the property of their respective owners.

Copyright 2015 Alcatel-Lucent.

All rights reserved.

### **Disclaimers**

Alcatel-Lucent products are intended for commercial uses. Without the appropriate network design engineering, they must not be sold, licensed or otherwise distributed for use in any hazardous environments requiring fail-safe performance, such as in the operation of nuclear facilities, aircraft navigation or communication systems, air traffic control, direct life-support machines, or weapons systems, in which the failure of products could lead directly to death, personal injury, or severe physical or environmental damage. The customer hereby agrees that the use, sale, license or other distribution of the products for any such application without the prior written consent of Alcatel-Lucent, shall be at the customer's sole risk. The customer hereby agrees to defend and hold Alcatel-Lucent harmless from any claims for loss, cost, damage, expense or liability that may arise out of or in connection with the use, sale, license or other distribution of the products in such applications.

This document may contain information regarding the use and installation of non-Alcatel-Lucent products. Please note that this information is provided as a courtesy to assist you. While Alcatel-Lucent tries to ensure that this information accurately reflects information provided by the supplier, please refer to the materials provided with any non-Alcatel-Lucent product and contact the supplier for confirmation. Alcatel-Lucent assumes no responsibility or liability for incorrect or incomplete information provided about non-Alcatel-Lucent products.

However, this does not constitute a representation or warranty. The warranties provided for Alcatel-Lucent products, if any, are set forth in contractual documentation entered into by Alcatel-Lucent and its customers.

This document was originally written in English. If there is any conflict or inconsistency between the English version and any other version of a document, the English version shall prevail.

# Table of Contents

<b>Preface</b>	19
About This Guide	19
Audience	19
List of Technical Publications	19
Technical Support	21
<b>Getting Started</b>	
In This Chapter	23
Alcatel-Lucent 7450 ESS-Series Services Configuration Process	23
<b>Introduction to Triple Play</b>	
In This Section	25
Alcatel-Lucent's Triple Play Service Delivery Architecture	26
Introduction to Triple Play	26
Blueprint for Optimizing Triple Play Service Infrastructures	27
Architectural Foundations	28
Optimizing Triple Play Service Infrastructures	30
Distributed Service Edges	31
Service Differentiation, QoS Enablement	33
Virtual MAC Subnetting for VPLS	37
Services	40
Service Types	41
Service Policies	42
Alcatel-Lucent Service Model	43
Introduction	43
Service Entities	44
Customers	44
Service Access Points (SAPs)	45
SAP Encapsulation Types and Identifiers	46
Ethernet Encapsulations	46
SAP Considerations	47
Service Distribution Points (SDPs)	49
SDP Binding	49
Spoke and Mesh SDPs	50
SDP Encapsulation Types	51
SDP Keepalives	52
Epipe Service Overview	53
VPLS Service Overview	54
Split Horizon SAP Groups and Split Horizon Spoke SDP Groups	54
Residential Split Horizon Groups	55
IES Service Overview	56
IP Interface	57
Deploying and Provisioning Services	58
Phase 1: Core Network Construction	58
Phase 2: Service Administration	58

## Table of Contents

Phase 3: Service Provisioning .....	58
Configuration Notes .....	59
General .....	59
Configuring Triple Play Services with CLI .....	61
Configuring VPLS Residential Split Horizon Groups .....	61
Configuring Static Hosts .....	62
Triple Play Services Command Reference .....	65
Configuration Commands .....	65
Triple Play Service Configuration Commands .....	79

## DHCP Management

In This Chapter .....	313
DHCP Principles .....	314
DHCP Features .....	316
DHCP Relay .....	316
DHCP Relay Enhancements .....	316
Subscriber Identification Using Option 82 Field .....	319
Trusted and Untrusted .....	321
DHCP Snooping .....	322
DHCP Lease State Table .....	324
DHCP and Layer 3 Aggregation .....	326
Local DHCP Servers .....	328
Overview .....	328
Local DHCP Server Support .....	330
DHCPv6 .....	331
DHCPv6 Relay Agent .....	331
DHCPv6 Prefix Options .....	331
Neighbor Resolution via DHCPv6 Relay .....	331
DHCPv6 Lease Persistency .....	332
Local Proxy Neighbor Discovery .....	332
IPv6oE Hosts Behind Bridged CPEs .....	333
IPv6 Link-Address Based Pool Selection .....	333
IPv6 Address/Prefix Stickiness .....	333
IPv4/v6 Linkage for Dual-Stack Hosts or Layer 3 RGs .....	333
Host Connectivity Checks for IPv6 .....	334
DHCP Relay Enhancements .....	335
Flexible Host Identification in LUDB Based on DHCPv4/v6 Options .....	336
DHCP Caching .....	336
Flexible Creation of DHCPv4/6 Host Parameters Utilizing Python and Internal Caching .....	337
Python DTC Variables and API .....	338
DTC Debugging Facility .....	342
Virtual Subnet for DHCPv4 Hosts .....	343
Proxy DHCP Server .....	344
Local DHCP Servers .....	348
Terminology .....	348
Overview .....	350
DHCP Lease Synchronization .....	353
Intercommunication Link Failure Detection .....	354
DHCP Server Failover States .....	355



Lease Time Synchronization	356
Maximum Client Lead Time (MCLT)	357
Sharing IPv4 Address-Range or IPv6 Prefix Between Redundant 7x50 DHCP Servers in Access-Driven Mode	360
Fast-Switchover of IP Address/Prefix Delegation For Remote IP Address/Prefix Range	364
Local address assignment	366
Stateless Address Auto-configuration	366
Configuring DHCP with CLI	367
Common Configuration Tasks	368
Enabling DHCP Snooping	368
Configuring Option 82 Handling	370
Enabling DHCP Relay	371
Configuring Local User Database Parameters	372
Triple Play DHCP Command Reference	379
Configuration Commands	379
Triple Play DHCP Configuration Commands	405

## Triple Play Security

In This Chapter	545
Triple Play Security Features	546
Anti-Spoofing Filters	546
Anti-spoofing Filter Types	547
Filtering Packets	547
Layer 2 Triple Play Security Features	548
MAC Pinning	548
MAC Protection	548
DoS Protection	549
VPLS Redirect Policy	551
ARP Handling	552
ARP Reply Agent	552
Dynamic ARP Table Population	553
Local Proxy ARP	553
Web Portal Redirect	554
Configuring Triple Play Security with CLI	557
Common Configuration Tasks	558
Configuring Anti-Spoofing Filters	558
Configuring Triple Play Security features	559
Configuring ARP Handling	565
Configuring Web Portal Redirect	571
Triple Play Security Command Reference	573
Command Hierarchies	573
Triple Play Security Configuration Commands	575
Show Commands	586
Configuring Triple Play Multicast Services with CLI	587
Configuring IGMP Snooping in the BSA	588
Enabling IGMP Snooping in a VPLS Service	588
With IGMPv3 Multicast Routers	588
With IGMPv1/2 Multicast Routers	589
Modifying IGMP Snooping Parameters	590
Modifying IGMP Snooping Parameters for a SAP or SDP	591

## Table of Contents

Configuring Static Multicast Groups on a SAP or SDP . . . . .	593
Enabling IGMP Group Membership Report Filtering . . . . .	594
Configuring Multicast VPLS Registration (MVR) . . . . .	597
<b>Triple Play Multicast</b>	
In This Chapter . . . . .	599
Introduction to Multicast . . . . .	601
Multicast in the Broadband Service Router . . . . .	602
Internet Group Management Protocol . . . . .	602
IGMP Versions and Interoperability Requirements . . . . .	602
IGMP Version Transition . . . . .	603
Multicast Listener Discovery . . . . .	604
MLD Versions and Interoperability Requirements . . . . .	604
Source Specific Multicast Groups . . . . .	605
Protocol Independent Multicast Sparse Mode (PIM-SM) . . . . .	606
Ingress Multicast Path Management (IMPM) Enhancements . . . . .	607
Multicast in the BSA . . . . .	608
IGMP Snooping . . . . .	608
IGMP/MLD Message Processing . . . . .	609
IGMP Message Processing . . . . .	609
MLD Message Processing . . . . .	610
IGMP/MLD Filtering . . . . .	611
Multicast VPLS Registration (MVR) . . . . .	612
Layer 3 Multicast Load Balancing . . . . .	613
IGMP State Reporter . . . . .	615
IGMP Data Records . . . . .	616
Transport Mechanism . . . . .	619
HA Compliance . . . . .	619
QoS Awareness . . . . .	619
Hardware Support . . . . .	619
IGMP Reporting Caveats . . . . .	620
Multicast Support over Subscriber Interfaces in Routed CO Model . . . . .	621
Hardware Support . . . . .	623
Multicast Over IPoE . . . . .	624
Per SAP Replication Mode . . . . .	624
Per Subscriber Host Replication Mode . . . . .	633
Multicast Over PPPoE . . . . .	638
IGMP Flooding Containment . . . . .	639
IGMP/MLD Timers . . . . .	639
IGMP/MLD Query Intervals . . . . .	640
HQoS Adjustment . . . . .	640
Host Tracking (HT) Considerations . . . . .	646
HQoS Adjust Per Vport . . . . .	647
Redirection . . . . .	650
Hierarchical Multicast CAC (H-MCAC) . . . . .	652
MCAC Bundle Bandwidth Limit Considerations . . . . .	655
Determining MCAC Policy in Effect . . . . .	660
Multicast Filtering . . . . .	661
Joining the Multicast Tree . . . . .	662
Wholesale/Retail Requirements . . . . .	662

QoS Considerations .....	664
Redundancy Considerations .....	664
Redirection Considerations .....	666
Triple Play Multicast Command Reference .....	669
Command Hierarchies .....	669
Multicast Management Configuration Commands .....	681
Configuring Enhanced Subscriber Management with CLI .....	717
Configuring RADIUS Authentication of DHCP Sessions .....	718
Configuring Enhanced Subscriber Management .....	719
Basic Configurations .....	719
Subscriber Interface Configuration .....	719
Configuring Enhanced Subscriber Management Entities .....	720
Applying the Profiles and Policies .....	726

## **Triple Play Enhanced Subscriber Management**

In This Section .....	729
Uniform RADIUS Server Configuration .....	730
RADIUS Server Configuration .....	730
Uniform RADIUS Server Configuration (Preferred) .....	730
Legacy RADIUS Server Configuration .....	734
RADIUS Authentication of Subscriber Sessions .....	735
RADIUS Authentication Extensions .....	737
Triple Play Network with RADIUS Authentication .....	738
RADIUS Authorization Extensions .....	740
Calling-Station-ID .....	741
Subscriber Session Timeout .....	741
RADIUS Reply Message for PPPoE PAP/CHAP .....	743
radius-server-policy retry Attempt Overview .....	744
Provisioning of Enhanced Subscriber Management (ESM) Objects .....	745
RADIUS-Based Accounting .....	751
Accounting Modes Of Operation .....	755
Per Session Accounting .....	758
RADIUS Per Host Accounting .....	760
No Host-Accounting .....	760
Host-Accounting Enabled .....	760
Accounting Interim Update Message Interval .....	763
Class Attribute .....	763
User Name .....	763
Accounting-On and Accounting Off .....	764
RADIUS Accounting Message Buffering .....	768
Sending an Accounting Stop Message upon a RADIUS Authentication Failure of a PPPoE Session .....	771
Enhanced Subscriber Management Overview .....	773
Enhanced Subscriber Management Basics .....	773
Standard and Enhanced Subscriber Management .....	773
ESM for IPv6 .....	780
Models .....	780
Setup .....	783
Behavior .....	784
Delegated-Prefix-Length .....	785

## Table of Contents

DHCPv6 Relay Agent	788
DHCPv6 Local Server	791
Dynamic Subscriber Host Processing	793
Dynamic Tables	793
Enhanced Subscriber Management Entities	797
Instantiating a New Host	798
Packet Processing for an Existing Host	799
ESM Host Lockout	800
ANCP and GSMP	802
ANCP	802
General Switch Management Protocol Version 3 (GSMPv3)	806
DHCP Release Messages	807
DHCP Client Mobility	807
DHCP Lease Control	807
Using Scripts for Dynamic Recognition of Subscribers	808
Python Language and Programmable Subscriber Configuration Policy (PSCP)	808
Determining the Subscriber Profile and SLA Profile of a Host	809
Determining the Subscriber Profile	810
Determining the SLA Profile	812
Auto-Sub ID	816
Sub-id Identifiers	819
Dual Stack Hosts	819
Mixing Hosts with Auto-Generated IDs and non Auto-Generated IDs	819
PPPoA/PPPoEoA Considerations	820
Deployment Considerations	820
Caveats	822
Limiting Subscribers and Hosts on a SAP	823
Static Subscriber Hosts	823
QoS for Subscribers and Hosts	824
QoS Parameters in Different Profiles	824
QoS Policy Overrides	824
ESM Subscriber Hierarchical Traffic Control	825
Subscriber HQoS	825
Subscriber CFHP	829
ATM/Ethernet Last-Mile Aware QoS for Broadband Network Gateway	831
Subscriber Volume Statistics	853
IP (Layer 3) Volume Accounting	853
Separate IPv4 and IPv6 Counters	854
Configuring IP and IPv6 Filter Policies for Subscriber Hosts	860
ESM PPPoA/PPPoEoA	873
PPPoA	873
PPPoEoA	875
Hardware Support	877
Termination Points within 7x50	878
PPPoA Encapsulation	878
Encapsulation Summary	882
Concurrent Support for Different Service Types on the Same Port	883
Restrictions in Scaled ATM MDA Mode	883
QoS Implementation	884
Association Between the Subscriber and ATM VC Traffic Descriptor (QoS)	887

Per VP Shaping .....	890
ATM/IOM QoS Integration .....	892
Subscriber Instantiation Use Cases .....	902
Authentication .....	906
LUDB Access via Capture SAP .....	907
Encapsulation Autosensing .....	908
SAP Autoprovisioning .....	908
PPP Nodes and ppp-policy .....	909
MTU Considerations .....	910
Multi-Chassis Synchronization .....	913
Overview .....	913
Subscriber Routed Redundancy Protocol (SRRP) .....	916
SRRP Messaging .....	922
SRRP and Multi-Chassis Synchronization .....	923
SRRP Instance .....	924
Subscriber Subnet Owned IP Address Connectivity .....	927
Subscriber Subnet SRRP Gateway IP Address Connectivity .....	927
Receive SRRP Advertisement SAP and Anti-Spoof .....	927
PPPoE MC Redundancy .....	928
Hardware Support .....	928
SRRP Considerations for PPPoE .....	929
State Synchronization .....	930
Traffic Control and Redundant Interface .....	933
MSAP Considerations .....	936
Unnumbered Interface Support .....	937
Compatibility with MC-LAG .....	937
IPv6 Support .....	937
Considerations with Local DHCP Server .....	939
Redundant Interface Considerations .....	940
Routed Central Office (CO) .....	941
Layer 3 Subscriber Interfaces .....	942
Wholesale Retail Routed CO .....	949
Routed Subscriber Hosts .....	955
Dual Homing .....	963
Dual Homing to Two PEs (Redundant-Pair Nodes) in Triple Play Aggregation .....	963
Steady-State Operation of Dual-homed Ring .....	966
Broken-Ring Operation and the Transition to this State .....	968
Transition from Broken to Closed Ring State .....	970
Provisioning Aspects and Error Cases .....	970
Dual Homing to Two BSR Nodes .....	971
MC Services .....	972
Routed CO Dual Homing .....	974
SRRP and Multi-Chassis Synchronization .....	981
Dual Homing and ANCP .....	981
SRRP Enhancement .....	982
SRRP Aware Routing - IPv4/IPv6 Route Advertisement Based on SRRP State .....	991
SRRP in Conjunction with a PW in ESM Environment – Use Case .....	996
Group-monitor .....	996
Subscriber Override .....	1001
Dual Stack Lite .....	1003

## Table of Contents

IP-in-IP .....	1004
Configuring Dual Stack Lite .....	1005
L2TP over IPv6 .....	1006
L2TP Tunnel RADIUS Accounting .....	1007
Accounting Packets List .....	1008
RADIUS Attributes Value Considerations .....	1012
Other Optional RADIUS Attributes .....	1012
RADIUS VSA to Enable L2TP Tunnel Accounting .....	1013
MLPPP on the LNS Side .....	1013
RADIUS Route Download .....	1014
Managed SAP (MSAP) .....	1016
ESM Identification Process .....	1020
Default-Subscriber .....	1020
Multicast Management .....	1021
Subscriber Mirroring .....	1021
Volume and Time Based Accounting .....	1022
Metering .....	1022
Categories Map and Categories .....	1023
Quota Consumption .....	1024
RADIUS VSA Credit-Control-Quota .....	1024
Credit Negotiation Mechanisms .....	1025
Action on Credit Exhaustion .....	1026
Action on Error-Conditions .....	1026
Applicability of Volume and Time Based Accounting .....	1027
Subscriber Host Idle Timeout .....	1028
Web Authentication Protocol (WPP) .....	1030
WPP Configurations .....	1031
WPP Triggered Host Creation .....	1033
LUDB Support For WPP .....	1033
WPP Multi-Chassis Redundancy Support .....	1034
One-time HTTP Redirection Overview .....	1035
ESM over MPLS Pseudowires .....	1036
Encapsulation .....	1038
ESM Configuration with PW-Ports and PW-SAPs .....	1038
QoS Support .....	1041
BNG Redundancy with ESM over Pseudowire .....	1043
EPIPE Based Aggregation Service .....	1043
VPLS Based Aggregation Service .....	1047
Show Commands Related to Active/Standby Pseudowire on Dual BNGs .....	1051
On-Demand Subnet Allocation (ODSA) .....	1054
DHCP pool subnet-binding-key .....	1054
ODSA Subnet Advertisement and Routing .....	1055
ODSA with SRRP .....	1056
ODSA SRRP Failover DHCP Behavior .....	1056
ODSA SRRP Recovery DHCP Behavior .....	1057
Logical Link Identifier (LLID) .....	1058
Open Authentication Model for DHCP and PPPoE Hosts .....	1059
Terminology .....	1059
LUDB and RADIUS Access Models .....	1059
No Authentication .....	1060

LUDB Only Access . . . . .	1060
LUDB Access via DHCPv4 Server . . . . .	1060
RADIUS Only Access . . . . .	1061
Consecutive Access to LUDB and RADIUS . . . . .	1061
RADIUS Fallback . . . . .	1062
Subscriber Services . . . . .	1063
Flexible Subscriber-Interface Addressing (Unnumbered Subscriber-Interfaces) . . . . .	1064
Terminology . . . . .	1064
Flexible Subscriber-Interface Addressing for IPOE/PPPOE v4/v6 Subscribers . . . . .	1064
Default Gateway in IPv4 Flexible Addressing . . . . .	1065
IPv4 Subnet Sharing . . . . .	1067
IPv4 Subnet Mask Auto-Generation . . . . .	1068
Local-proxy-arp and arp-populate . . . . .	1069
Gi-address Configuration Consideration . . . . .	1070
PPPoE Considerations . . . . .	1070
IPoEv6 Considerations . . . . .	1070
General Configuration Guidelines for Flexible IP Address Assignment . . . . .	1070
Caveats . . . . .	1072
uRPF for Subscriber Management . . . . .	1073
IPoE Sessions . . . . .	1074
Enabling IPoE Sessions . . . . .	1075
IPoE Session Authentication . . . . .	1075
IPoE Session Accounting . . . . .	1076
IPoE Session Mid-Session Changes . . . . .	1077
IPoE Session Termination . . . . .	1077
Limiting the Number of IPoE sessions . . . . .	1078
SAP Session Index . . . . .	1078
Resiliency . . . . .	1078
Configuration steps . . . . .	1080
Subscriber Management Command Reference . . . . .	1083
Configuration Commands . . . . .	1083
Subscriber Management Service Commands . . . . .	1116
Triple Play Subscriber Management Configuration Commands . . . . .	1137
Show Commands . . . . .	1359
Clear Commands . . . . .	1415
Tools Commands . . . . .	1420
Debug Commands . . . . .	1426
Monitor Commands . . . . .	1435

## **RADIUS Triggered Dynamic Data Services**

In This Section . . . . .	1441
Introduction to RADIUS Triggered Dynamic Data Services . . . . .	1442
RADIUS Triggered Dynamic Data Services Command Reference . . . . .	1443
Configuration Commands . . . . .	1443
WIFI Aggregation and Offload Command Reference . . . . .	1447
Command Hierarchies . . . . .	1447

## **Python Script Support for ESM**

In This Chapter . . . . .	1465
Python Script Support for ESM . . . . .	1466

## Table of Contents

Python in SR-OS Overview .....	1467
Python Changes .....	1467
Python Support in sub-ident-policy .....	1468
Configuration .....	1470
Operator Debugging .....	1472
Python Scripts .....	1473
Sample Python Scripts .....	1474
Example .....	1474
Example .....	1475
Example .....	1476
Limitations .....	1478
RADIUS Script Policy Overview .....	1479
Python RADIUS API .....	1480
Sample Script .....	1480
Python Policy Overview .....	1481
Python Policy – RADIUS API .....	1482
Python Policy – DHCPv4 API .....	1482
Python Policy – DHCPv6 API .....	1486
Python Policy – Diameter API .....	1493
Python Policy – DHCP Transaction Cache API .....	1500
Python Cache Support .....	1502
Applying a Python Policy .....	1504
Python Script Protection .....	1504
Tips and Tricks .....	1505
<b>Standards and Protocol Support .....</b>	<b>1507</b>



# List of Tables

## Getting Started

Table 1: Configuration Process .....	23
--------------------------------------	----

## Introduction to Triple Play

Table 2: Alcatel-Lucent's TPSDA .....	30
Table 3: Downstream QoS Enablement .....	34
Table 4: Upstream QoS Enablement .....	35
Table 5: Default QinQ and TopQ SAP Dot1P Evaluation .....	144
Table 6: Top Position QinQ and TopQ SAP Dot1P Evaluation .....	144
Table 7: Bottom Position QinQ and TopQ SAP Dot1P Evaluation .....	145

## DHCP Management

Table 8: ESM-Related Python Variables .....	339
Table 9: DHCP Failover Scenarios .....	361

## Triple Play Security

### Triple Play Multicast

Table 10: Data Record Field Description .....	617
---	-----

## Triple Play Enhanced Subscriber Management

Table 11: Subscriber Session Timeout .....	741
Table 12: Accounting Modes of Operation .....	755
Table 13: SRRP State Effect on Subscriber Hosts Associated with Group IP Interface .....	918
Table 14: RADIUS VSA Attributes to Setup Dynamic BGP Peering .....	958
Table 15: L2TP Tunnel Accounting Behavior .....	1008
Table 16: Optional RADIUS Attributes .....	1012
Table 17: Supported RADIUS VSAs .....	1013
Table 18: IPoE Session Authentication Trigger Packets .....	1075

## RADIUS Triggered Dynamic Data Services

### Python Script Support for ESM

Table 19: DHCP Object Members .....	1469
Table 20: alc.dhcpv4 Attributes .....	1482
Table 21: DHCPv6 Header Fields .....	1486
Table 22: Diameter Message Header alc.diameter Attributes .....	1494
Table 23: Message and AVP Manipulation Functionality alc.diameter Methods .....	1494



# List of Figures

## Introduction to Triple Play

Figure 1: Triple Play Service Delivery Architecture	29
Figure 2: Alcatel-Lucent's Triple Play Service Delivery Architecture	31
Figure 3: Downstream QoS Enablement	33
Figure 4: Upstream QoS Enablement	34
Figure 5: Distributed Multicasting in TPSDA	36
Figure 6: Subnetting Topology	37
Figure 7: VMAC Subnetting Topology	38
Figure 8: Alcatel-Lucent's Service Entities	44
Figure 9: Service Access Point (SAP)	45
Figure 10: Multiple SAPs on a Single Port	47
Figure 11: A GRE Service Distribution Point (SDP) pointing from ALA-A to ALA-B	50
Figure 12: Epipe/VLL Service	53
Figure 13: Internet Enhanced Service	56

## DHCP Management

Figure 14: IP Address Assignment with DHCP	314
Figure 15: DHCPv4 Server Routing Instance	317
Figure 16: relay-unicast-msg Command in the DHCPv4 Relay	318
Figure 17: DHCP Lease State Table	324
Figure 18: CMTS/WAC Network Configuration Example	327
Figure 19: Typical DHCP Deployment Scenarios	344
Figure 20: Example of Triple Play Aggregation Network With DHCP to RADIUS Authentication	346
Figure 21: Redundancy Model	354
Figure 22: Potential Expiration Time	357
Figure 23: Address/Prefix Allocation without Synchronization	358
Figure 24: Failover Scenario with SRRP and DHCP in Access-Driven Mode	363

## Triple Play Security

Figure 25: IP Illustration of Message Flow in Web Portal Redirect	555
Figure 26: VPLS Redirect Policy Example	562

## Triple Play Multicast

Figure 27: IGMP Message Processing	609
Figure 28: MVR and MVR by Proxy	612
Figure 29: Common IGMP Data Record Header	616
Figure 30: Data Record Field TLV Structure	617
Figure 31: A Typical Business Connectivity Model	622
Figure 32: 1:1 Model	626
Figure 33: - N:1 Model - AN in IGMP Snooping Mode	628
Figure 34: N:1 Model - AN in Proxy mode	632
Figure 35: 1:1 Model	635
Figure 36: N:1 Model — No IGMP/MLD in the AN	637
Figure 37: Multicast IPv4 Address and Unicast MAC Address in PPPoE Subscriber Multicast	638
Figure 38: HQoS Adjustment per Subscriber and Vport	647

## List of Figures

Figure 39: MCAC Policy Inheritance in Per-SAP Replication Mode .....	658
Figure 40: MCAC Policy Inheritance in Per-HOST Replication Mode .....	659
Figure 41: Wholesale/Retail Multicast Support .....	662

### Triple Play Enhanced Subscriber Management

Figure 42: Triple Play Aggregation Network with RADIUS-Based DHCP Host Authentication .....	738
Figure 43: Purging Message from Buffer .....	769
Figure 44: IPv6 Prefix .....	785
Figure 45: Enhanced Subscriber Management Dynamic Tables .....	793
Figure 46: Relationship Between Enhanced Subscriber Management Entities .....	797
Figure 47: Static ANCP Management Example .....	803
Figure 48: ESM Dynamic ANCP Example .....	804
Figure 49: Data Flow in Determining Subscriber Profile and SLA Profile .....	809
Figure 50: 7450 ESS Determining the Subscriber Profile .....	810
Figure 51: 7450 ESS Determining the SLA Profile .....	813
Figure 52: Ingress Scheduling Hierarchy Options .....	826
Figure 53: Egress Scheduling Hierarchy Options .....	827
Figure 54: Ingress Policing Hierarchy Options .....	829
Figure 55: Egress Policing Hierarchy Options .....	830
Figure 56: BNG Application .....	831
Figure 57: BNG Queuing and Scheduling Model .....	833
Figure 58: Subscriber Host Session Encapsulation Types .....	846
Figure 59: Access-Loop-Encapsulation Sub-TLV .....	847
Figure 60: Insert Shared Filters .....	864
Figure 61: PPPoA Architecture and Packet Encapsulation .....	874
Figure 62: PPOEoA Host Terminated Session .....	875
Figure 63: PPPoEoA DSL CPE Terminated Session .....	876
Figure 64: PPPoA LLC Encapsulation .....	879
Figure 65: PPPoA AAL5MUX Encapsulation .....	879
Figure 66: PPPoEoA Bridged LLC/SNAP Encapsulation .....	880
Figure 67: PPPoA AAL5MUX Encapsulation .....	881
Figure 68: LLC/SNAP Encapsulation .....	882
Figure 69: Scheduling on ATM MDA .....	885
Figure 70: ATM Traffic Descriptor Association with Subscriber .....	887
Figure 71: VP Shaper .....	890
Figure 72: Tier HQoS .....	894
Figure 73: VPI Based V-Port <-> Subscriber Association .....	896
Figure 74: QoS Adjustment .....	898
Figure 75: ATM Wire Overhead .....	901
Figure 76: Subhost per VC .....	902
Figure 77: Multiple VCs per Subscriber .....	903
Figure 78: Multiple Hosts per Subscriber, Single VC .....	904
Figure 79: VP Shaping .....	905
Figure 80: Dual-Homing Configuration .....	913
Figure 81: Fully Redundant "Statefull 1:1" Model .....	932
Figure 82: Shared Subscriber IP Space .....	934
Figure 83: Option 'B' – IP Subnet per Active SRRP Group .....	935
Figure 84: DSLAM Connection .....	941
Figure 85: Subscriber Interface in an IES/VRN Service .....	942

Figure 86: Details of a Group Interface . . . . .	943
Figure 87: Aggregation Network with Direct DSLAM-BSE Connection . . . . .	943
Figure 88: Detailed View of Configurable Objects Related to Layer 3 Subscriber Interfaces . . . . .	944
Figure 89: Wholesale Retail Model . . . . .	949
Figure 90: Wholesale Retail – Hub and Spoke Forwarding . . . . .	954
Figure 91: Router Subscriber Hosts. . . . .	955
Figure 92: Dual-Homing to Two PEs . . . . .	963
Figure 93: Layer 2 CO Dual Homing - Network Diagram. . . . .	965
Figure 94: Dual Homing Ring Under Steady-State Condition . . . . .	966
Figure 95: Broken Ring State. . . . .	968
Figure 96: Low . . . . .	971
Figure 97: MC Services in a Layer 3-Ring Topology (a) . . . . .	972
Figure 98: MC Services on a Layer 3-Ring Topology (b). . . . .	973
Figure 99: Dual Homing Example . . . . .	977
Figure 100: IP Subnet Per SRRP Master Group . . . . .	983
Figure 101: FSG — Single Network Failure . . . . .	984
Figure 102: Multiple Network Failures . . . . .	984
Figure 103: SRRP Fate Sharing . . . . .	986
Figure 104: Scenario 1. . . . .	987
Figure 105: Scenario 2. . . . .	987
Figure 106: Scenario 3. . . . .	988
Figure 107: Scenario 4. . . . .	988
Figure 108: Pseudowire Example . . . . .	997
Figure 109: Dual-Stack Lite . . . . .	1003
Figure 110: IP-in-IP . . . . .	1004
Figure 111: L2TP over IPv6 . . . . .	1006
Figure 112: L2TP Tunnel Accounting . . . . .	1007
Figure 113: Threshold Configured/Not Configured . . . . .	1025
Figure 114: WPP Authentication . . . . .	1030
Figure 115: ESM over MPLS Pseudowire Example . . . . .	1036
Figure 116: Group Interface Example . . . . .	1037
Figure 117: Subscriber Frame with PWE Encapsulation . . . . .	1038
Figure 118: BNG Redundancy Based on Active/Standby PW Signaling . . . . .	1043
Figure 119: BNG Redundancy with VPLS Based Aggregation Service. . . . .	1047
Figure 120: Subscriber Interfaces Sharing a DHCP Pool . . . . .	1055
Figure 121: Use Case for Flexible IP Addressing Model . . . . .	1064
Figure 122: IPoE Session . . . . .	1074

## List of Figures

# Preface

---

## About This Guide

This guide describes details pertaining to Triple Play Services Delivery Architecture (TPSDA) support provided by the SR OS and presents examples to configure and implement various protocols and services.

This guide is organized into functional chapters and provides concepts and descriptions of the implementation flow, as well as Command Line Interface (CLI) syntax and command usage.

---

## Audience

This guide is intended for network administrators who are responsible for configuring the 7450 ESS routers. It is assumed that the network administrators have an understanding of networking principles and configurations. Concepts described in this guide include the following:

- Triple Play concepts
  - Triple Play Security
  - Enhanced subscriber management
  - Multicast
  - Operation, Administration and Maintenance (OAM) operations
- 

## List of Technical Publications

The 7450 ESS documentation set is composed of the following guides:

- 7450 ESS Basic System Configuration Guide  
This guide describes basic system configurations and operations.
- 7450 ESS System Management Guide  
This guide describes system security and access configurations as well as event logging and accounting logs.
- 7450 ESS Interface Configuration Guide

This guide describes card, Media Dependent Adapter (MDA) and port provisioning.

- 7450 ESS Router Configuration Guide

This guide describes logical IP routing interfaces and associated attributes such as an IP address, as well as IP and MAC-based filtering, and VRRP and Cflowd.

- 7450 ESS Routing Protocols Guide

This guide provides an overview of routing concepts and provides configuration examples for RIP, OSPF, IS-IS, BGP, and route policies.

- 7450 ESS MPLS Guide

This guide describes how to configure Multiprotocol Label Switching (MPLS) and Label Distribution Protocol (LDP).

- 7450 ESS Services Overview Guide

This guide describes how to configure service parameters such as service distribution points (SDPs), customer information, and user services.

- 7450 ESS Layer 2 Services and EVPN Guide

This guide describes Virtual Leased Lines (VLL), Virtual Private LAN Service (VPLS), Provider Backbone Bridging (PBB), and Ethernet VPN (EVPN).

- 7450 ESS Layer 3 Services Guide

This guide describes Internet Enhanced Services (IES) and Virtual Private Routed Network (VPRN) services.

- 7450 ESS Versatile Service Module Guide

This guide describes how to configure service parameters for the Versatile Service Module (VSM).

- 7450 ESS OAM and Diagnostics Guide

This guide describes how to configure features such as service mirroring and Operations, Administration and Management (OAM) tools.

- 7450 ESS Triple Play Guide

This guide describes Triple Play services and support provided by the 7450 ESS and presents examples to configure and implement various protocols and services.

- 7450 ESS Quality of Service Guide

This guide describes how to configure Quality of Service (QoS) policy management.

- 7450 ESS Gx AVPs Reference Guide

This guide describes Gx Attribute Value Pairs (AVP).

- OS Multi-Service ISA Guide

This guide describes services provided by integrated service adapters such as Application Assurance, ad insertion (ADI) and Network Address Translation (NAT).



## Technical Support

If you purchased a service agreement for your 7450 ESS router and related products from a distributor or authorized reseller, contact the technical support staff for that distributor or reseller for assistance. If you purchased an Alcatel-Lucent service agreement, follow this link to contact an Alcatel-Lucent support representative and to access product manuals and documentation updates:

<http://support.alcatel-lucent.com>



# Getting Started

---

## In This Chapter

This book provides process flow information to configure provision protocols and services pertaining to Triple Play Services Delivery Architecture (TPSDA).

---

## Alcatel-Lucent 7450 ESS-Series Services Configuration Process

[Table 1](#) lists the tasks necessary to configure TPSDA entities. Each section describes a software area and provides CLI syntax and command usage to configure parameters for a functional area.

**Table 1: Configuration Process**

Area	Task	Chapter
Introduction	Overview	<a href="#">Introduction to Triple Play on page 25</a>
DHCP	Theory, configuration	<a href="#">DHCP Management on page 313</a>
Security	Anti-spoofing, MAC pinning, MAC protection, DoS protection VPLS redirect policies, Web portal redirect, ARP handling	<a href="#">Triple Play Security on page 545</a>
PPPoE	Theory, configuration	<a href="#">Point-to-Point Protocol over Ethernet (PPPoE) Management on page 561</a>
Multicast	Multicast, IGMP, SSM, PIM	<a href="#">Triple Play Multicast on page 599</a>
Enhanced Subscriber management	RADIUS authentication, ESM entities, Routed CO, M-SAP	<a href="#">Triple Play Enhanced Subscriber Management on page 729</a>

**Table 1: Configuration Process**

Area	Task	Chapter (Continued)
	Python Scripting	<a href="#">Python Script Support for ESM on page 1465</a>
Reference	List of IEEE, IETF, and other proprietary entities.	<a href="#">Standards and Protocol Support on page 1507</a>

**Note:** In SR OS 12.0.R4 any function that displays an IPv6 address or prefix changes to reflect rules described in RFC 5952, *A Recommendation for IPv6 Address Text Representation*. Specifically, hexadecimal letters in IPv6 addresses are now represented in lowercase, and the correct compression of all leading zeros is displayed. This changes visible display output compared to previous SR OS releases. Previous SR OS behavior can cause issues with operator scripts that use standard IPv6 address expressions and with libraries that have standard IPv6 parsing as per RFC 5952 rules. See the section on IPv6 Addresses in the Router Configuration Guide for more information.

# Introduction to Triple Play

---

## In This Section

This section provides an overview of the ESS-Series services, service model and service entities used in conjunction with Triple Play services only to the relevant service types. Details about services, configurations, and CLI syntax can be found in the 7450 ESS Services Overview Guide. Topics in this section include:

- [Alcatel-Lucent's Triple Play Service Delivery Architecture on page 26](#)
  - [Introduction to Triple Play on page 26](#)
  - [Blueprint for Optimizing Triple Play Service Infrastructures on page 27](#)
  - [Architectural Foundations on page 28](#)
  - [Optimizing Triple Play Service Infrastructures on page 30](#)
- [Services on page 40](#)
  - [Service Types on page 41](#)
  - [Service Policies on page 42](#)
- [Alcatel-Lucent Service Model on page 43](#)
  - [Service Entities on page 44](#)
  - [Customers on page 44](#)
  - [Service Access Points \(SAPs\) on page 45](#)
  - [Service Distribution Points \(SDPs\) on page 49](#)
- [Epipe Service Overview on page 53](#)
- [VPLS Service Overview on page 54](#)
  - [Split Horizon SAP Groups and Split Horizon Spoke SDP Groups on page 54](#)
- [IES Service Overview on page 56](#)
  - [IP Interface on page 57](#)

# Alcatel-Lucent's Triple Play Service Delivery Architecture

---

## Introduction to Triple Play

For more than a decade, telephony service providers have considered offering video services to residential customers. However, in the past it was not economically nor technically feasible to launch the implementation on a large scale.

Recently, several technical trends and evolutions have propelled video delivery to the foreground, including:

- Technical improvements in areas such as real-time MPEG encoding and compression.
- Widespread deployment of High Speed Internet (HSI) over broadband access (ADSL and cable modems).
- Decreased cost of high-bandwidth infrastructure (typically Ethernet-based) as well as storing, converting, and delivering video content.
- Increased competition between telephony and cable operators. This is partly due to changes in regulations.

Traditional cable operators began offering television services and later added Internet access and telephony to their offerings. Conversely, traditional telephony operators such as RBOCS, PTTs, have also added Internet access, and many are now in the process of also adding video delivery.

This bundling of video, voice, and data services to residential subscribers is now commonly known as Triple Play services. The video component always includes linear programming (broadcast television), but often also has a non-linear Video on Demand (VoD) component.

## Blueprint for Optimizing Triple Play Service Infrastructures

Alcatel-Lucent's TPSDA allows network operators to progressively integrate their HSI, voice, and video services within a unified and homogeneous Ethernet-based aggregation network environment. The key benefits of the proposed service infrastructure include cost optimization, reduced risk, and accelerated time to market for new services.

At a high level, TPSDA implements:

- Ethernet-based service architecture — Solves bandwidth bottlenecks and exponential capital expenditure and operating expenses issues in the second mile by leveraging the efficiency of this technology.
- Multiple distributed service edges — Allows service providers to achieve faster times to market for new services while retaining the existing Broadband Remote Access Server (BRAS) / Point-to-Point Protocol over Ethernet (PPPoE) mode of operation for wholesale and retail HSI.
- Distributed multicasting functions in access and aggregation networks — Enables service providers to optimize bandwidth and content delivery mechanisms, based on densities and penetration rates. It is also essential to subscriber and service scaling, and optimizes the bandwidth required in the aggregation network.
- Carrier video and Voice over Internet Protocol (VoIP) services using Dynamic Host Configuration Protocol (DHCP) — Enables service providers to introduce plug-and-play services delivered through set-top boxes and VoIP devices, which are designed for use with the DHCP.
- Flexible deployment models — The architecture allows data, video, and VoIP services to be rapidly rolled out without any lock-in to specific operational models. It allows service providers to maximize flexibility and minimize financial and technological risks by allowing all modes of operation, including:
  - Copper (DSL/DSLAM) and fiber-based (FTTx) deployments in the first mile.
  - Single or multiple last mile circuits.
  - Bridged or routed home gateways.
  - Single or multiple IP address deployment models.

## Architectural Foundations

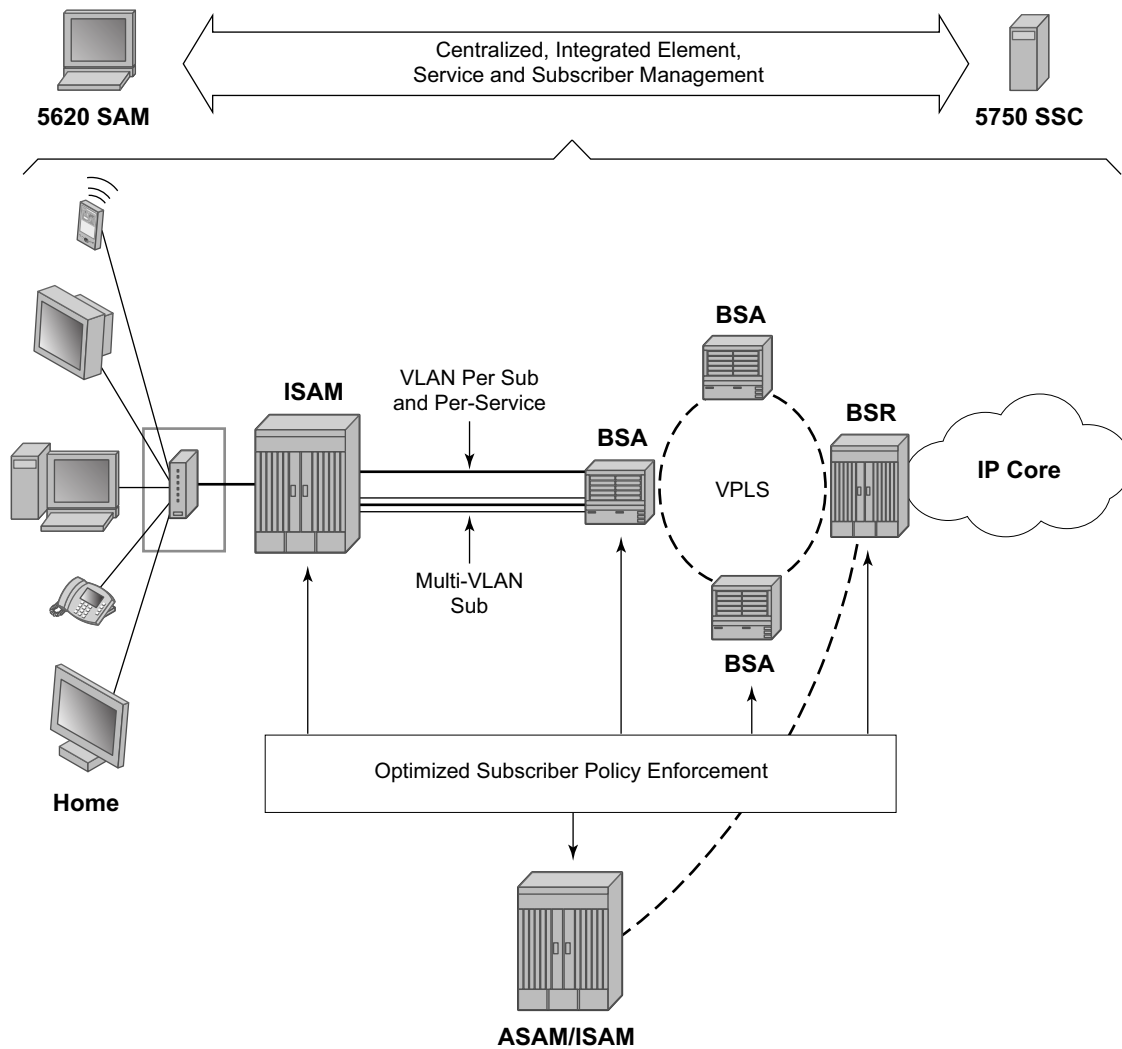
With SR OS, the architectural foundations of Alcatel-Lucent's TPSDA established in previous releases is reinforced while its applicability is expanded to encompass many new deployment models and support Any Mode of Operation (AMO). Through these enhancements, TPSDA becomes more universally deployable and flexible in addressing the specifics of any provider's network/rollout.

Alcatel-Lucent has defined new terminologies that have been adopted industry-wide, including:

- Broadband Service Access Node (BSAN)
- Broadband Service Aggregator (BSA)
- Broadband Service Router (BSR)

[Figure 1](#) depicts TPSDA's centralized, integrated element, service and subscriber management architecture.





OSSG091

**Figure 1: Triple Play Service Delivery Architecture**

# Optimizing Triple Play Service Infrastructures

More than a branding exercise, new terminologies signal a significant shift from “best effort” and traditional DSLAMs, Ethernet switches and BRASs, in the sense that they capture a shift in required characteristics and capabilities for a new generation of service rollouts, including:

- High-availability for non-stop service delivery (non-stop unicast and multicast routing, non-stop services, etc.).
- Multi-dimensional scale (such as the ability to scale performance, bandwidth, services, and subscribers concurrently).
- Ethernet — Optimization (leading density, capacity, scaling, performance).
- Optimal system characteristics (optimal delay/jitter/loss characteristics, etc.).
- Rich service capabilities with uncompromised performance.

Alcatel-Lucent’s Triple Play Service Delivery Architecture (TPSDA) advocates the optimal distribution of service intelligence over the BSAN, BSA and BSR, rather than concentrating on fully centralized or decentralized BRAS models which artificially define arbitrary policy enforcement points in the network. With SR OS, the optimized enforcement of subscriber policies across nodes or over a single node (as dictated by evolving traffic patterns), allows a more flexible, optimized, and cost-effective deployment of services in a network, guaranteeing high quality and reliable delivery of all services to the user.

**Table 2: Alcatel-Lucent’s TPSDA**

Entity	Description
Subscriber Management	Centralized and fully integrated with element and services management across the infrastructure end-to-end solution (through the Alcatel-Lucent 5750 SSC).
Policy Enforcement	Optimally distributed, based on actual traffic patterns. Maximized flexibility, minimized risk of architectural lock-in. Optimized cost structure.
Support for “Any Mode of Operation”	With TPSDA, network economics, subscriber density, network topologies and subscriber viewership patterns define the optimal policy enforcement point for each policy type (security, QoS, multicasting, anti-spoofing, filtering etc.). The SR OS capabilities allow service providers to support any mode of operation, including any combination of access methods, home gateway type, and policy enforcement point (BSAN, BSA or BSR or a combination of the three).

All of the SR OS and Alcatel-Lucent’s 5750 SSC’s subscriber policy enforcement and management capabilities described in this section build upon Alcatel-Lucent’s TPSDA extensive capabilities and provide key capabilities in the following areas:

- Operationalization of Triple Play Services (AAA, subscriber policy enforcement, etc.)
- Service Assurance and Control
- Non-stop Video Service Delivery

## Distributed Service Edges

The TPSDA architecture (Figure 2), is based on two major network elements optimized for their respective roles, the Broadband Service Aggregator (BSA) and the Broadband Service Router (BSR). An important characteristic of BSAs and BSRs is that they effectively form a distributed virtual node with the BSAs performing subscriber-specific functions where the various functions scale, and the BSRs providing the routing intelligence where it is most cost-effective.

The Alcatel-Lucent 7450 ESS and 7750 SR Series, respectively, provide the BSA and BSR functionalities in TPSDA. Both are managed as a single virtual node using Alcatel-Lucent's 5620 Service Aware manager (SAM), which provides a unified interface for streamlined service and policy activation across the distributed elements of the TPSDA architecture, including VPLS, QoS, multicasting, security, filtering and accounting.

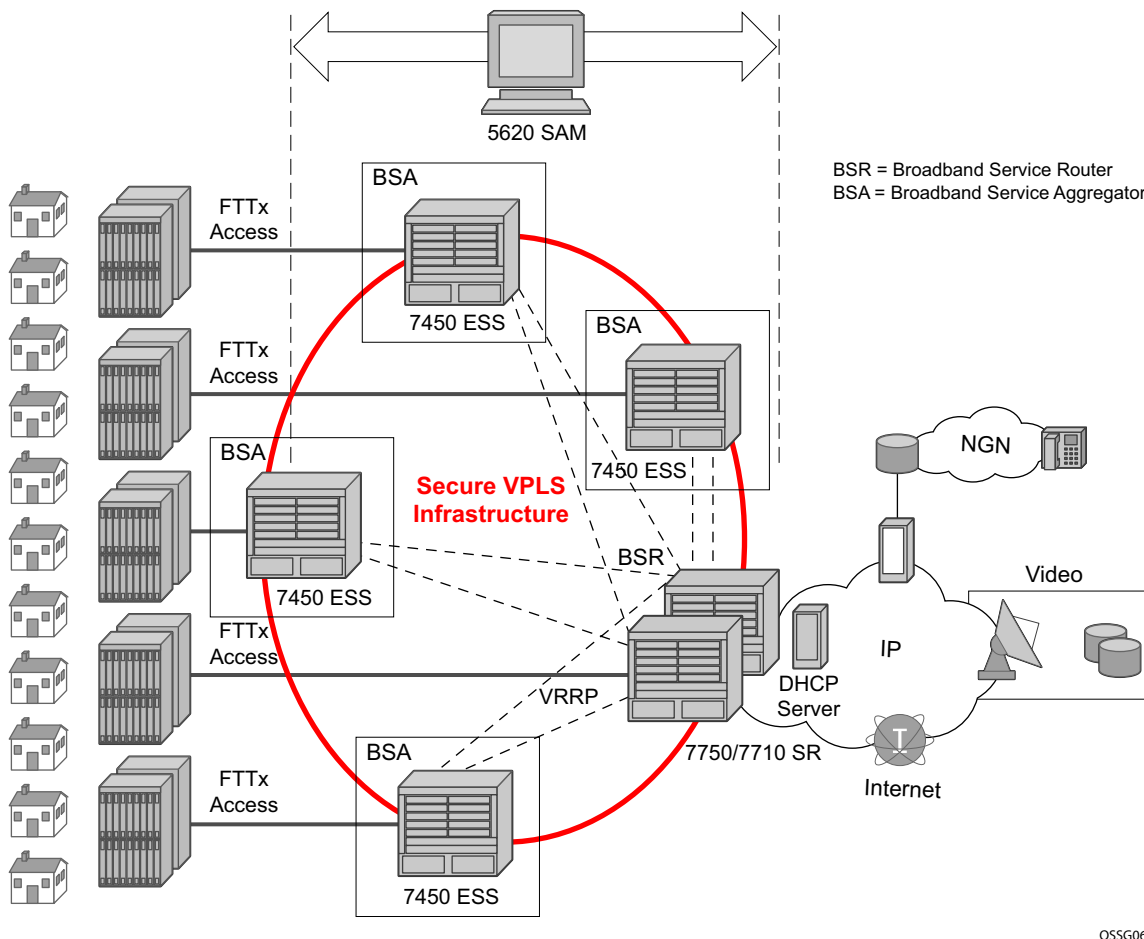


Figure 2: Alcatel-Lucent's Triple Play Service Delivery Architecture

Digital subscriber line access multiplexers (DSLAMs) or other access nodes are connected to Ethernet access ports on the BSA. Typically a single VLAN per subscriber is configured between the access node and the BSA. A VLAN per subscriber provides a persistent context against which per-subscriber policies (QoS, filtering, accounting) can be applied in the BSA.

Scaling of traffic and services is achieved by dividing the Layer 2 and Layer 3 functions between the BSA and BSR and by distributing key service delivery functions. BSAs are more distributed than BSRs, cost-effectively scaling per-subscriber policy enforcement.

The BSA is a high-capacity Ethernet-centric aggregation device that supports hundreds of Gigabit Ethernet ports, tens of thousands of filter policies, and tens of thousands of queues. The BSA incorporates wire speed security, per-subscriber service queuing, scheduling, accounting, and filtering.

BSAs aggregate traffic for all services towards the BSR. The BSR terminates the Layer 2 access and routes over IP/MPLS (Multi Protocol Label Switching) with support for a full set of MPLS and IP routing protocols, including multicast routing. The BSR supports hundreds of ports and sophisticated QoS for per-service and per-content/source differentiation.

The connectivity between BSAs and BSRs is a Layer 2 forwarding model shown in [Figure 2](#) above as a secure VPLS infrastructure. This refers to the fact that the BSA-BSR interconnections form a multipoint Ethernet network with security extensions to prevent unauthorized communication, denial of service, and theft of service. One of the advantages of using VPLS for this application is that VPLS instances can be automatically established over both ‘hub and spoke’ and ring topologies providing sub-50 ms resilience. Regardless of the fiber plant layout, VPLS enables a full mesh to be created between BSA and BSR nodes, ensuring efficient traffic distribution and resilience to node or fiber failure.

Other unique features of the BSA and BSR that contribute to this secure VPLS infrastructure are:

1. Using Residential Split Horizon Groups (RSHG), direct user-user bridging is automatically prohibited, without the need for address-specific ACLs;
2. RSHG combined with the ARP reply agent perform ARP and broadcast suppression to ensure that addressing information is restricted;
3. Protection against theft of service and denial of service is provided by MAC and/or IP filters automatically populated using DHCP snooping, and by MAC pinning;
4. Using the RADIUS interface, is possible to perform RADIUS authentication of users before allowing a DHCP discover to progress into the network.

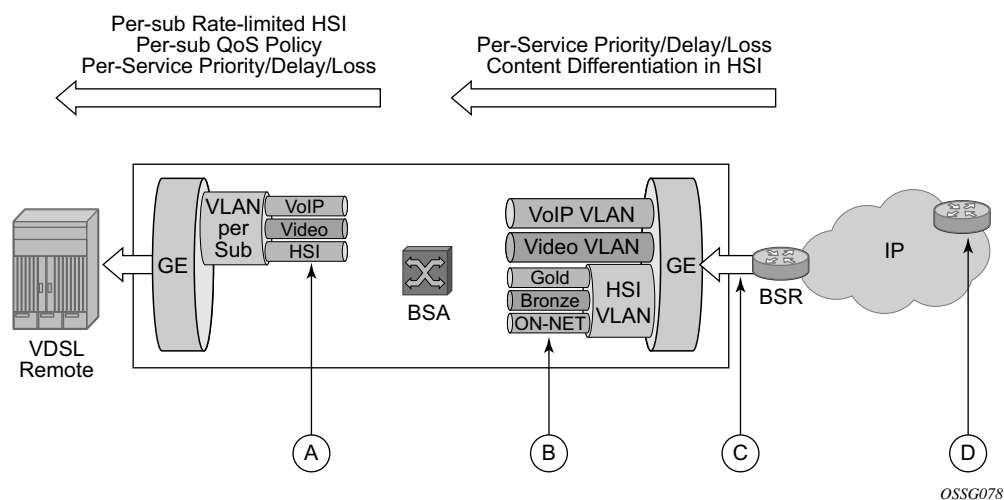
## Service Differentiation, QoS Enablement

Alcatel-Lucent's TPSDA approach provides a model based on call admission for video and VoIP, with the need to guarantee delay/jitter/loss characteristics once the service connection is accepted. The architecture also meets the different QoS needs of HSI, namely per-subscriber bandwidth controls, including shaping and policing functions that have little or no value for video and VoIP services. In conjunction with the architecture's support for content differentiation, this enables differentiated service pricing within HSI.

The distribution of QoS policy and enforcement across BSA and BSR allows the service provider to implement meaningful per-subscriber service level controls. Sophisticated and granular QoS in the BSA allows the service provider to deliver truly differentiated IP services based on the subscriber as well as on the content.

In the BSR to BSA downstream direction ([Figure 3](#)), IP services rely on IP layer classification of traffic from the network to queue traffic appropriately towards the BSA. Under extreme loading (only expected to occur under network fault conditions), lower priority data services and/or HSI traffic will be compromised in order to protect video and voice traffic. Classification of HSI traffic based on source network address or IEEE 802.1p marking allows the QoS information to be propagated to upstream or downstream nodes by network elements. Refer to [Table 3](#) for the descriptions.

The BSR performs service distribution routing based on guarantees required to deliver the service and associated content, rather than on individual subscribers. The BSR only needs to classify content based on the required forwarding class for a given BSA to ensure that each service's traffic receives the appropriate treatment towards the BSA.



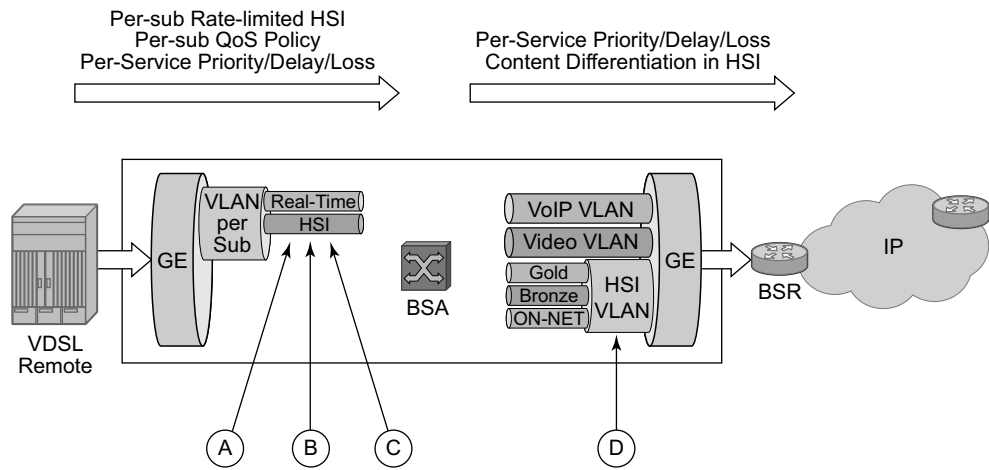
### Figure 3: Downstream QoS Enablement

**Table 3: Downstream QoS Enablement**

Key	Description
A	Per-subscriber queueing and PIR/CIR policing/shaping for HSI. HSI service classified on source IP range. Per-service prioritization for VoIP and video. VoIP is prioritized over video. Destination IP and/or DSCP classification. 802.1 marking for prioritization in the access and home.
B	VoIP and video queued and prioritized on per-VLAN QoS policy basis. HSI content differentiation based on DSCP. Each queue may have individual CIR/PIR and shaping. Optical overall subscriber rate limiting on VLAN (H-QoS).
C	For HSI, content differentiation queueing for gold/silver/bronze based on DSCP classification. Optional overall subscriber rate limiting on VLAN.
D	Preferred content marked (DSCP) of trusted ingress points of IP network.

In the upstream direction (BSA to BSR, [Figure 4](#)), traffic levels are substantially lower. Class-based queuing is used on the BSA network interface to ensure that video control traffic is propagated with a minimal and consistent delay, and that preferred data/HSI services receive better treatment for upstream/peering service traffic than the best effort Internet class of service.

Note that the IP edge is no longer burdened with enforcing per-subscriber policies for hundreds of thousands of users. This function is now distributed to the BSAs, and the per-subscriber policies can be implemented on the interfaces directly facing the access nodes.



OSSG079

**Figure 4: Upstream QoS Enablement**

**Table 4: Upstream QoS Enablement**

Key	Description
A	HSI: Per-subscriber queueing with PIR/CIR policy/shaping.
B	VoIP/Video: Shared queueing for prioritization of real-time traffic over HSI. Upstream video is negligible.
C	Per-subscriber QoS/Content classification for content differentiation.
D	Video/VoIP: Policy defines priority aggregate CIR/PIR. HSI: QoS policy defines priority and aggregate CIR/PIR. Content differentiation based on ingress classification. DSCP is marked.

The BSA is capable of scheduling and queuing functions on a per-service, per-subscriber basis, in addition to performing wire-speed packet classification and filtering based on both Layer 2 and Layer 3 fields.

Each subscriber interface provides at least three dedicated queues. TPSDA makes it possible to configure these queues such that the forwarding classes defined for all services can all be mapped to one service VLAN upstream. In the BSA, assuming hundreds of subscribers per Gigabit Ethernet interface, this translates to a thousand or more queues per port.

In addition to per-service rate limiting for HSI services, each subscriber's service traffic can be rate limited as an aggregate using a bundled service policy. This allows different subscribers to receive different service levels independently and simultaneously.

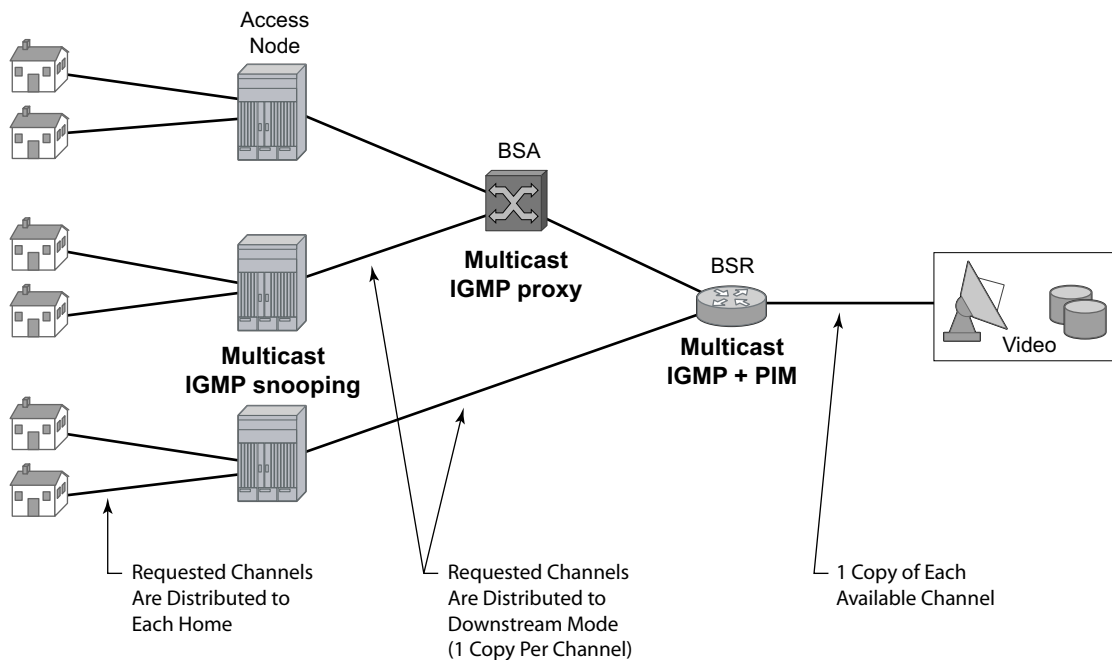
Distributed multicasting today's predominant video service is broadcast TV, and will likely remain significant for a long time. As video services are introduced, it is sensible to optimize investments by matching resources to the service model relevant at the time. Consequently, the objective of the service infrastructure should be to incorporate sufficient flexibility to optimize for broadcast TV in the short term, yet scale to support a full unicast (VoD) model as video service offerings evolve.

Optimizing for broadcast TV means implementing multicast packet replication throughout the network. Multicast improves the efficiency of the network by reducing the bandwidth and fiber needed to deliver broadcast channels to the subscriber. A multicasting node can receive a single copy of a broadcast channel and replicate it to any downstream nodes that require it, substantially reducing the required network resources. This efficiency becomes increasingly important closer to the subscriber. Multicast should be performed at each or either of the access, aggregation, and video edge nodes.

Multicasting as close as possible to the subscriber has other benefits since it enables a large number of users to view the content concurrently. The challenges of video services are often encountered in the boundary cases, such as live sports events and breaking news, for which virtually all the subscribers may be watching just a few channels. These exceptional cases generally involve live content, which is true broadcast content. Multicasting throughout the

network makes it possible to deliver content under these circumstances while simplifying the engineering of the network.

Efficient multicasting requires the distribution of functions throughout the access and the aggregation network to avoid overloading the network capacity with unnecessary traffic. TPSDA realizes efficient multicasting by implementing IGMP snooping in the access nodes, IGMP snooping in the BSA and multicast routing in the BSR (Figure 5).



OSSG079

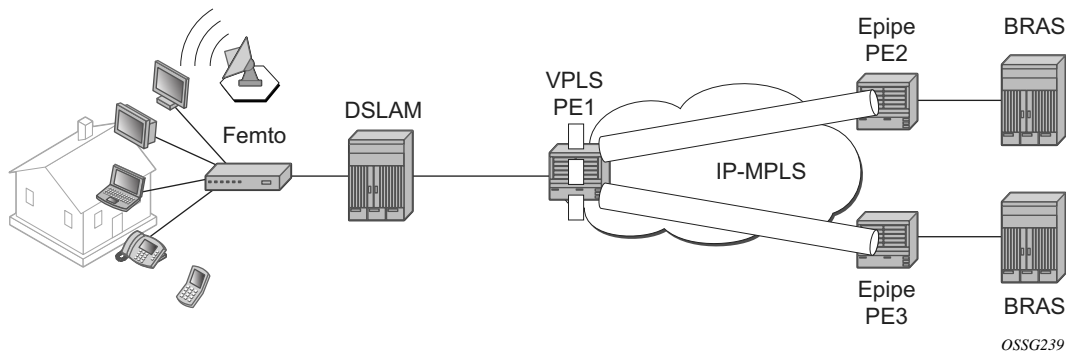
**Figure 5: Distributed Multicasting in TPSDA**



## Virtual MAC Subnetting for VPLS

This feature allows, at the VPLS instance level, MAC subnetting, such as learning and switching based on a configurable number of bits from the source MAC address and from the destination MAC, respectively. This considerably reduces the VPLS FIB size.

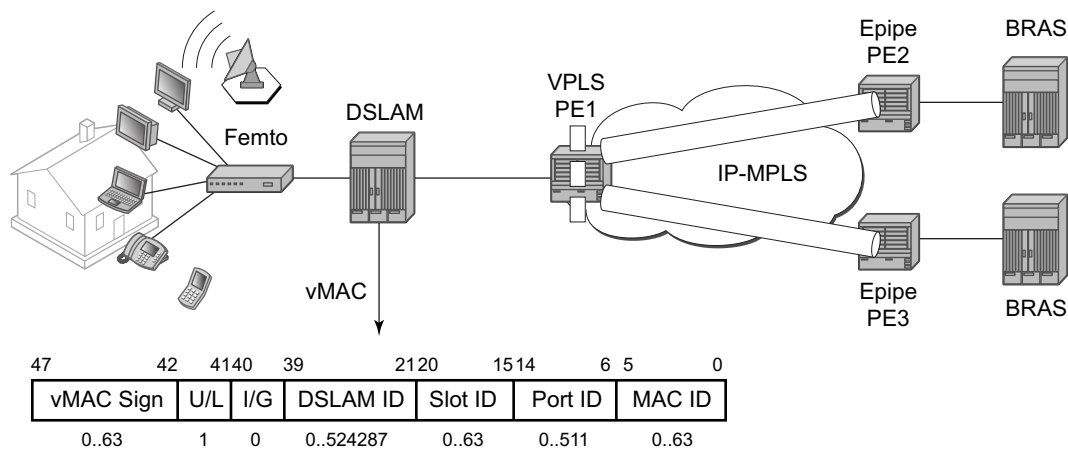
MAC scalability involving MAC learning and switching based on the first x bits of a virtual MAC address is suitable in an environment where some MAC addresses can be aggregated based on a common first x bits, for example 28 out of 48. This can be deployed in a TPSDA environment where the VPLS is used for pure aggregation (there is no subscriber management) between the DSLAM and BRAS devices. The DSLAMs must be able to map customer MAC addresses to a pool of internal virtual MAC addresses where the first bits (28, for example) identify the DSLAM with the next 20 bits identifying, the DSLAM slot, port number, and customer MAC station on that port. The VPLS instance(s) in the PE distinguishes only between different DSLAMs connected to it. They need to learn and switch based only on the first 28 bits of the MAC address allowing scaling of the FIB size in the PE.



**Figure 6: Subnetting Topology**

Figure 6 displays a Layer 2 PE network (such as the ESS-Series) aggregating traffic from DSLAMs (Alcatel-Lucent) to BRAS devices. The VPLS service is running in the PEs directly connected to the DSLAMs (VPLS PE1) while the PEs connected to the BRAS devices are running a point-to-point Layer 2 service (Epipe).

Alcatel-Lucent DSLAMs have the capability to map every customer MAC to a service provider MAC using the virtual MAC addressing scheme depicted in Figure 7.



OSSG240

**Figure 7: VMAC Subnetting Topology**

As the packet ingresses the DSLAM from the residential customer, the source MAC address (a customer MAC for one of its terminals/routers) is replaced by the DSLAM with a virtual MAC using the format depicted in [Figure 7](#).

- The U/L bit is the seventh bit of the first byte and is used to determine whether the address is universally or locally administered. The U/L bit is set to 1 to indicate the address is locally administered.
- The following bits are used to build the VMAC address: DSLAM ID bits 39 — 21, slot ID bits 20 — 15, port ID bits 14-6 and the customer station ID bits 5 — 0.

Based on this scheme, it is apparent that the VMACs from one DSLAM have bits 47-21 in common.

The VPLS instance in PE1 only learns the first part of the MAC (bits 47 — 21) and, as the packets arrive from the BRAS device, switches based only on these bits in the destination MAC address to differentiate between the connected DSLAMs. Once the packet arrives at the DSLAM, the entire destination MAC is checked to determine the slot, port and which specific customer station the packet is destined to. As the packet is sent to the customer, the DSLAM replaces the destination MAC address with the actual customer MAC corresponding to the customer station.

The following are VPLS features not supported when the VMAC subnetting feature is enabled:

- Blocked features — CLI consistency checked provided
- Residential Split Horizon Groups
- BGP AD
- TPSDA (subscriber management) features
- PBB
- VPLS OAM (MAC populate, MAC ping, MAC trace, CPE Ping)

## Services

A service is a globally unique entity that refers to a type of connectivity service for either Internet or VPN connectivity. Each service is uniquely identified by a service ID within a service area. The ESS-Series service model uses logical service entities to construct a service. In the service model, logical service entities are provide a uniform, service-centric configuration, management, and billing model for service provisioning.

Services can provide Layer 2/bridged service or Layer3/IP routed connectivity between a service access point (SAP) on one ESS-Series router and another service access point (a SAP is where traffic enters and exits the service) on the same (local) or another ESS-Series router (distributed). A distributed service spans more than one router.

Distributed services use service distribution points (SDPs) to direct traffic to another ESS-Series through a service tunnel. SDPs are created on each participating ESS-Series, specifying the origination address (the ESS-Series router participating in the service communication) and the destination address of another ESS-Series. SDPs are then bound to a specific customer service. Without the binding process, far-end ESS-Series devices are not able to participate in the service (there is no service without associating an SDP with a service).

## Service Types

The ESS-Series offers the following types of subscriber services which are described in more detail in the referenced chapters:

- Virtual Leased Line (VLL) services:
  - Ethernet pipe (Epipe) — A Layer 2 point-to-point VLL service for Ethernet frames.
  - IP Pipe (Ipipe) — Provides IP connectivity between a host attached to a point-to-point access circuit (FR, ATM, PPP) with routed IPv4 encapsulation and a host attached to an Ethernet interface.
- Virtual Private LAN Service (VPLS) — A Layer 2 multipoint-to-multipoint VPN. VPLS includes Hierarchical VPLS (H-VPLS) which is an enhancement of VPLS which extends Martini-style signaled or static virtual circuit labeling outside the fully meshed VPLS core.
- Internet Enhanced Service (IES) — A direct Internet access service where the customer is assigned an IP interface for Internet connectivity.

## Service Policies

Common to all ESS-Series connectivity services are policies that are assigned to the service. Policies are defined at a global level and then applied to a service on the router. Policies are used to define ESS-Series service enhancements. The types of policies that are common to all ESS-Series connectivity services are:

- SAP Quality of Service (QoS) policies which allow for different classes of traffic within a service at SAP ingress and SAP egress.

QoS ingress and egress policies determine the QoS characteristics for a SAP. A QoS policy applied to a SAP specifies the number of queues, queue characteristics (such as forwarding class, committed, and peak information rates, etc.) and the mapping of traffic to a forwarding class. A QoS policy must be created before it can be applied to a SAP. A single ingress and a single egress QoS policy can be associated with a SAP.

- Filter policies allow selective blocking of traffic matching criteria from ingressing or egressing a SAP.

Filter policies, also referred to as access control lists (ACLs), control the traffic allowed in or out of a SAP based on MAC or IP match criteria. Associating a filter policy on a SAP is optional. Filter policies are identified by a unique filter policy ID. A filter policy must be created before it can be applied to a SAP. A single ingress and single egress filter policy can be associated with a SAP.

- Scheduler policies define the hierarchy and operating parameters for virtual schedulers. Schedulers are divided into groups based on the tier each scheduler is created under. A tier is used to give structure to the schedulers within a policy and define rules for parent scheduler associations.
- Accounting policies define how to count the traffic usage for a service for billing purposes.

The ESS-Series routers provide a comprehensive set of service-related counters. Accounting data can be collected on a per-service, per-forwarding class basis, which enables network operators to accurately measure network usage and bill each customer for each individual service using any of a number of different billing models.

# Alcatel-Lucent Service Model

Topics in this section:

- [Service Entities on page 44](#)
  - [Customers on page 44](#)
  - [Service Access Points \(SAPs\) on page 45](#)
  - [Service Distribution Points \(SDPs\) on page 49](#)
- 

## Introduction

In the ESS-Series service model, the ESS-Series service edge routers are deployed at the provider edge. Services are provisioned on ESS-Series and transported across an IP and/or IP/MPLS provider core network in encapsulation tunnels created using Generic Router Encapsulation (GRE) or MPLS Label Switched Paths (LSPs).

The service model uses logical service entities to construct a service. The logical service entities are designed to provide a uniform, service-centric configuration, management, and billing model for service provisioning. Some benefits of this service-centric design include:

- Many services can be bound to a single customer.
- Many services can be bound to a single tunnel.
- Tunnel configurations are independent of the services they carry.
- Changes are made to a single logical entity rather than multiple ports on multiple devices. It is easier to change one tunnel rather than several services.
- The operational integrity of a logical entity (such as a service tunnel and service end points) can be verified rather than dozens of individual services improving management scaling and performance.
- A failure in the network core can be correlated to specific subscribers and services.
- QoS policies, filter policies, and accounting policies are applied to each service instead of correlating parameters and statistics from ports to customers to services.

Service provisioning uses logical entities to provision a service where additional properties can be configured for bandwidth provisioning, QoS, security filtering, accounting/billing to the appropriate entity.

## Service Entities

The basic logical entities in the service model used to construct a service are:

- [Customers](#) (see page 44)
- [Service Access Points \(SAPs\)](#) (see page 45)
- [Service Distribution Points \(SDPs\)](#) (see page 49) (for distributed services only)




Figure 8: Alcatel-Lucent's Service Entities

## Customers

The terms customers and subscribers are used synonymously. The most basic required entity is the customer ID value which is assigned when the customer account is created. To provision a service, a customer ID must be associated with the service at the time of service creation.




## Service Access Points (SAPs)

Each subscriber service type is configured with at least one service access point (SAP). A SAP identifies the customer interface point for a service on an Alcatel-Lucent ESS-Series router (Figure 9). The SAP configuration requires that slot, MDA, and port information be specified. The slot, MDA, and port parameters must be configured prior to provisioning a service (see the [Cards, MDAs, and Ports](#) section of the 7450 ESS OS Interface Configuration Guide).

A SAP is a local entity to the 7450 ESS and is uniquely identified by:

- The physical Ethernet port or SONET/SDH port
- The encapsulation type
- The encapsulation identifier (ID)

Depending on the encapsulation, a physical port can have more than one SAP associated with it. SAPs can only be created on ports designated as “access” in the physical port configuration. SAPs cannot be created on ports designated as core-facing “network” ports as these ports have a different set of features enabled in software.



OSSG002

**Figure 9: Service Access Point (SAP)**

## SAP Encapsulation Types and Identifiers

The encapsulation type is an access property of a service Ethernet port or SONET/SDH. The appropriate encapsulation type for the port depends on the requirements to support multiple services on a single port on the associated SAP and the capabilities of the downstream equipment connected to the port. For example, a port can be tagged with IEEE 802.1Q (referred to as dot1q) encapsulation in which each individual tag can be identified with a service. A SAP is created on a given port by identifying the service with a specific encapsulation ID.


---

## Ethernet Encapsulations

The following lists encapsulation service options on Ethernet ports:

- 1 Null — Supports a single service on the port. For example, where a single customer with a single service customer edge (CE) device is attached to the port. The encapsulation ID is always 0 (zero).
- 2 Dot1q — Supports multiple services for one customer or services for multiple customers (Figure 10). For example, the port is connected to a multi-tenant unit (MTU) device with multiple downstream customers. The encapsulation ID used to distinguish an individual service is the VLAN ID in the IEEE 802.1Q header.
- 3 Q-in-Q — The q-in-q encapsulation type adds a IEEE 802.1Q tag to the 802.1Q tagged packets entering the network to expand the VLAN space by tagging tagged packets, producing a double tagged frame.  
Note that the SAP can be defined with a wildcard for the inner label. (e.g. “100:\*”). In this situation all packets with an outer label of 100 will be treated as belonging to the SAP. If, on the same physical link there is also a SAP defined with q-in-q encap of 100:1 then traffic with 100:1 will go to that SAP and all other traffic with 100 as the first label will go to the SAP with the 100:\* definition.

In the dot1q and q-in-q options, traffic encapsulated with tags for which there is no definition are discarded.



**Figure 10: Multiple SAPs on a Single Port**

## SAP Considerations

When configuring a SAP, consider the following:

- A SAP is a local entity and only locally unique to a given device. The same SAP ID value can be used on another ESS-Series.
- There are no default SAPs. All SAPs must be created.
- The default administrative state for a SAP at creation time is administratively enabled.
- When a SAP is deleted, all configuration parameters for the SAP will also be deleted. For Internet Ethernet Service (IES), the IP interface must be shutdown before the SAP on that interface may be removed.
- A SAP is owned by and associated with the service in which it is created in each ESS-Series.
- A port with a dot1q or BCP-dot1q encapsulation type means the traffic for the SAP is identified based on a specific IEEE 802.1Q VLAN ID value. The VLAN ID is stripped off at SAP ingress and the appropriate VLAN ID placed on at SAP egress. As a result, VLAN IDs only have local significance, so the VLAN IDs for the SAPs for a service need not be the same at each SAP.
- If a port is administratively shutdown, all SAPs on that port will be operationally out of service.
- A SAP cannot be deleted until it has been administratively disabled (shutdown).

## Service Access Points (SAPs)

- Each SAP can be configured with only the following:
  - Ingress or egress filter policy
  - Ingress or egress QoS policy
  - Accounting policy
  - Ingress or egress scheduler policy

## Service Distribution Points (SDPs)

A service distribution point (SDP) acts as a logical way to direct traffic from one ESS-Series to another ESS-Series through a uni-directional (one-way) service tunnel. The SDP terminates at the far-end ESS-Series which directs packets to the correct service egress SAPs on that device. A distributed service consists of a configuration with at least one SAP on a local node, one SAP on a remote node, and an SDP binding the service to the service tunnel.

An SDP has the following characteristics:

- An SDP is locally unique to a participating ESS-Series. The same SDP ID can appear on other ESS-Series routers.
- An SDP uses the system IP address to identify the far-end ESS-Series edge router.
- An SDP is not specific to any one service or any type of service. Once an SDP is created, services are bound to the SDP. An SDP can also have more than one service type associated with it.
- All services mapped to an SDP use the same transport encapsulation type defined for the SDP (either GRE or MPLS).
- An SDP is a management entity. Even though the SDP configuration and the services carried within are independent, they are related objects. Operations on the SDP affect all the services associated with the SDP. For example, the operational and administrative state of an SDP controls the state of services bound to the SDP.


An SDP from the local device to a far-end ESS-Series requires a return path SDP from the far-end ESS-Series back to the local ESS-Series. Each device must have an SDP defined for every remote router to which it wants to provide service. SDPs must be created first, before a distributed service can be configured.

---

## SDP Binding

To configure a distributed service from ALA-A to ALA-B, the SDP ID (4) ([Figure 11](#)) must be specified in the service creation process in order to “bind” the service to the tunnel (the SDP). Otherwise, service traffic is not directed to a far-end point and the far-end ESS-Series device(s) cannot participate in the service (there is no service).

## Service Distribution Points (SDPs)



**Figure 11: A GRE Service Distribution Point (SDP) pointing from ALA-A to ALA-B**

## Spoke and Mesh SDPs

When an SDP is bound to a service, it is bound as either a spoke SDP or a mesh SDP. The type of SDP indicates how flooded traffic is transmitted.

A spoke SDP is treated like the equivalent of a traditional bridge “port” where flooded traffic received on the spoke SDP is replicated on all other “ports” (other spoke and mesh SDPs or SAPs) and not transmitted on the port it was received.

All mesh SDPs bound to a service are logically treated like a single bridge “port” for flooded traffic where flooded traffic received on any mesh SDP on the service is replicated to other “ports” (spoke SDPs and SAPs) and not transmitted on any mesh SDPs.

## SDP Encapsulation Types

The Alcatel-Lucent service model uses encapsulation tunnels through the core to interconnect ESS-Series service edge routers. An SDP is a logical way of referencing the entrance to an encapsulation tunnel.

The following encapsulation types are supported:

- L2 within Generic Routing Encapsulation ([GRE](#))
- L2 within RSVP signaled, loose hop non-reserved MPLS LSP
- L2 within RSVP signaled, strict hop non-reserved MPLS LSP
- L2 within RSVP-TE signaled, bandwidth reserved MPLS LSP

---

## GRE

GRE encapsulated tunnels have very low overhead and are best used for Best-Effort class of service. Packets within the GRE tunnel follow the Interior Gateway Protocol (IGP) shortest path from edge to edge. If a failure occurs within the service core network, the tunnel will only converge as fast as the IGP itself. If Equal Cost Multi-Path (ECMP) routing is used in the core, many loss-of-service failures can be minimized to sub-second timeframes.

---

## MPLS

Multi-Protocol Label Switching (MPLS) encapsulation has the following characteristics:

- LSPs (label switched paths) are used through the network, for example, primary, secondary, loose hop, etc. These paths define how traffic traverses the network from point A to B. If a path is down, depending on the configuration parameters, another path is substituted.

Paths can be manually defined or a constraint-based routing protocol (e.g., OSPF-TE or CSPF) can be used to determine the best path with specific constraints.

- A ESS-Series router supports both signaled and non-signaled LSPs through the network.
- Non-signaled paths are defined at each hop through the network.
- Signaled paths are communicated via protocol from end to end using Resource Reservation Protocol (RSVP).

Because services are carried in encapsulation tunnels and an SDP is an entrance to the tunnel, an SDP has an implicit Maximum Transmission Unit (MTU) value. The MTU for the service tunnel can affect and interact with the MTU supported on the physical port where the SAP is defined.

## SDP Keepalives

SDP keepalives are a way of actively monitoring the SDP operational state using periodic Alcatel-Lucent SDP Ping Echo Request and Echo Reply messages. Alcatel-Lucent SDP Ping is a part of Alcatel-Lucent's suite of Service Diagnostics built on an Alcatel-Lucent service-level OA&M protocol. When SDP Ping is used in the SDP keepalive application, the SDP Echo Request and Echo Reply messages are a mechanism for exchanging far-end SDP status.

Configuring SDP keepalives on a given SDP is optional. SDP keepalives for a particular SDP have the following configurable parameters:

- AdminUp/AdminDown State
- Hello Time
- Message Length
- Max Drop Count
- Hold Down Time

SDP keepalive Echo Request messages are only sent when the SDP is completely configured and administratively up and SDP keepalives is administratively up. If the SDP is administratively down, keepalives for the SDP are disabled.

SDP keepalive Echo Request messages are sent out periodically based on the configured Hello Time. An optional Message Length for the Echo Request can be configured. If Max Drop Count Echo Request messages do not receive an Echo Reply, the SDP will immediately be brought operationally down.

If a keepalive response is received that indicates an error condition, the SDP will immediately be brought operationally down.

Once a response is received that indicates the error has cleared and the Hold Down Time interval has expired, the SDP will be eligible to be put into the operationally up state. If no other condition prevents the operational change, the SDP will enter the operational state.




## Epipe Service Overview

An Epipe service is Alcatel-Lucent's implementations of an Ethernet VLL based on the IETF "Martini Drafts" (draft-martini-l2circuit-trans-mpls-08.txt and draft-martini-l2circuit-encapmpls-04.txt) and the IETF Ethernet Pseudo-wire Draft (draft-so-pwe3-ethernet-00.txt).

An Epipe service is a Layer 2 point-to-point service where the customer data is encapsulated and transported across a service provider's IP or MPLS network. An Epipe service is completely transparent to the subscriber's data and protocols. The 7450 ESS Epipe service does not perform any MAC learning. A local Epipe service consists of two SAPs on the same node, whereas a distributed Epipe service consists of two SAPs on different nodes. SDPs are not used in local Epipe services.

Each SAP configuration includes a specific port on which service traffic enters the 7450 ESS from the customer side (also called the access side). Each port is configured with an encapsulation type. If a port is configured with an IEEE 802.1Q (referred to as dot1q) encapsulation, then a unique encapsulation value (ID) must be specified.



**Figure 12: Epipe/VLL Service**

## VPLS Service Overview

Virtual Private LAN Service (VPLS) as described in Internet Draft *draft-ietf-ppvpn-vpls-ldp-01.txt*, is a class of virtual private network service that allows the connection of multiple sites in a single bridged domain over a provider-managed IP/MPLS network. The customer sites in a VPLS instance appear to be on the same LAN, regardless of their location. VPLS uses an Ethernet interface on the customer-facing (access) side which simplifies the LAN/WAN boundary and allows for rapid and flexible service provisioning.

VPLS offers a balance between point-to-point Frame Relay service and outsourced routed services. VPLS enables each customer to maintain control of their own routing strategies. All customer routers in the VPLS service are part of the same subnet (LAN) which simplifies the IP addressing plan, especially when compared to a mesh constructed from many separate point-to-point connections. The VPLS service management is simplified since the service is not aware of nor participates in the IP addressing and routing.

A VPLS service provides connectivity between two or more SAPs on one (which is considered a local service) or more (which is considered a distributed service) 7450 ESS routers. The connection appears to be a bridged domain to the customer sites so protocols, including routing protocols, can traverse the VPLS service.

Other VPLS advantages include:

- VPLS is a transparent, protocol-independent service.
- There is no Layer 2 protocol conversion between LAN and WAN technologies.
- There is no need to design, manage, configure, and maintain separate WAN access equipment, thus, eliminating the need to train personnel on WAN technologies such as Frame Relay.

For details on VPLS, including a packet walkthrough, refer to VPLS section in the SR-OS Services Guide.

---

## Split Horizon SAP Groups and Split Horizon Spoke SDP Groups

Within the context of VPLS services, a loop-free topology within a fully meshed VPLS core is achieved by applying split-horizon forwarding concept that packets received from a mesh SDP are never forwarded to other mesh SDPs within the same service. The advantage of this approach is that no protocol is required to detect loops within the VPLS core network.

In applications such as DSL aggregation, it is useful to extend this split-horizon concept also to groups of SAPs and/or spoke SDPs. This extension is referred to as a split horizon SAP group or residential bridging.

Traffic arriving on a SAP or a spoke SDP within a split horizon group will not be copied to other SAPs and spoke SDPs in the same split horizon group (but will be copied to SAPs / spoke SDPs in other split horizon groups if these exist within the same VPLS).

---

## Residential Split Horizon Groups

To improve the scalability of a SAP-per-subscriber model in the broadband services aggregator (BSA), the ESS-Series supports a variant of split horizon groups called residential split horizon groups (RSHG).

A RSHG is a group of split horizon group SAPs with following limitations:

- Downstream broadcast traffic is not allowed.
- Downstream multicast traffic is allowed when IGMP snooping is configured in the VPLS.
- STP is not supported.


Spoke SDPs can also be members of a RSHG VPLS. The downstream multicast traffic restriction does not apply to spoke SDPs.

## IES Service Overview

Internet Enhanced Service (IES) is a routed connectivity service where the subscriber communicates with an IP router interface to send and receive Internet traffic. An IES has one or more logical IP routing interfaces each with a SAP which acts as the access point to the subscriber's network. IES allow customer-facing IP interfaces to participate in the same routing instance used for service network core routing connectivity. IES services require that the IP addressing scheme used by the subscriber be unique between other provider addressing schemes and potentially the entire Internet.

While IES is part of the routing domain, the usable IP address space may be limited. This allows a portion of the service provider address space to be reserved for service IP provisioning, and be administered by a separate but subordinate address authority.

IP interfaces defined within the context of an IES service must have a SAP associated as the access point to the subscriber network. Multiple IES services are created to segregate subscriber-owned IP interfaces.



**Figure 13: Internet Enhanced Service**

The IES service provides Internet connectivity. Other features include:

- Multiple IES services are created to separate customer-owned IP interfaces.
- More than one IES service can be created for a single customer ID.
- More than one IP interface can be created within a single IES service ID. All IP interfaces created within an IES service ID belong to the same customer.

## IP Interface

IES customer IP interfaces can be configured with most of the same options found on the core IP interfaces. The advanced configuration options supported are:

- VRRP (for IES services with more than one IP interface)
- Secondary IP addresses
- ICMP options

Configuration options found on core IP interfaces not supported on IES IP interfaces are:

- Unnumbered interfaces
- NTP broadcast receipt

## Deploying and Provisioning Services

The service model provides a logical and uniform way of constructing connectivity services. The basic steps for deploying and provisioning services can be broken down into three phases.

---

### Phase 1: Core Network Construction

Before the services are provisioned, the following tasks should be completed:

- Build the IP or IP/MPLS core network.
  - Configure routing protocols.
  - Configure MPLS LSPs (if MPLS is used).
  - Construct the core SDP service tunnel mesh for the services.
- 

### Phase 2: Service Administration

Perform preliminary policy and SDP configurations to control traffic flow, operator access, and to manage fault conditions and alarm messages, the following tasks should be completed:

- Configure group and user access privileges.
  - Build templates for QoS, filter and/or accounting policies needed to support the core services.
- 

### Phase 3: Service Provisioning

- Provision customer account information.
- If necessary, build any customer-specific QoS, filter or accounting policies.
- Provision the customer services on the ESS-Series service edge routers by defining SAPs, binding policies to the SAPs, and then binding the service to appropriate SDPs as necessary.

# Configuration Notes

This section describes service configuration caveats.

---

## General

Service provisioning tasks can be logically separated into two main functional areas, core tasks and subscriber tasks and are typically performed prior to provisioning a subscriber service.

Core tasks include the following:

- Create customer accounts
- Create template QoS, filter, scheduler, and accounting policies
- Create LSPs
- Create SDPs

Subscriber services tasks include the following:

- Create VLL, VPLS, or IES services
- Configure interfaces (where required) and SAPs
- Bind SDPs
- Create exclusive QoS and filter policies





## Configuring Triple Play Services with CLI

This section provides information to configure Residential Broadband Aggregation services using the command line interface. It is assumed that the reader is familiar with basic configuration of VPLS and IES services.

Topics in this section include:

- [Configuring VPLS Residential Split Horizon Groups on page 61](#)
- [Configuring Static Hosts on page 62](#)
- [Configuring Static Hosts on an IES SAP on page 64](#)
- [Configuring Static Hosts on an VPRN SAP on page 64](#)

---

## Configuring VPLS Residential Split Horizon Groups

To configure a group of SAPs in a VPLS service as a Residential Split Horizon Group (RSHG), add the residential-group parameter when creating the split horizon group. Traffic arriving on a SAP within an RSHG will not be copied to other SAPs in the same split horizon group. Note that the split horizon group must be created before it can be applied.

The following example displays a VPLS configuration with split horizon enabled:

```
*A:ALA-48>config>service>vpls# info
-----
      split-horizon-group "DSL-group2" residential-group create
      description "split horizon group for DSL - no broadcast supported"
      exit
      stp
      shutdown
      exit
      sap 2/1/4:100 split-horizon-group "DSL-group2" create
      description "SAP in RSHG"
      exit
      sap 2/1/4:200 split-horizon-group "DSL-group2" create
      description "another SAP in the RSHG"
      exit
      no shutdown
-----
*A:ALA-48>config>service>vpls#
```

## Configuring Static Hosts

In order for the static host to be operational, forwarding traffic bi-directional, the mac address of the host must be learned or configured. Learning the MAC of the static host is different for IPv4 or IPv6.

If an IPv4 static host MAC is not specified

- The system will learn respective MAC address dynamically from ARP packets (arp-request or gratuitous-arp) generated by the host with the specified IP address.
- On a VPLS service, this can occur if arp-reply-agent function is enabled on a given SAP. On Layer 3 services, such as IES or VPRN, the ARP packets are always examined so no further conditions are applicable.

If an IPv6 static host MAC is not specified

- The system learns the MAC address depending on the type of host configured such as:  
IPv6 prefix host or IPv6 address host.
  - A SAP can be specified as a single-MAC and it implies that there is only a single device attached to the SAP. It changes the MAC learning behavior on the SAP for IPv6 host only. Firstly, all IPv6 hosts will share the same learned MAC. Secondly, the MAC address is learned from the host's router solicits and neighbor discoveries.
  - For static host with an address, upon shutdown, a RS for the clients IPv6 address is sent towards the host and the MAC is learned upon the RA reply
  - For static host with an address, SHCV will send the RS, and the MAC is learned from the RA
  - For static host with an address, the OAM command can trigger a RS and the MAC is learned from the RA
  - For static host with either a prefix or an address, linking the IPv6 host to an IPv4 host will copy the IPv4 host MAC address to the IPv6 host and vice versa.

The learned MAC address will be handled as a MAC address of static host with explicitly defined mac-address. Meaning:

- The MAC address will not be aged by the mac-aging or any other aging timers.
- The MAC address will not be moved to another SAP as a consequence of re-learning event (= event when learning request for the same MAC address comes from another SAP).
- The MAC address will not be flushed from FDB due to SAP failure or STP flush messages.

Every time the given static-host uses different MAC address in its ARP request, the dynamic MAC learning process will be performed. The old MAC address will be overwritten by a new MAC address.

The learned MAC address will not be made persistent (a static host is not a part of the persistency file). A service discontinuity of such a host could be proportional to its arp-cache timeout.

The following interactions are described:

- Antispoof (all services) — In case a static IP-only host is configured on a given SAP, anti-spoof types, IP, NH MAC, and IP MAC are supported. Static hosts for which MAC address is not known will not have any antispoof entry. This will be added only after the corresponding MAC has been learned. As a consequence, all traffic generated by the host before the MAC is learned are dropped.
- MAC-linking (IES and VPRN service only) — the MAC address can be learned from either the IPv4 or IPv6 host. Once learned it is copied over to the host of the other address family
- Single-MAC (IES and VPRN service only) — This specifies that there is only one single subscriber (MAC) on the SAP and any ICMP6 message from the SAP can be assumed to be the subscriber MAC address. This does not apply to IPv4 host.
- Enhanced subscriber management (all services) — ESM is supported in a combination with a static ip-only host. It is assumed that ip-mac antispoofing is enabled. The resources (queues, etc.) are allocated at the time such a host is configured, although they will be effectively used only after antispoof entry has been installed.
- Dual-homing (for IPv4 host only) — It is assumed that static host is configured on both chassis. The dynamic mac-address learning event will be then synchronized (also, if the members are on two different nodes) and corresponding anti-spoof entries will be installed on both chassis.
- MAC-pinning (for VPLS services only) — The dynamically learned MAC address of the static-host will be considered as a static-mac and is not affected by the no mac-pinning command.
- ARP-reply-agent (VPLS services only) — It is possible to enable arp-reply-agent on a SAP where static host with ip-only configured. Besides the regular arp-reply-agent functionality (reply to all arp-requests targeting the given host's IP address) learning of the host's MAC address will be performed. As long as no MAC address have been learned no ARP replies on behalf of such host should be expected. Enabling of arp-reply-agent is optional for SAP with ip-only static hosts.

### Configuring Static Hosts on an VPLS SAP

The following example displays a static host on a VPLS SAP configuration:

```
*A:ALA-48>config>service# info
-----
vpls 800 customer 6001 create
  description "VPLS with residential split horizon for DSL"
  stp
    shutdown
  exit
  sap 1/2/7:100 split-horizon-group "DSL-group2" create
    description "SAP for RSHG"
    static-host ip 10.1.1.1
  exit
  no shutdown
-----
*A:ALA-48>config>service#
```

---

### Configuring Static Hosts on an IES SAP

The following displays a static host on an IES SAP:

```
*A:ALA-49>config>service>ies>sub-if>grp-if# info
-----
  sap 7/1/5 create
  description "IES with static host"
    static-host ip 10.1.1.1 create
    static-host ip 2001::1/128 create
    static-host ip 2001:1::/64 create
  exit
-----
*A:ALA-49>config>service>ies#
```

---

### Configuring Static Hosts on an VPRN SAP

The following displays a static host on a VPRN SAP:

```
*A:ALA-49>config>service>vprn>sub-if>grp-if# info
-----
  description "VPRN service with static host"
  sap 7/1/5 create
    description "IES with static host"
    static-host ip 10.1.1.1 create
    static-host ip 2001::1/128 create
    static-host ip 2001:1::/64 create
  exit
-----
*A:ALA-49>config>service>vprn#
```

# Triple Play Services Command Reference

## Configuration Commands

Note: The command trees in this section are limited to those commands specific to Triple Play services. For the full command trees for a specific service type refer to the appropriate section in the 7450 ESS Services Overview Guide.

- [Generic VPLS Triple Play Commands on page 65](#)
- [Generic IES Triple Play Commands on page 69](#)
- [Service DHCP and Anti-Spoof Filtering Commands on page 71](#)
- [Triple Play ARP Commands on page 73](#)
- [Triple Play Multicast Commands on page 74](#)
- [Show Commands on page 76](#)
- [Clear Commands on page 78](#)

## Generic VPLS Triple Play Commands

```

config
— service
    — [no] vpls service-id [customer customer-id] [vpn vpn-id] [m-vpls]
        — description description-string
        — no description
        — [no] shutdown
        — [no] disable-aging
        — [no] disable-learning
        — [no] discard-unknown
        — [no] fdb-table-high-wmark
        — [no] fdb-table-low-wmark
        — fdb-table-size table-size
        — no fdb-table-size [table-size]
        — igmp-snooping
            — mvr
                — description description-string
                — no description
                — group-policy policy-name
                — no group-policy
                — [no] shutdown
            — query-interval seconds
            — no query-interval
            — query-src-ip ip-address
            — no query-src-ip
            — report-src-ip ip-address
            — no report-src-ip
            — robust-count robust-count

```

```

— no robust-count
— [no] shutdown
— local-age aging-timer
— no local-age
— mac-protect
— [no] mac ieee-address
— mac-subnet-length subnet-length
— no mac-subnet-length
— mcr-default-gtw
— ip address
— no ip
— mac ieee-address
— no mac
— remote-age seconds
— no remote-age
— service-mtu octets
— no service-mtu
— service-name service-name
— no service-name
— [no] split-horizon-group [group-name] [residential-group]
— description description-string
— no description
— sap sap-id [split-horizon-group group-name]
— no sap sap-id
— accounting-policy acct-policy-id
— no accounting-policy [acct-policy-id]
— arp-host
— host-limit max-num-hosts
— no host-limit
— min-auth-interval min-auth-interval
— no min-auth-interval
— [no] shutdown
— description description-string
— no description
— [no] collect-stats
— [no] disable-aging
— [no] disable-learning
— egress
— filter ip ip-filter-id
— filter mac mac-filter-id
— no filter
— no filter [ip ip-filter-id] [mac mac-filter-id]
— qos policy-id
— no qos
— queue-override
— [no] queue queue-id
— adaptation-rule [pir {max|min|closest}] [cir
{max | min | closest}]
— no adaptation-rule
— avg-frame-overhead percentage
— no avg-frame-overhead
— cbs size-in-kbytes
— no cbs
— high-prio-only percent
— no high-prio-only
— mbs size-in-kbytes
— no mbs

```

```

— rate pir-rate [cir cir-rate]
— no rate
— [no] scheduler-override
— [no] scheduler scheduler-name
— rate pir-rate [cir cir-rate]
— no rate
— scheduler-policy scheduler-policy-name
— no scheduler-policy
— host {[ip ip-address] [mac mac-address]} [subscriber-sap-id | subscriber
sub-ident-string] [sub-profile sub-profile-name] [sla-profile sla-profile-
name] [ancp-string ancp-string] [inter-dest-id intermediate-destination-
id]
— no host {[ip ip-address] [mac mac-address]}
— no host all
— ingress
— filter ip ip-filter-id
— filter mac mac-filter-id
— no filter
— no filter [ip ip-filter-id] [mac mac-filter-id]
— qos policy-id
— no qos
— queue-override
— [no] queue queue-id
— adaptation-rule [pir {max|min|closest}] [cir
{max | min | closest}]
— no adaptation-rule
— cbs size-in-kbytes
— no cbs
— high-prio-only percent
— no high-prio-only
— mbs size-in-kbytes
— no mbs
— rate pir-rate [cir cir-rate]
— no rate
— [no] scheduler-override
— [no] scheduler scheduler-name
— rate pir-rate [cir cir-rate]
— no rate
— scheduler-policy scheduler-policy-name
— no scheduler-policy
— scheduler-policy scheduler-policy-name
— no scheduler-policy
— max-nbr-mac-addr table-size
— no max-nbr-mac-addr
— multi-service-site customer-site-name
— no multi-service-site
— static-host ip ip/did-address [mac ieee-address] [create]
— static-host mac ieee-address [create]
— no static-host [ip ip-address] mac ieee-address
— no static-host all [force]
— no static-host ip ip-address
— ancp-string ancp-string
— no ancp-string
— app-profile app-profile-name
— no app-profile

```

```

— inter-dest-id intermediate-destination-id
— no inter-dest-id
— [no] shutdown
— sla-profile sla-profile-name
— no sla-profile
— sub-profile sub-profile-name
— no sub-profile
— subscriber sub-ident
— no subscriber
— [no] subscriber-sap-id
— [no] shutdown
— [no] static-mac
— mesh-sdp sdp-id[:vc-id] [vc-type {ether | vlan}]
— no mesh-sdp sdp-id[:vc-id]
— accounting-policy acct-policy-id
— no accounting-policy [acct-policy-id]
— [no] collect-stats
— egress
— filter {ip ip-filter-name | mac mac-filter-id}
— no filter
— vc-label egress-vc-label
— no vc-label [egress-vc-label]
— ingress
— filter {ip ip-filter-name | mac mac-filter-id}
— no filter
— vc-label ingress-vc-label
— no vc-label [ingress-vc-label]
— [no] shutdown
— [no] static-mac
— vlan-vc-tag 0..4094
— no vlan-vc-tag [0..4094]
— spoke-sdp sdp-id[:vc-id] [vc-type {ether | vlan}] [split-horizon-group group-name]
— no spoke-sdp sdp-id[:vc-id]
— accounting-policy acct-policy-id
— no accounting-policy [acct-policy-id]
— [no] block-on-mesh-failure
— [no] collect-stats
— egress
— filter {ip ip-filter-name | mac mac-filter-id}
— no filter
— vc-label egress-vc-label
— no vc-label [egress-vc-label]
— ingress
— filter {ip ip-filter-name | mac mac-filter-id}
— no filter
— vc-label ingress-vc-label
— no vc-label [ingress-vc-label]
— max-nbr-mac-addr table-size
— no max-nbr-mac-addr
— [no] shutdown
— [no] static-mac
— vlan-vc-tag 0..4094
— no vlan-vc-tag [0..4094]

```



## Generic IES Triple Play Commands

```

config
  — service
    — ies service-id [customer customer-id] [vpn vpn-id]
      — description description-string
      — no description
      — service-name service-name
      — no service-name
      — [no] shutdown
      — [no] interface ip-int-name
        — address ip-address/mask [netmask] [broadcast {all-ones | host-ones}]
        — no address
        — [no] allow-directed-broadcast
        — description description-string
        — no description
        — [no] loopback
        — mac ieee-address
        — no mac
        — [no] sap sap-id
          — accounting-policy acct-policy-id
          — no accounting-policy [acct-policy-id]
          — anti-spoof [ip | mac | ip-mac]
          — no anti-spoof
          — [no] collect-stats
          — description description-string
          — no description
          — egress
            — filter {ip ip-filter-name | mac mac-filter-id}
            — no filter
            — qos policy-id
            — no qos
            — [no] queue-override
            — [no] queue queue-id
              — adaptation-rule [pir {max|min|closest}]
                [cir {max | min | closest}]
              — no adaptation-rule
              — avg-frame-overhead percentage
              — no avg-frame-overhead
              — cbs size-in-kbytes
              — no cbs
              — high-prio-only percent
              — no high-prio-only
              — mbs size-in-kbytes
              — no mbs
              — rate pir-rate [cir cir-rate]
              — no rate
            — scheduler-policy scheduler-policy-name
            — no scheduler-policy
          — host {[ip ip-address] [mac mac-address]} [subscriber sub-ident-string] [sub-profile sub-profile-name] [sla-profile sla-profile-name]
          — no host {[ip ip-address] [mac mac-address]}
          — no host all
          — ingress

```

```

— filter {ip ip-filter-name | mac mac-filter-id}
— no filter
— match-qinq-dot1p {top | bottom}
— no match-qinq-dot1p
— qos policy-id
— no qos
— scheduler-policy scheduler-policy-name
— no scheduler-policy
— multi-service-site customer-site-name
— no multi-service-site
— static-host ip ip/did-address [mac ieee-address] [create]
— static-host mac ieee-address [create]
— no static-host [ip ip-address] mac ieee-address
— no static-host all [force]
— no static-host ip ip-address
   — ancp-string ancp-string
   — no ancp-string
   — app-profile app-profile-name
   — no app-profile
   — inter-dest-id intermediate-destination-id
   — no inter-dest-id
   — [no] shutdown
   — sla-profile sla-profile-name
   — no sla-profile
   — sub-profile sub-profile-name
   — no sub-profile
   — subscriber sub-ident
   — no subscriber
   — [no] subscriber-sap-id
— [no] shutdown

```

## Service DHCP and Anti-Spoof Filtering Commands

```

config
  — service
    — [no] vpls service-id [customer customer-id] [vpn vpn-id] [m-vpls]
      — sap sap-id [split-horizon-group group-name] [capture-sap]
      — no sap sap-id
        — anti-spoof {ip | mac | ip-mac}
        — no anti-spoof
        — authentication-policy auth-plcy-name
        — no authentication-policy
        — diameter-auth-policy name
        — no diameter-auth-policy
        — dhcp
          — description description-string
          — no description
          — lease-populate [nbr-of-leases]
          — no lease-populate
          — [no] option
            — action [dhcp-action]
            — no action
            — [no] circuit-id
            — [no] remote-id
          — [no] shutdown
          — [no] snoop
      — mesh-sdp sdp-id[:vc-id] [vc-type {ether | vlan}]
      — no mesh-sdp sdp-id[:vc-id]
        — dhcp
          — description description-string
          — no description
          — [no] snoop
      — spoke-sdp sdp-id[:vc-id] [vc-type {ether | vlan}] [split-horizon-group group-name]
      — no spoke-sdp sdp-id[:vc-id]
        — dhcp
          — description description-string
          — no description
          — [no] snoop

config
  — service
    — ies service-id [customer customer-id] [vpn vpn-id]
      — [no] interface ip-int-name
        — anti-spoof {ip | mac | ip-mac}
        — no anti-spoof
        — dhcp
          — description description-string
          — no description
          — gi-address ip-address [src-ip-addr]
          — no gi-address
          — lease-populate nbr-of-leases
          — no lease-populate
          — option
            — action {replace | drop | keep}

```

```

— no action
— circuit-id [ascii-tuple | ifindex]
— no circuit-id
— [no] remote-id
— [no] relay-plain-bootp
— relay-unicast-msg [release-update-src-ip]
— no relay-unicast-msg
— server server1 [server2...(up to 8 max)]
— no server
— [no] shutdown
— [no] trusted

config
— service
— vprn
— [no] interface ip-int-name
— dhcp
— description description-string
— no description
— gi-address ip-address [src-ip-addr]
— no gi-address
— lease-populate [nbr-of-leases]
— no lease-populate
— [no] option
— action {replace | drop | keep}
— no action
— circuit-id [ascii-tuple | ifindex]
— no circuit-id
— [no] remote-id
— [no] relay-plain-bootp
— relay-unicast-msg [release-update-src-ip]
— no relay-unicast-msg
— server server1 [server2...(up to 8 max)]
— no server
— [no] shutdown
— [no] trusted

```

## Triple Play ARP Commands

```

config
— service
    — [no] vpls service-id [customer customer-id] [vpn vpn-id] [m-vpls]
        — sap sap-id [split-horizon-group group-name]
        — no sap sap-id
            — arp-reply-agent [sub-ident]
            — no arp-reply-agent

```

```

config
— service
    — ies service-id [customer customer-id] [vpn vpn-id]
        — [no] interface ip-int-name
            — [no] arp-populate
            — arp-timeout seconds
            — no arp-timeout
            — [no] local-proxy-arp
            — [no] proxy-arp-policy
            — [no] remote-proxy-arp
            — [no] sap sap-id

```

## Triple Play Multicast Commands

```

config
— service
— [no] vpls service-id [customer customer-id] [vpn vpn-id] [m-vpls]
— igmp-snooping
— mvr
— group-policy policy-name
— [no] shutdown
— query-interval seconds
— no query-interval
— robust-count robust-count
— no robust-count
— [no] shutdown
— mfib-table-high-wmark high-water-mark
— no mfib-table-high-wmark
— mfib-table-low-wmark low-water-mark
— no mfib-table-low-wmark
— mfib-table-size table-size
— no mfib-table-size
— sap sap-id [split-horizon-group group-name]
— no sap ap-id
— igmp-snooping
— [no] fast-leave
— import policy-name
— no import
— last-member-query-interval interval
— no last-member-query-interval
— max-num-groups max-num-groups
— no max-num-groups
— [no] mrrouter-port
— mvr
— from-vpls vpls-id
— no from-vpls
— to-sap sap-id
— no to-sap
— query-interval interval
— no query-interval
— query-response-interval interval
— no query-response-interval
— robust-count count
— no robust-count
— [no] send-queries
— static
— [no] group group-address
— [no] source ip-address
— [no] starg
— mesh-sdp sdp-id[:vc-id] [vc-type {ether | vlan | vpls}]
— no mesh-sdp sdp-id[:vc-id]
— igmp-snooping
— [no] fast-leave
— import policy-name
— no import
— last-member-query-interval interval
— no last-member-query-interval
— max-num-groups max-num-groups

```

```

— no max-num-groups
— query-interval interval
— no query-interval
— query-response-interval interval
— no query-response-interval
— robust-count count
— no robust-count
— [no] send-queries
— static
    — [no] group group-address
    — [no] source ip-address
    — [no] starg
— spoke-sdp sdp-id[:vc-id] [vc-type {ether | vlan | vpls}] [split-horizon-group
group-name]
— no spoke-sdp sdp-id[:vc-id]
    — igmp-snooping
        — [no] fast-leave
        — import policy-name
        — no import
        — last-member-query-interval interval
        — no last-member-query-interval
        — max-num-groups max-num-groups
        — no max-num-groups
        — [no] mrouter-port
        — query-interval interval
        — no query-interval
        — query-response-interval interval
        — no query-response-interval
        — robust-count count
        — no robust-count
        — [no] send-queries
        — static
            — [no] group group-address
            — [no] source ip-address
            — [no] starg

```

## Show Commands

```

show
  — service
    — active-subscribers detail
    — active-subscribers mirror
    — active-subscribers [summary]
      — credit-control credit-control [subscriber sub-ident-string]
      — credit-control out-of-credit [action action] [summary]
      — filter [subscriber sub-ident-string] [origin origin]
      — hierarchy [subscriber sub-ident-string]
      — host-tracking [subscriber sub-ident-string]
      — host-tracking [subscriber sub-ident-string] detail
      — host-tracking [subscriber sub-ident-string] summary
      — host-tracking [subscriber sub-ident-string] statistics
        — groups [group group-ip-address]
        — groups group group-ip-address detail
        — groups group group-ip-address summary
      — igmp [subscriber sub-ident-string][detail]
      — subscriber sub-ident-string
      — subscriber sub-ident-string detail
      — subscriber sub-ident-string mirror
      — subscriber sub-ident-string sap sap-id sla-profile sla-profile-name
      — subscriber sub-ident-string sap sap-id sla-profile sla-profile-name detail
      — subscriber sub-ident-string sap sap-id sla-profile sla-profile-name mirror
  — id service-id
    — arp [ip-address] | [mac ieee-address] | [sap port-id:encap] | [interface ip-int-name]
    — base
    — authentication
      — statistics
    — dhcp
      — 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]}] [detail]
      — statistics [sap sap-id] | [sdp [sdp-id[:vc-id]]]
      — summary
    — gsm
      — neighbors group [name] [ip-address]
      — sessions [group name] neighbor ip-address [port port-number] [association] [statistics]
    — host [sap sap-id] [wholesaler service-id] [port port-id] [inter-dest-id intermediate-destination-id] [detail]
    — host [sap sap-id] [wholesaler service-id] [port port-id] no-inter-dest-id [detail]
    — host summary
    — retailers
    — split-horizon-group [group-name]
    — static-host [sap sap-id] [wholesaler service-id] [port port-id] [inter-dest-id intermediate-destination-id] [detail]
    — static-host [sap sap-id] [wholesaler service-id] [port port-id] no-inter-dest-id [detail]
    — static-host summary
    — wholesalers
  — subscriber-using [service-id service-id] [sap-id sap-id] [interface ip-int-name] [ip ip-address[/mask]] [mac ieee-address] [sub-profile sub-profile-name] [sla-profile sla-profile-

```



*name*] [**app-profile** *app-profile-name*] [**port** *port-id*] [**no-inter-dest-id** | **inter-dest-id** *intermediate-destination-id*]

- **router**
  - **dhcp**
    - **lease-state** [**interface** *ip-int-name* | *ip-address*]
    - **statistics** [*ip-int-name* | *ip-address*]
    - **summary**

**show**

- **service**
  - **id** *service-id*
    - **igmp-snooping**
      - **all**
      - **base**
      - **mrouters** [**detail**]
      - **mvr**
      - **port-db** {**sap** *sap-id* | **sdp** *sdp-id:vc-id*} [**group** *grp-address*] | **detail**
      - **proxy-db** [**group** *grp-address* | **detail**]
      - **querier**
      - **static** [**sap** *sap-id* | **sdp** *sdp-id:vc-id*]
      - **statistics** [**sap** *sap-id* | **sdp** *sdp-id:vc-id*]
    - **mfib** [**brief** | **group** *grp-address* | **statistics** [**group** *grp-address*]]

## Clear Commands

```
clear
— router
— dhcp
— lease-state [interface ip-int-name | ip-addr | ip-address ip-address | mac ieee-address]
— statistics [ip-int-name | ip-address]
— service
— id service-id
— authentication
— statistics
— dhcp
— lease-state [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 sap sap-id [no-dhcp-release]
— lease-state sdp sdp-id[:vc-id] [no-dhcp-release]
— igmp-snooping
— port-db sap sap-id [group grp-address [source ip-address]]
— port-db sdp sdp-id[:vc-id] [group grp-address [source ip-address]]
— querier
— statistics [all | sap sap-id | sdp sdp-id[:vc-id]]
— mfib
— statistics [all | group grp-address]
— mld-snooping
— port-db sap sap-id [group grp-ipv6-address]
— port-db sap sap-id group grp-ipv6-address source src-ipv6-address
— port-db sdp sdp-id[:vc-id] [group grp-ipv6-address]
— port-db sdp sdp-id[:vc-id] group grp-ipv6-address source src-ipv6-address
— querier
— statistics all
— statistics sap sap-id
— statistics sdp sdp-id[:vc-id]
```

---

# Triple Play Service Configuration Commands

---

## Global Commands

---

### shutdown

<b>Syntax</b>	<b>[no] shutdown</b>
<b>Context</b>	<pre> config&gt;service&gt;vpls&gt;sap&gt;dhcp config&gt;service&gt;vpls&gt;igmp-snooping&gt;mvr config&gt;service&gt;vpls config&gt;service&gt;vpls&gt;sap config&gt;service&gt;vpls&gt;sap&gt;arp-host config&gt;service&gt;vpls&gt;mesh-sdp config&gt;service&gt;vpls&gt;spoke-sdp config&gt;service&gt;ies&gt;if&gt;sap config&gt;service&gt;ies&gt;if&gt;dhcp </pre>
<b>Description</b>	<p>This command administratively disables an entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics.</p> <p>The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they may be deleted.</p> <p>Services are created in the administratively down (<b>shutdown</b>) state. When a <b>no shutdown</b> command is entered, the service becomes administratively up and then tries to enter the operationally up state. Default administrative states for services and service entities is described below in Special Cases.</p> <p>The <b>no</b> form of this command places the entity into an administratively enabled state.</p>
<b>Special Cases</b>	<p><b>Service Admin State</b> — Bindings to an SDP within the service will be put into the out-of-service state when the service is shutdown. While the service is shutdown, all customer packets are dropped and counted as discards for billing and debugging purposes.</p> <p><b>Service Operational State</b> — A service is regarded as operational providing that two SAPs or if one SDP are operational.</p> <p><b>SDP (global)</b> — When an SDP is shutdown at the global service level, all bindings to that SDP are put into the out-of-service state and the SDP itself is put into the administratively and operationally down states. Packets that would normally be transmitted using this SDP binding will be discarded and counted as dropped packets.</p> <p><b>SDP (service level)</b> — Shutting down an SDP within a service only affects traffic on that service from entering or being received from the SDP. The SDP itself may still be operationally up for other services.</p> <p><b>SDP Keepalives</b> — Enables SDP connectivity monitoring keepalive messages for the SDP ID. Default state is disabled (shutdown) in which case the operational state of the SDP-ID is not affected by the keepalive message state.</p>

**VPLS SAPs and SDPs** — SAPs are created in a VPLS and SDPs are bound to a VPLS in the administratively up default state. The created SAP will attempt to enter the operationally up state. An SDP will attempt to go into the in-service state once bound to the VPLS.

## description

<b>Syntax</b>	<b>description</b> <i>description-string</i> <b>no description</b>
<b>Context</b>	config>service>vpls config>service>vpls>igmp-snooping>mvr config>service>vpls>split-horizon-group config>service>vpls>sap config>service>ies>if>sap config>service>vpls>sap>dhcp config>service>vpls>mld-snooping>mvr
<b>Description</b>	This command creates a text description stored in the configuration file for a configuration context.  The <b>description</b> command associates a text string with a configuration context to help identify the content in the configuration file.  The <b>no</b> form of this command removes the string from the configuration.
<b>Default</b>	No description associated with the configuration context.
<b>Parameters</b>	<i>description-string</i> — The description character string. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

## Service Commands

### vpls

<b>Syntax</b>	<b>vpls</b> <i>service-id</i> <b>customer</b> <i>customer-id</i> <b>vpn</b> <i>vpn-id</i> [ <b>m-vpls</b> ] <b>vpls</b> <i>service-id</i> <b>no vpls</b> <i>service-id</i>
<b>Context</b>	config>service
<b>Description</b>	<p>This command creates or edits a Virtual Private LAN Services (VPLS) instance.</p> <p>The <b>vpls</b> command is used to create or maintain a VPLS service. If the <i>service-id</i> does not exist, a context for the service is created. If the <i>service-id</i> exists, the context for editing the service is entered.</p> <p>A VPLS service connects multiple customer sites together acting like a zero-hop, layer 2 switched domain. A VPLS is always a logical full mesh.</p> <p>When a service is created, the <b>create</b> keyword must be specified if the <b>create</b> command is enabled in the <b>environment</b> context.</p> <p>When a service is created, the <b>customer</b> keyword and <i>customer-id</i> must be specified and associates the service with a customer. The <i>customer-id</i> must already exist having been created using the <b>customer</b> command in the service context. Once a service has been created with a customer association, it is not possible to edit the customer association. The service must be deleted and recreated with a new customer association.</p> <p>To create a management VPLS, the m-vpls keyword must be specified. See section <b>Hierarchical VPLS Redundancy</b> for an introduction to the concept of management VPLS.</p> <p>Once a service is created, the use of the <b>customer</b> <i>customer-id</i> is optional for navigating into the service configuration context. Attempting to edit a service with the incorrect <i>customer-id</i> specified will result in an error.</p> <p>More than one VPLS service may be created for a single customer ID.</p> <p>By default, no VPLS instances exist until they are explicitly created.</p> <p>The <b>no</b> form of this command deletes the VPLS service instance with the specified <i>service-id</i>. The service cannot be deleted until all SAPs and SDPs defined within the service ID have been shutdown and deleted, and the service has been shutdown.</p> <p><i>service-id</i> — The unique service identification number identifying the service in the service domain. This ID must be unique to this service and may not be used for any other service of any type. The <i>service-id</i> must be the same number used for every router on which this service is defined.</p> <p><b>Values</b>      service-id: 1 — 214748364                         svc-name: A string up to 64 characters in length.</p> <p><b>customer</b> <i>customer-id</i> — Specifies the customer ID number to be associated with the service. This parameter is required on service creation and optional for service editing or deleting.</p> <p><b>Values</b>      1 — 2147483647</p>

**vpn** *vpn-id* — Specifies the VPN ID number which allows you to identify virtual private networks (VPNs) by a VPN identification number.

**Values** 1 — 2147483647

**Default** null (0)

**m-vpls** — Specifies a managed VPLS.

### service-name

<b>Syntax</b>	<b>service-name</b> <i>service-name</i> <b>no service-name</b>
<b>Context</b>	config>service>epipe config>service>ies config>service>vpls config>service>vpn
<b>Description</b>	This command configures an optional service name, up to 64 characters in length, which adds a name identifier to a given service to then use that service name in configuration references as well as display and use service names in show commands throughout the system. This helps the service provider/administrator to identify and manage services within the 7750 SR and 7450 ESS platforms.  All services are required to assign a service ID to initially create a service. However, either the service ID or the service name can be used to identify and reference a given service once it is initially created.
<b>Parameters</b>	<i>service-name</i> — Specifies a unique service name to identify the service. Service names may not begin with an integer (0-9).

### ies

<b>Syntax</b>	<b>ies</b> <i>service-id</i> <b>customer</b> <i>customer-id</i> <b>vpn</b> <i>vpn-id</i> <b>ies</b> <i>service-id</i> <b>no ies</b> <i>service-id</i>
<b>Context</b>	config>service
<b>Description</b>	This command creates or edits an IES service instance.  The <b>ies</b> command is used to create or maintain an Internet Ethernet Service (IES). If the <i>service-id</i> does not exist, a context for the service is created. If the <i>service-id</i> exists, the context for editing the service is entered.  IES services allow the creation of customer facing IP interfaces in the same routing instance used for service network core routing connectivity. IES services require that the IP addressing scheme used by the subscriber must be unique between it and other addressing schemes used by the provider and potentially the entire Internet.  While IES is part of the routing domain, the usable IP address space may be limited. This allows a portion of the service provider address space to be set aside for service IP provisioning, becoming administered by a separate but subordinate address authority. This feature is defined using the <b>config router service-prefix</b> command.

IP interfaces defined within the context of an IES service ID must have a SAP created as the access point to the subscriber network. This allows a combination of bridging and IP routing for redundancy purposes.

When a service is created, the **customer** keyword and *customer-id* must be specified and associates the service with a customer. The *customer-id* must already exist having been created using the **customer** command in the service context. Once a service has been created with a customer association, it is not possible to edit the customer association. The service must be deleted and recreated with a new customer association.

Once a service is created, the use of the **customer** *customer-id* is optional for navigating into the service configuration context. Attempting to edit a service with the incorrect *customer-id* specified will result in an error.

Multiple IES services are created to separate customer owned IP interfaces. More than one IES service may be created for a single customer ID. More than one IP interface may be created within a single IES service ID. All IP interfaces created within an IES service ID belongs to the same customer.

By default, no IES service instances exist until they are explicitly created.

The **no** form of this command deletes the IES service instance with the specified *service-id*. The service cannot be deleted until all the IP interfaces defined within the service ID have been shutdown and deleted.

<b>Parameters</b>	<i>service-id</i> — The unique service identification number identifying the service in the service domain. This ID must be unique to this service and may not be used for any other service of any type. The <i>service-id</i> must be the same number used for every router on which this service is defined.
	<b>Values</b> service-id: 1 — 214748364 svc-name: A string up to 64 characters in length.
	<b>customer</b> <i>customer-id</i> — Specifies the customer ID number to be associated with the service. This parameter is required on service creation and optional for service editing or deleting.
	<b>Values</b> 1 — 2147483647
	<b>vpn</b> <i>vpn-id</i> — Specifies the VPN ID number which allows you to identify virtual private networks (VPNs) by a VPN identification number.
	<b>Values</b> 1 — 2147483647
	<b>Default</b> null (0)

## vpn

<b>Syntax</b>	<b>vpn</b> <i>service-id</i> [ <b>customer</b> <i>customer-id</i> ] <b>no vpn</b> <i>service-id</i>
<b>Context</b>	config>service
<b>Description</b>	<p>This command creates or edits a Virtu9al Private Routed Network (VPRN) service instance.</p> <p>If the <i>service-id</i> does not exist, a context for the service is created. If the <i>service-id</i> exists, the context for editing the service is entered.</p>

VPRN services allow the creation of customer-facing IP interfaces in the same routing instance used for service network core routing connectivity. VPRN services require that the IP addressing scheme used by the subscriber must be unique between it and other addressing schemes used by the provider and potentially the entire Internet.

IP interfaces defined within the context of an VPRN service ID must have a SAP created as the access point to the subscriber network.

When a service is created, the **customer** keyword and *customer-id* must be specified and associates the service with a customer. The *customer-id* must already exist having been created using the customer command in the service context. When a service is created with a customer association, it is not possible to edit the customer association. The service must be deleted and re-created with a new customer association.

When a service is created, the use of the **customer** *customer-id* is optional to navigate into the service configuration context. If attempting to edit a service with the incorrect *customer-id* results in an error.

Multiple VPRN services are created to separate customer-owned IP interfaces. More than one VPRN service can be created for a single customer ID. More than one IP interface can be created within a single VPRN service ID. All IP interfaces created within an VPRN service ID belongs to the same customer.

The **no** form of the command deletes the VPRN service instance with the specified *service-id*. The service cannot be deleted until all the IP interfaces and all routing protocol configurations defined within the service ID have been shutdown and deleted.

**Default**      None — No VPRN service instances exist until they are explicitly created.

**Parameters**      *service-id* — The unique service identification number identifying the service in the service domain. This ID must be unique to this service and may not be used for any other service of any type. The *service-id* must be the same number used for every 7750 SR on which this service is defined.

**Values**      *service-id:*      1 — 2147483648  
                  *svc-name:*      64 characters maximum

**customer** *customer-id* — Specifies an existing customer identification number to be associated with the service. This parameter is required on service creation and optional for service editing or deleting.

**Values**      1 — 2147483647

sap

**Syntax**      **sap** *sap-id*  
                  **no sap** *sap-id*

**Context**      config>service>vprn>if  
                  config>service>vprn>subscriber-interface>sap

**Description**      This command creates a Service Access Point (SAP) within a service. A SAP is a combination of port and encapsulation parameters which identifies the service access point on the interface and within the router. Each SAP must be unique.



All SAPs must be explicitly created. If no SAPs are created within a service or on an IP interface, a SAP will not exist on that object.

Enter an existing SAP without the **create** keyword to edit SAP parameters. The SAP is owned by the service in which it was created.

A SAP can only be associated with a single service. A SAP can only be defined on a port that has been configured as an access port using the **config interface port-type port-id mode access** command. Channelized TDM ports are always access ports.

If a port is shutdown, all SAPs on that port become operationally down. When a service is shutdown, SAPs for the service are not displayed as operationally down although all traffic traversing the service will be discarded. The operational state of a SAP is relative to the operational state of the port on which the SAP is defined.

The **no** form of this command deletes the SAP with the specified port. When a SAP is deleted, all configuration parameters for the SAP will also be deleted.

**Default** No SAPs are defined.

**Parameters** *sap-id* — Specifies the physical port identifier portion of the SAP definition. See [Common CLI Command Descriptions on page 2061](#) for command syntax.

*port-id* — Specifies the physical port ID in the *slot/mda/port* format.

If the card in the slot has Media Dependent Adapters (MDAs) installed, the *port-id* must be in the *slot\_number/MDA\_number/port\_number* format. For example 6/2/3 specifies port 3 on MDA 2 in slot 6.

The *port-id* must reference a valid port type. When the *port-id* parameter represents SONET/SDH and TDM channels the port ID must include the channel ID. A period “.” separates the physical port from the *channel-id*. The port must be configured as an access port.

If the SONET/SDH port is configured as clear-channel then only the port is specified.

**create** — Keyword used to create a SAP instance.

## disable-aging

**Syntax** **[no] disable-aging**

**Context** config>service>vpls  
config>service>vpls>spoke-sdp  
config>service>vpls>sap

**Description** This command disables MAC address aging across a VPLS service or on a VPLS service SAP or spoke SDP.

Like in a Layer 2 switch, learned MACs can be aged out if no packets are sourced from the MAC address for a period of time (the aging time). In each VPLS service instance, there are independent aging timers for local learned MAC and remote learned MAC entries in the VPLS forwarding database (FDB). The **disable-aging** command turns off aging for local and remote learned MAC addresses.

When **no disable-aging** is specified for a VPLS, it is possible to disable aging for specific SAPs and/or spoke SDPs by entering the **disable-aging** command at the appropriate level.

When the **disable-aging** command is entered at the VPLS level, the **disable-aging** state of individual SAPs or SDPs will be ignored.

The **no** form of this command enables aging on the VPLS service.

**Default** no disable-aging

### disable-learning

**Syntax** **disable-learning**  
**no disable-learning**

**Context** config>service>vpls  
config>service>vpls>sap  
config>service>vpls>spoke-sdp

**Description** This command enables learning of new MAC addresses in the VPLS forwarding database (FDB) for the service instance, SAP instance or spoke SDP instance.

When **disable-learning** is enabled, new source MAC addresses will not be entered in the VPLS service forwarding database. This is true for both local and remote MAC addresses.

When **no disable-learning** is specified for a VPLS, it is possible to disable learning for specific SAPs and/or spoke SDPs by entering the **disable-learning** command at the appropriate level.

When disabled, new source MAC addresses will be learned and entered into the VPLS forwarding database. When the **no disable-learning** command is entered on VPLS level, the **disable-learning** state of individual SAPs or spoke SDPs will be ignored.

This parameter is mainly used in conjunction with the **discard-unknown** command.

The **no** form of this command enables learning of MAC addresses meaning that normal MAC learning is enabled.

**Default** no disable-learning

### discard-unknown

**Syntax** [no] **discard-unknown**

**Context** config>service>vpls

**Description** By default, packets with unknown destination MAC addresses are flooded. If discard-unknown is enabled at the VPLS level, packets with unknown destination MAC address will be dropped instead (even when configured FIB size limits for VPLS or SAP are not yet reached).

The **no** form of this command allows flooding of packets with unknown destination MAC addresses in the VPLS.

**Default** no discard-unknown

## fdb-table-high-wmark

<b>Syntax</b>	<b>[no] fdb-table-high-wmark</b> <i>high-water-mark</i>
<b>Context</b>	config>service>vpls
<b>Description</b>	This command specifies the value to send logs and traps when the threshold is reached.
<b>Default</b>	95
<b>Parameters</b>	<i>high-water-mark</i> — When to send logs and traps.
<b>Values</b>	0 — 100

## fdb-table-low-wmark

<b>Syntax</b>	<b>[no] fdb-table-low-wmark</b> <i>low-water-mark</i>
<b>Context</b>	config>service>vpls
<b>Description</b>	This command specifies the value to send logs and traps when the threshold is reached.
<b>Default</b>	90
<b>Parameters</b>	<i>low-water-mark</i> — When to send logs and traps.
<b>Values</b>	0 — 100

## fdb-table-size

<b>Syntax</b>	<b>fdb-table-size</b> <i>table-size</i> <b>no fdb-table-size</b> [ <i>table-size</i> ]
<b>Context</b>	config>service>vpls
<b>Description</b>	<p>This command specifies the maximum number of MAC entries in the forwarding database (FDB) for the VPLS instance on this node.</p> <p>The <b>fdb-table-size</b> specifies the maximum number of forwarding database entries for both learned and static MAC addresses for the VPLS instance.</p> <p>The <b>no</b> form of this command returns the maxium FDB table size to default.</p>
<b>Default</b>	<b>250</b> — Forwarding table of 250 MAC entries.
<b>Parameters</b>	<i>table-size</i> — The number of entries permitted in the forwarding database for this VPLS instance.

## local-age

<b>Syntax</b>	<b>local-age</b> <i>seconds</i> <b>no local-age</b>
<b>Context</b>	config>service>vpls
<b>Description</b>	<p>Specifies the aging time for locally learned MAC addresses in the forwarding database (FDB) for the Virtual Private LAN Service (VPLS) instance.</p> <p>In a VPLS service, MAC addresses are associated with a Service Access Point (SAP) or with a Service Distribution Point (SDP). MACs associated with a SAP are classified as local MACs, and MACs associated with an SDP are remote MACs.</p> <p>Like in a Layer 2 switch, learned MACs can be aged out if no packets are sourced from the MAC address for a period of time (the aging time). In each VPLS service instance, there are independent aging timers for local learned MAC and remote learned MAC entries in the FDB. The <b>local-age</b> timer specifies the aging time for local learned MAC addresses.</p> <p>The <b>no</b> form of this command returns the local aging timer to the default value.</p>
<b>Default</b>	<b>local age 300</b> — Local MACs aged after 300 seconds.
<b>Parameters</b>	<i>seconds</i> — The aging time for local MACs expressed in seconds. <b>Values</b> 60 — 86400

## remote-age

<b>Syntax</b>	<b>remote-age</b> <i>seconds</i> <b>no remote-age</b>
<b>Context</b>	config>service>vpls
<b>Description</b>	<p>Specifies the aging time for remotely learned MAC addresses in the forwarding database (FDB) for the Virtual Private LAN Service (VPLS) instance.</p> <p>In a VPLS service, MAC addresses are associated with a Service Access Point (SAP) or with a Service Distribution Point (SDP). MACs associated with a SAP are classified as local MACs, and MACs associated with an SDP are remote MACs.</p> <p>Like in a layer 2 switch, learned MACs can be aged out if no packets are sourced from the MAC address for a period of time (the aging time). In each VPLS service instance, there are independent aging timers for local learned MAC and remote learned MAC entries in the FDB. The <b>remote-age</b> timer specifies the aging time for remote learned MAC addresses. To reduce the amount of signaling required between switches configure this timer larger than the <b>local-age</b> timer.</p> <p>The <b>no</b> form of this command returns the remote aging timer to the default value.</p>
<b>Default</b>	<b>remote age 900</b> — Remote MACs aged after 900 seconds.
<b>Parameters</b>	<i>seconds</i> — The aging time for remote MACs expressed in seconds. <b>Values</b> 60 — 86400

## service-mtu

**Syntax** **service-mtu** *octets*  
**no service-mtu**

**Context** config>service>vpls

**Description** This command configures the service payload (Maximum Transmission Unit – MTU), in bytes, for the service. This MTU value overrides the service-type default MTU.

The **service-mtu** defines the payload capabilities of the service. It is used by the system to validate the SAP and SDP binding's operational state within the service.

The service MTU and a SAP's service delineation encapsulation overhead (i.e., 4 bytes for a Dot1q tag) is used to derive the required MTU of the physical port or channel on which the SAP was created. If the required payload is larger than the port or channel MTU, then the SAP will be placed in an inoperative state. If the required MTU is equal to or less than the port or channel MTU, the SAP will be able to transition to the operative state.

When binding an SDP to a service, the service MTU is compared to the path MTU associated with the SDP. The path MTU can be administratively defined in the context of the SDP. The default or administrative path MTU can be dynamically reduced due to the MTU capabilities discovered by the tunneling mechanism of the SDP or the egress interface MTU capabilities based on the next hop in the tunnel path. If the service MTU is larger than the path MTU, the SDP binding for the service will be placed in an inoperative state. If the service MTU is equal to or less than the path MTU, then the SDP binding will be placed in an operational state.

In the event that a service MTU, port or channel MTU, or path MTU is dynamically or administratively modified, then all associated SAP and SDP binding operational states are automatically re-evaluated.

The **no** form of this command returns the default **service-mtu** for the indicated service type to the default value.

**Default** VPLS: 1514

The following table displays MTU values for specific VC types.

VC-Type	Example Service MTU	Advertised MTU
Ethernet	1514	1500
Ethernet (with preserved dot1q)	1518	1504
VPLS	1514	1500
VPLS (with preserved dot1q)	1518	1504
VLAN (dot1p transparent to MTU value)	1514	1500
VLAN (QinQ with preserved bottom Qtag)	1518	1504

*octets* — The size of the MTU in octets, expressed as a decimal integer.

**Values** 1 — 9194

## split-horizon-group

<b>Syntax</b>	<b>[no] split-horizon-group</b> <i>[group-name]</i> <i>[residential-group]</i>
<b>Context</b>	config>service>vpls
<b>Description</b>	<p>This command creates a new split horizon group for the VPLS instance. Traffic arriving on a SAP or spoke SDP within this split horizon group will not be copied to other SAPs or spoke SDPs in the same split horizon group.</p> <p>A split horizon group must be created before SAPs and spoke SDPs can be assigned to the group.</p> <p>The split horizon group is defined within the context of a single VPLS. The same group-name can be re-used in different VPLS instances.</p> <p>Up to 30 split horizon groups can be defined per VPLS instance.</p> <p>The <b>no</b> form of the command removes the group name from the configuration.</p>
<b>Parameters</b>	<p><i>group-name</i> — Specifies the name of the split horizon group to which the SDP belongs.</p> <p><i>residential-group</i> — Defines a split horizon group as a residential split horizon group (RSHG). Doing so entails that:</p> <ul style="list-style-type: none"> <li>a) SAPs which are members of this Residential Split Horizon Group will have: <ul style="list-style-type: none"> <li>– Double-pass queuing at ingress as default setting (can be disabled)</li> <li>– STP disabled (can <u>not</u> be enabled)</li> <li>– ARP reply agent enabled by default (can be disabled)</li> <li>– MAC pinning enabled by default (can be disabled)</li> <li>– broadcast packets are discarded, blocking unknown, flooded traffic</li> </ul> </li> <li>b) Spoke SDPs which are members of this Residential Split Horizon Group will have: <ul style="list-style-type: none"> <li>– Downstream multicast traffic supported</li> <li>– Double-pass queuing is not applicable</li> <li>– STP is disabled (can be enabled)</li> <li>– ARP reply agent is not applicable (dhcp-lease-states are not supported on spoke SDPs)</li> <li>– MAC pinning enabled per default (can be disabled)</li> </ul> </li> </ul>
<b>Default</b>	A split horizon group is by default not created as a residential-group.

## sap

<b>Syntax</b>	<b>sap</b> <i>sap-id</i> <b>[split-horizon-group</b> <i>group-name</i> <b>]</b> <b>[create]</b> <b>[capture-sap]</b> <b>no sap</b> <i>sap-id</i>
<b>Context</b>	config>service>vpls
<b>Description</b>	This command creates a Service Access Point (SAP) within a service. A SAP is a combination of port and encapsulation parameters which identifies the service access point on the interface and within the

7450. Each SAP must be unique. All SAPs must be explicitly created. If no SAPs are created within a service or on an IP interface, a SAP will not exist on that object.

Enter an existing SAP without the **create** keyword to edit SAP parameters. The SAP is owned by the service in which it was created.

A SAP can only be associated with a single service. A SAP can only be defined on a port that has been configured as an access port using the **config interface port-type port-id mode access** command.

If a port is shutdown, all SAPs on that port become operationally down. When a service is shutdown, SAPs for the service are not displayed as operationally down although all traffic traversing the service will be discarded. The operational state of a SAP is relative to the operational state of the port on which the SAP is defined.

The **no** form of this command deletes the SAP with the specified port. When a SAP is deleted, all configuration parameters for the SAP will also be deleted. For Internet Ethernet Service (IES), the IP interface must be shutdown before the SAP on that interface may be removed.

<b>Default</b>	No SAPs are defined.
<b>Special Cases</b>	<p>A VPLS SAP can be defined with Ethernet or SONET/SDH ports. ESS-Series ESS-Series</p> <p>A default SAP has the following format: port-id:*. This type of SAP is supported only on Ethernet MDAs and its creation is allowed only in the scope of Layer 2 services (Epipe and VPLS). This type of SAP is mutually exclusive with a SAP defined by explicit null encapsulation (for example, 1/1/1:0).</p>
<b>Parameters</b>	<p><i>sap-id</i> — Specifies the physical port identifier portion of the SAP definition. See <a href="#">Common Service Commands on page 2062</a> for <i>sap-id</i> command syntax.</p> <p><i>port-id</i> — Specifies the physical port ID in the <i>slot/mda/port</i> format.</p> <p>If the card in the slot has Media Dependent Adapters (MDAs) installed, the <i>port-id</i> must be in the slot_number/MDA_number/port_number format. For example 6/2/3 specifies port 3 on MDA 2 in slot 6.</p> <p>The <i>port-id</i> must reference a valid port type. When the <i>port-id</i> parameter represents SONET/SDH, the port ID must include the channel ID. A period “.” separates the physical port from the <i>channel-id</i>. The port must be configured as an access port.</p> <p>If the SONET/SDH port is configured as clear-channel then only the port is specified.</p> <p><b>split-horizon-group</b> <i>group-name</i> — Specifies an existing split horizon group name.</p> <p><b>capture-sap</b> — Keyword to create a capture SAP.</p> <p><b>create</b> — Keyword used to create a SAP instance. The <b>create</b> keyword requirement can be enabled/disabled in the <b>environment&gt;create</b> context.</p>

## accounting-policy

<b>Syntax</b>	<b>accounting-policy</b> <i>acct-policy-id</i> <b>no accounting-policy</b>
<b>Context</b>	config>service>vpls>sap config>service>vpls>spoke-sdp

```
config>service>vpls>mesh-sdp
config>service>ies>if>sap
```

<b>Description</b>	<p>This command creates the accounting policy context that can be applied to a SAP or SDP.</p> <p>An accounting policy must be defined before it can be associated with a SAP or SDP. If the <i>policy-id</i> does not exist, an error message is generated.</p> <p>A maximum of one accounting policy can be associated with a SAP or SDP at one time. Accounting policies are configured in the <b>config&gt;log</b> context.</p> <p>The <b>no</b> form of this command removes the accounting policy association from the SAP or SDP, and the accounting policy reverts to the default.</p>
<b>Default</b>	Default accounting policy.
<b>Parameters</b>	<i>acct-policy-id</i> — Enter the accounting <i>policy-id</i> as configured in the <b>config&gt;log&gt;accounting-policy</b> context.
<b>Values</b>	1 — 99

## collect-stats

<b>Syntax</b>	<b>[no] collect-stats</b>
<b>Context</b>	<pre>config&gt;service&gt;vpls&gt;sap config&gt;service&gt;vpls&gt;spoke-sdp config&gt;service&gt;vpls&gt;mesh-sdp config&gt;service&gt;ies&gt;if&gt;sap</pre>
<b>Description</b>	<p>This command enables accounting and statistical data collection for either the SAP or SDP, network port, or IP interface. When applying accounting policies the data, by default, is collected in the appropriate records and written to the designated billing file.</p> <p>When the <b>no collect-stats</b> command is issued the statistics are still accumulated by the IOM cards. However, the CPU will not obtain the results and write them to the billing file. If a subsequent <b>collect-stats</b> command is issued then the counters written to the billing file include all the traffic while the <b>no collect-stats</b> command was in effect.</p>
<b>Default</b>	collect-stats

## limit-mac-move

<b>Syntax</b>	<b>limit-mac-move [blockable   non-blockable]</b> <b>no limit-mac-move</b>
<b>Context</b>	<pre>config&gt;service&gt;vpls&gt;sap config&gt;service&gt;vpls&gt;spoke-sdp</pre>
<b>Description</b>	<p>This command indicates whether or not the mac-move agent, when enabled using <b>config&gt;service&gt;vpls&gt;mac-move</b> or <b>config&gt;service&gt;epipe&gt;mac-move</b>, will limit the MAC re-learn (move) rate on this SAP.</p>



<b>Default</b>	<b>SAPs and spoke SDPs are blockable</b>
<b>Parameters</b>	<p><b>blockable</b> — the agent will monitor the MAC re-learn rate on the SAP, and it will block it when the re-learn rate is exceeded.</p> <p><b>non-blockable</b> — When specified, this SAP will not be blocked, and another blockable SAP will be blocked instead.</p>

## mac-pinning

<b>Syntax</b>	<b>[no] mac-pinning</b>
<b>Context</b>	<pre>config&gt;service&gt;vpls&gt;sap config&gt;service&gt;vpls&gt;spoke-sdp config&gt;service&gt;vpls&gt;mesh-sdp</pre>
<b>Description</b>	<p>Enabling this command will disable re-learning of MAC addresses on other SAPs within the VPLS. The MAC address will remain attached to a given SAP for the duration of its age-timer.</p> <p>The age of the MAC address entry in the FIB is set by the age timer. If <b>mac-aging</b> is disabled on a given VPLS service, any MAC address learned on a SAP/SDP with <b>mac-pinning</b> enabled will remain in the FIB on this SAP/SDP forever.</p> <p>Every event that would otherwise result in re-learning will be logged (MAC address, original-SAP, new-SAP).</p> <p>Note that MAC addresses learned during DHCP address assignment (DHCP snooping enabled) are not impacted by this command. MAC-pinning for such addresses is implicit.</p>
<b>Default</b>	When a SAP or spoke SDP is part of a Residential Split Horizon Group (RSHG), MAC pinning is activated at creation of the SAP. Otherwise MAC pinning is not enabled by default.

## managed-vlan-list

<b>Syntax</b>	<b>managed-vlan-list</b>
<b>Context</b>	<pre>config&gt;service&gt;vpls&gt;sap</pre>
<b>Description</b>	<p>This command enters the context for configuring VLAN ranges to be managed by a management VPLS. The list indicates, for each SAP, the ranges of associated VLANs that will be affected when the SAP changes state.</p> <p>This command is only valid when the VPLS in which it is entered was created as a management VPLS.</p>

## range

<b>Syntax</b>	<b>[no] range <i>vlan-range</i></b>
<b>Context</b>	<pre>config&gt;service&gt;vpls&gt;sap&gt;managed-vlan-list</pre>

<b>Description</b>	<p>This command configures a range of VLANs on an access port that are to be managed by an existing management VPLS.</p> <p>This command is only valid when the VPLS in which it is entered was created as a management VPLS, and when the SAP in which it was entered was created on an Ethernet port with encapsulation type of dot1q or qinq, or on a Sonet/SDH port with encapsulation type of bcp-dot1q.</p> <p>To modify the range of VLANs, first the new range should be entered and afterwards the old range removed.</p>
<b>Default</b>	None
<b>Parameters</b>	<i>vlan-range</i> — Specify the VLAN start value and VLAN end value. The end-vlan must be greater than start-vlan. The format is <start-vlan>-<end-vlan>
<b>Values</b>	start-vlan: 1 — 4094 end-vlan: 1 — 4094

---

## Service Billing Commands

### authentication-policy

<b>Syntax</b>	<b>authentication-policy</b> <i>name</i> <b>no authentication-policy</b>
<b>Context</b>	config>service>vppls>sap config>service>vprn>sub-if>grp-if
<b>Description</b>	This command defines which subscriber authentication policy must be applied when a DHCP message is received on the interface. The authentication policies must already be defined. The policy will only be applied when DHCP snooping is enabled on the SAP on Layer 2 interfaces.
<b>Parameters</b>	<i>name</i> — Specifies a unique authentication policy name.

### root-guard

<b>Syntax</b>	<b>[no] root-guard</b>
<b>Context</b>	config>service>vppls>sap>stp
<b>Description</b>	This command specifies whether this port is allowed to become an STP root port. It corresponds to the restrictedRole parameter in 802.1Q. If set, it can cause lack of spanning tree connectivity.
<b>Default</b>	no root-guard

## SAP Subscriber Management Commands

Subscriber management commands are also described in the [Triple Play Services Command Reference on page 65](#) section.

### sub-sla-mgmt

<b>Syntax</b>	<b>[no] sub-sla-mgmt</b>
<b>Context</b>	config>service>vpls>sap config>service>ies>if>sap
<b>Description</b>	This command entables the context to configure subscriber management parameters for this SAP.
<b>Default</b>	no sub-sla-mgmt

### def-sla-profile

<b>Syntax</b>	<b>def-sla-profile</b> <i>default-sla-profile-name</i> <b>no def-sla-profile</b>
<b>Context</b>	config>service>vpls>sap>sub-sla-mgmt
<b>Description</b>	<p>This command specifies a default SLA profile for this SAP. The SLA profile must be defined prior to associating the profile with a SAP in the <b>config&gt;subscriber-mgmt&gt;sla-profile</b> context.</p> <p>An SLA profile is a named group of QoS parameters used to define per service QoS for all subscriber hosts common to the same subscriber within a provider service offering. A single SLA profile may define the QoS parameters for multiple subscriber hosts. SLA profiles are maintained in two locations, the subscriber identification policy and the subscriber profile templates. After a subscriber host is associated with an SLA profile name, either the subscriber identification policy used to identify the subscriber or the subscriber profile associated with the subscriber host must contain an SLA profile with that name. If both the subscriber identification policy and the subscriber profile contain the SLA profile name, the SLA profile in the subscriber profile is used.</p> <p>The <b>no</b> form of the command removes the default SLA profile from the SAP configuration.</p>
<b>Default</b>	no def-sla-profile
<b>Parameters</b>	<i>default-sla-profile-name</i> — Specifies a default SLA profile for this SAP. The SLA profile must be defined prior to associating the profile with a SAP in the <b>config&gt;subscriber-mgmt&gt;sla-profile</b> context.

## def-sub-profile

<b>Syntax</b>	<b>def-sub-profile</b> <i>default-subscriber-profile-name</i>
<b>Context</b>	config>service>vpls>sap>sub-sla-mgmt
<b>Description</b>	<p>This command specifies a default subscriber profile for this SAP. The subscriber profile must be defined prior to associating the profile with a SAP in the <b>config&gt;subscriber-mgmt&gt;sub-profile</b> context.</p> <p>A subscriber profile defines the aggregate QoS for all hosts within a subscriber context. This is done through the definition of the egress and ingress scheduler policies that govern the aggregate SLA for subscriber using the subscriber profile. Subscriber profiles also allow for specific SLA profile definitions when the default definitions from the subscriber identification policy must be overridden.</p> <p>The <b>no</b> form of the command removes the default SLA profile from the SAP configuration.</p>
<b>Parameters</b>	<i>default-sub-profile</i> — Specifies a default subscriber profile for this SAP. The subscriber profile must be defined prior to associating the profile with a SAP in the <b>config&gt;subscriber-mgmt&gt;sub-profile</b> context.

## sub-ident-policy

<b>Syntax</b>	<b>sub-ident-policy</b> <i>sub-ident-policy-name</i>
<b>Context</b>	config>service>vpls>sap>sub-sla-mgmt
<b>Description</b>	<p>This command associates a subscriber identification policy to this SAP. The subscriber identification policy must be defined prior to associating the profile with a SAP in the <b>config&gt;subscriber-mgmt&gt;sub-ident-policy</b> context.</p> <p>Subscribers are managed by the system through the use of subscriber identification strings. A subscriber identification string uniquely identifies a subscriber. For static hosts, the subscriber identification string is explicitly defined with each static subscriber host.</p> <p>For dynamic hosts, the subscriber identification string must be derived from the DHCP ACK message sent to the subscriber host. The default value for the string is the content of Option 82 CIRCUIT-ID and REMOTE-ID fields interpreted as an octet string. As an option, the DHCP ACK message may be processed by a subscriber identification policy which has the capability to parse the message into an alternative ASCII or octet string value.</p> <p>When multiple hosts on the same port are associated with the same subscriber identification string they are considered to be host members of the same subscriber.</p> <p>The <b>no</b> form of the command removes the default subscriber identification policy from the SAP configuration.</p>
<b>Default</b>	no sub-ident-policy
<b>Parameters</b>	<i>sub-ident-policy-name</i> — Specifies a subscriber identification policy for this SAP. The subscriber profile must be defined prior to associating the profile with a SAP in the <b>config&gt;subscriber-mgmt&gt;sub-ident-policy</b> context.

**subscriber** *sub-ident-string* — Specifies a subscriber identification profile to be associated with the static subscriber host. The subscriber information is used by the SAP arp-reply-agent to determine the proper handling of received ARP requests from subscribers.

- For VPLS SAPs with **arp-reply-agent** enabled with the optional *sub-ident* parameter, the static subscriber host's *sub-ident-string* is used to determine whether an ARP request received on the SAP is sourced from a host belonging to the same subscriber as the destination host. When both the destination and source hosts from the ARP request are known on the SAP and the subscriber identifications do not match, the ARP request may be forwarded to the rest of the destinations.

If the static subscriber host's *sub-ident* string is not defined, the host is not considered to belong to the same subscriber as another host on the SAP.

If source or destination host is unknown, the hosts are not considered to belong to the same subscriber. (ARP messages from unknown hosts are subject to anti-spoof filtering rules applied at the SAP.)

If *sub-ident* is not enabled on the SAP arp-reply-agent, subscriber identification matching is not performed on ARP requests received on the SAP.

ARP requests are never forwarded back to the same SAP or within the receiving SAP's Split Horizon Group.

## profiled-traffic-only

<b>Syntax</b>	<b>[no] profiled-traffic-only</b>
<b>Context</b>	config>service>vpls>sap>sub-sla-mgmt
<b>Description</b>	<p>This command enables profiled traffic only for this SAP. The profiled traffic refers to single subscriber traffic on a dedicated SAP (in the VLAN-per-subscriber model). When enabled, subscriber queues are instantiated through the QOS policy defined in the sla-profile and the associated SAP queues are deleted. This can increase subscriber scaling by reducing the number of queues instantiated per subscriber (in the VLAN-per-subscriber model). In order for this to be achieved, any configured multi-sub-sap limit must be removed (leaving the default of 1).</p> <p>The <b>no</b> form of the command disables the command.</p>

---

## Multicast Commands

### fast-leave

<b>Syntax</b>	<b>[no] fast-leave</b>
<b>Context</b>	<pre>config&gt;service&gt;vpls&gt;sap&gt;snooping config&gt;service&gt;vpls&gt;spoke-sdp&gt;snooping config&gt;service&gt;vpls&gt;mesh-sdp&gt;snooping config&gt;service&gt;vpls&gt;sap&gt;mld-snooping config&gt;service&gt;vpls&gt;spoke-sdp&gt;mld-snooping config&gt;service&gt;vpls&gt;mesh-sdp&gt;mld-snooping</pre>
<b>Description</b>	<p>This command enables fast leave.</p> <p>When IGMP fast leave processing is enabled, the ESS-Series will immediately remove a SAP or SDP from the IP multicast group when it detects an IGMP 'leave' on that SAP or SDP. Fast leave processing allows the switch to remove a SAP or SDP that sends a 'leave' from the forwarding table without first sending out group-specific queries to the SAP or SDP, and thus speeds up the process of changing channels ('zapping').</p> <p>Fast leave should only be enabled when there is a single receiver present on the SAP or SDP.</p> <p>When fast leave is enabled, the configured last-member-query-interval value is ignored.</p>
<b>Default</b>	no fast-leave

### from-vpls

<b>Syntax</b>	<b>from-vpls</b> <i>vpls-id</i> <b>no from-vpls</b>
<b>Context</b>	<pre>config&gt;service&gt;vpls&gt;sap&gt;snooping&gt;mvr config&gt;service&gt;vpls&gt;sap&gt;mld-snooping&gt;mvr</pre>
<b>Description</b>	<p>This command configures the VPLS from which multicast traffic is copied upon receipt of an IGMP join request.</p> <p>IGMP snooping must be enabled on the MVR VPLS.</p>
<b>Default</b>	no from-vpls
<b>Parameters</b>	<i>vpls-id</i> — Specifies the MVR VPLS from which multicast channels should be copied into this SAP.

## group

<b>Syntax</b>	<b>[no] group <i>grp-address</i></b>
<b>Context</b>	config>service>vpls>sap>snooping>static config>service>vpls>spoke-sdp>snooping>static config>service>vpls>mesh-sdp>snooping>static
<b>Description</b>	This command adds a static multicast group either as a (*, g) or as one or more (s,g) records. When a static IGMP group is added, multicast data for that (*,g) or (s,g) is forwarded to the specific SAP or SDP without receiving any membership report from a host.
<b>Default</b>	none
<b>Parameters</b>	<i>grp-address</i> — Specifies an IGMP multicast group address that receives data on an interface. The IP address must be unique for each static group.

## group-policy

<b>Syntax</b>	<b>group-policy <i>policy-name</i></b> <b>no group-policy</b>
<b>Context</b>	config>service>vpls>snooping>mvr config>service>vpls>mld-snooping>mvr
<b>Description</b>	Identifies filter policy of multicast groups to be applied to this MVR VPLS. The sources of the multicast traffic must be a member of the MVR VPLS  The <b>no</b> form of the command removes the MVR policy association from the VPLS.
<b>Default</b>	no import policy is specified.
<b>Parameters</b>	<i>policy-name</i> — The routing policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters excluding double quotes. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes. Routing policies are configured in the config>router>policy-options context. The router policy must be defined before it can be imported. For details on IGMP policies, see <a href="#">Enabling IGMP Group Membership Report Filtering on page 594</a> .

## igmp-snooping

<b>Syntax</b>	<b>igmp-snooping</b>
<b>Context</b>	config>service>vpls config>service>vpls>sap config>service>vpls>spoke-sdp config>service>vpls>mesh-sdp
<b>Description</b>	This command enables the Internet Group Management Protocol (IGMP) snooping context.
<b>Default</b>	none



## mld-snooping

<b>Syntax</b>	<b>mld-snooping</b>
<b>Context</b>	config>service>vpls config>service>vpls>sap
<b>Description</b>	This command configures MLD snooping parameters.

## import

<b>Syntax</b>	<b>import</b> <i>policy-name</i> <b>no import</b>
<b>Context</b>	config>service>vpls>sap>snooping config>service>vpls>spoke-sdp>snooping config>service>vpls>mesh-sdp>snooping config>service>vpls>sap>mld-snooping config>service>vpls>spoke-sdp>mld-snooping config>service>vpls>mesh-sdp>mld-snooping
<b>Description</b>	This command specifies the import routing policy to be used for IGMP packets to be used on this SAP or SDP. Only a single policy can be imported on a single SAP at any time.  The <b>no</b> form of the command removes the policy association from the SAP or SDP.
<b>Default</b>	no import (No import policy is specified)
<b>Parameters</b>	<i>policy-name</i> — The routing policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters excluding double quotes. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes. Routing policies are configured in the config>router>policy-options context. The router policy must be defined before it can be imported.

## last-member-query-interval

<b>Syntax</b>	<b>last-member-query-interval</b> <i>tenths-of-seconds</i> <b>no last-member-query-interval</b>
<b>Context</b>	config>service>vpls>sap>snooping config>service>vpls>spoke-sdp>snooping config>service>vpls>mesh-sdp>snooping config>service>vpls>sap>mld-snooping config>service>vpls>spoke-sdp>mld-snooping config>service>vpls>mesh-sdp>mld-snooping
<b>Description</b>	This command configures the maximum response time used in group-specific queries sent in response to 'leave' messages, and is also the amount of time between 2 consecutive group-specific queries. This value may be tuned to modify the leave latency of the network. A reduced value results in reduced time to detect the loss of the last member of a group.

The configured last-member-query-interval is ignored when fast-leave is enabled on the SAP or SDP.

<b>Default</b>	10
<b>Parameters</b>	<i>seconds</i> — Specifies the frequency, in tenths of seconds, at which query messages are sent.
<b>Values</b>	1 — 50

### max-num-groups

<b>Syntax</b>	<b>max-num-groups</b> <i>count</i> <b>no max-num-groups</b>
<b>Context</b>	config>service>vpls>sap>snooping config>service>vpls>spoke-sdp>snooping config>service>vpls>mesh-sdp>snooping config>service>vpls>sap>mld-snooping config>service>vpls>spoke-sdp>mld-snooping config>service>vpls>mesh-sdp>mld-snooping
<b>Description</b>	This command defines the maximum number of multicast groups that can be joined on this SAP or SDP. If the ESS-Series receives an IGMP join message that would exceed the configured number of groups, the request is ignored.
<b>Default</b>	no max-num-groups
<b>Parameters</b>	<i>count</i> — Specifies the maximum number of groups that can be joined on this SAP or SDP.
<b>Values</b>	1 — 1000

### mcac

<b>Syntax</b>	<b>mcac</b>
<b>Context</b>	config>service>pw-template>igmp-snooping config>service>vpls>mesh-sdp>snooping
<b>Description</b>	This command configures multicast CAC policy and constraints for this interface.
<b>Default</b>	none

### policy

<b>Syntax</b>	<b>policy</b> <i>policy-name</i> <b>no policy</b>
<b>Context</b>	config>service>pw-template>igmp-snooping>mcac config>service>vpls>mesh-sdp>snooping>mcac
<b>Description</b>	This command configures the multicast CAC policy name.

**Parameters** *policy-name* — The multicast CAC policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

## unconstrained-bw

**Syntax** **unconstrained-bw** *bandwidth* **mandatory-bw** *mandatory-bw*  
**no unconstrained-bw**

**Context** config>service>vpls>mesh-sdp>snooping>mcac  
 config>service>vpls>spoke-sdp>snooping>mcac  
 config>service>vpls>sap>igmp-snooping>mcac

**Description** This command configures the bandwidth for the interface's multicast CAC policy traffic. When disabled (**no unconstrained-bw**) there will be no checking of bandwidth constraints on the interface level. When enabled and a policy is defined, enforcement is performed. The allocated bandwidth for optional channels should not exceed the **unconstrained-bw** minus the **mandatory-bw** and the mandatory channels have to stay below the specified value for the **mandatory-bw**. After this interface check, the bundle checks are performed.

**Parameters** *bandwidth* — The bandwidth assigned for interface's MCAC policy traffic, in kilo-bits per second (kbps).

**Values** 0 — 2147483647

**mandatory-bw** *mandatory-bw* — Specifies the bandwidth pre-reserved for all the mandatory channels on a given interface in kilo-bits per second (kbps).

If the *bandwidth* value is 0, no mandatory channels are allowed. If *bandwidth* is not configured, then all mandatory and optional channels are allowed.

If the value of *mandatory-bw* is equal to the value of *bandwidth*, then all the unconstrained bandwidth on a given interface is allocated to mandatory channels configured through multicast CAC policy on that interface and no optional groups (channels) are allowed.

The value of *mandatory-bw* should always be less than or equal to that of *bandwidth*. An attempt to set the value of *mandatory-bw* greater than that of *bandwidth*, will result in inconsistent value error.

**Values** 0 — 2147483647

## mrrouter-port

**Syntax** [**no**] **mrrouter-port**

**Context** config>service>vpls>sap>snooping  
 config>service>vpls>spoke-sdp>snooping  
 config>service>vpls>mesh-sdp>snooping

**Description** This command specifies whether a multicast router is attached behind this SAP or SDP.

Configuring a SAP as an mrouter-port will have a double effect. Firstly, all multicast traffic received on another SAP or SDP will be copied to this SAP or SDP. Secondly, IGMP reports generated by the system as a result of someone joining or leaving a multicast group, will be sent to this SAP or SDP.

If two multicast routers exist in the network, one of them will become the active querier. While the other multicast router (non-querier) stops sending IGMP queries, it should still receive reports to keep its multicast trees up to date. To support this, the mrouter-port should be enabled on all SAPs or SDPs connecting to a multicast router.

Note that the IGMP version to be used for the reports (v1, v2 or v3) can only be determined after an initial query has been received. Until such time no reports are sent on the SAP or spoke SDP, even if mrouter-port is enabled.

If the **send-queries** command is enabled on this SAP or spoke SDP, the **mrouter-port** parameter can not be set.

**Default** no mrouter-port

### mvr

**Syntax** mvr

**Context** config>service>vpls>snooping  
config>service>vpls>mld-snooping  
config>service>vpls>sap>snooping

**Description** This command enables the context to configure Multicast VPLS Registration (MVR) parameters.

### query-interval

**Syntax** query-interval *seconds*  
no query-interval

**Context** config>service>vpls>snooping  
config>service>vpls>sap>snooping  
config>service>vpls>spoke-sdp>snooping  
config>service>vpls>mesh-sdp>snooping  
config>service>vpls>mld-snooping  
config>service>vpls>sap>mld-snooping  
config>service>vpls>spoke-sdp>mld-snooping  
config>service>vpls>mesh-sdp>mld-snooping

**Description** If the **send-queries** command is enabled, this parameter specifies the interval between two consecutive general queries sent by the system on this SAP or SDP.

The configured query-interval must be greater than the configured query-response-interval.

If send-queries is not enabled on this SAP or SDP, the configured query-interval value is ignored.

**Default** 125

**Parameters** *seconds* — The time interval, in seconds, that the router transmits general host-query messages.

**Values** 2 — 1024

## query-src-ip

<b>Syntax</b>	<b>query-src-ip</b> <i>ipv6-address</i> <b>no query-src-ip</b>
<b>Context</b>	config>service>vpls>mld-snooping config>service>vpls>igmp-snooping
<b>Description</b>	This command configures the IP source address used in IGMP or MLD queries.

## query-response-interval

<b>Syntax</b>	<b>query-response-interval</b> <i>seconds</i>
<b>Context</b>	config>service>vpls>sap>snooping config>service>vpls>spoke-sdp>snooping config>service>vpls>mesh-sdp>snooping config>service>vpls>sap>mld-snooping config>service>vpls>spoke-sdp>mld-snooping config>service>vpls>mesh-sdp>mld-snooping
<b>Description</b>	If the <b>send-queries</b> command is enabled, this parameter specifies the maximum response time advertised in IGMPv2/v3 queries.  The configured query-response-interval must be smaller than the configured query-interval.  If send-queries is not enabled on this SAP or SDP, the configured query-response-interval value is ignored.
<b>Default</b>	10
<b>Parameters</b>	<i>seconds</i> — Specifies the length of time to wait to receive a response to the host-query message from the host.
<b>Values</b>	1 — 1023

## report-src-ip

<b>Syntax</b>	<b>report-src-ip</b> <i>ip-address</i> <b>no report-src-ip</b>
<b>Context</b>	config>service>vpls>igmp-snooping
<b>Description</b>	This parameters specifies the source IP address used when generating IGMP reports. According the IGMPv3 standard, a zero source address is allowed in sending IGMP reports. However, for interoperability with some multicast routers, the source IP address of IGMP group reports can be configured using this command.

<b>Default</b>	0.0.0.0
<b>Parameters</b>	<i>ip-address</i> — The source IP source address in transmitted IGMP reports.

## robust-count

<b>Syntax</b>	<b>robust-count</b> <i>robust-count</i> <b>no robust-count</b>
<b>Context</b>	config>service>vpls>snooping config>service>vpls>sap>snooping config>service>vpls>spoke-sdp>snooping config>service>vpls>mesh-sdp>snooping config>service>vpls>sap>mld-snooping config>service>vpls>spoke-sdp>mld-snooping config>service>vpls>mesh-sdp>mld-snooping
<b>Description</b>	If the <b>send-queries</b> command is enabled, this parameter allows tuning for the expected packet loss on a SAP or SDP. The robust-count variable allows tuning for the expected packet loss on a subnet and is comparable to a retry count. If this SAP or SDP is expected to be 'lossy', this parameter may be increased. IGMP snooping on this SAP or SDP is robust to (robust-count-1) packet losses.  If send-queries is not enabled, this parameter will be ignored.
<b>Default</b>	2
<b>Parameters</b>	<i>robust-count</i> — Specifies the robust count for the SAP or SDP.
<b>Values</b>	2 — 7

## send-queries

<b>Syntax</b>	<b>[no] send-queries</b>
<b>Context</b>	config>service>vpls>sap>snooping config>service>vpls>spoke-sdp>snooping config>service>vpls>mesh-sdp>snooping config>service>vpls>sap>mld-snooping config>service>vpls>spoke-sdp>mld-snooping config>service>vpls>mesh-sdp>mld-snooping
<b>Description</b>	This command specifies whether to send IGMP general query messages on the SAP or SDP.  If mrouter-port is enabled on this SAP or spoke SDP, the <b>send-queries</b> command parameter can not be set.
<b>Default</b>	no send-queries

## source

**[no] source** *ip-address*

<b>Context</b>	config>service>vpls>sap>snooping>static>group config>service>vpls>spoke-sdp>snooping>static>group config>service>vpls>mesh-sdp>snooping>static>group
<b>Description</b>	This command adds a static (s,g) entry to allow multicast traffic for the corresponding multicast group from that specific source. For the same multicast group, more than one source can be specified.  Static (s,g) entries can not be entered when a starg is already created.  Use the no form of the command to remove the source from the configuration.
<b>Default</b>	none
<b>Parameters</b>	<i>ip-address</i> — Specifies the IPv4 unicast address.

## starg

<b>Syntax</b>	<b>[no] starg</b>
<b>Context</b>	config>service>vpls>sap>snooping>static>group config>service>vpls>spoke-sdp>snooping>static>group config>service>vpls>mesh-sdp>snooping>static>group
<b>Description</b>	This command adds a static (*,g) entry to allow multicast traffic for the corresponding multicast group from any source. This command can only be enabled if no existing source addresses for this group are specified.  Use the <b>no</b> form of the command to remove the starg entry from the configuration.
<b>Default</b>	no starg

## static

<b>Syntax</b>	<b>static</b>
<b>Context</b>	config>service>vpls>sap>snooping config>service>vpls>spoke-sdp>snooping config>service>vpls>mesh-sdp>snooping config>service>vpls>sap>mld-snooping config>service>vpls>spoke-sdp>mld-snooping config>service>vpls>mesh-sdp>mld-snooping
<b>Description</b>	This command enables access to the context to configure static group addresses. Static group addresses can be configured on a SAP or SDP. When present either as a (*, g) or a (s,g) entry, multicast packets matching the configuration will be forwarded even if no join message was registered for the specific group.
<b>Default</b>	none

### to-sap

<b>Syntax</b>	<b>to-sap</b> <i>sap-id</i> <b>no to-sap</b>
<b>Context</b>	config>service>vpls>sap>snooping>mvr
<b>Description</b>	<p>In some situations, the multicast traffic should not be copied from the MVR VPLS to the SAP on which the IGMP message was received (standard MVR behaviour) but to another SAP.</p> <p>This command configures the SAP to which the multicast data needs to be copied.</p>
<b>Default</b>	no to-sap
<b>Parameters</b>	<i>sap-id</i> — Specifies the SAP to which multicast channels should be copied. See <a href="#">Common Service Commands on page 2062</a> for <i>sap-id</i> command syntax.



---

## DHCP and Anti-Spoofing Commands

### anti-spoof

<b>Syntax</b>	<b>anti-spoof {ip   mac   ip-mac}</b> <b>no anti-spoof</b>
<b>Context</b>	config>service>vpls>sap config>service>ies>sap
<b>Description</b>	<p>This command enables anti-spoof filtering and optionally changes the anti-spoof matching type for the SAP.</p> <p>The type of anti-spoof filtering defines what information in the incoming packet is used to generate the criteria to lookup an entry in the anti-spoof filter table. The type parameter (<b>ip</b>, <b>mac</b>, <b>ip-mac</b>) defines the anti-spoof filter type enforced by the SAP when anti-spoof filtering is enabled.</p> <p>The <b>no</b> form of the command disables anti-spoof filtering on the SAP.</p>
<b>Default</b>	<b>no anti-spoof</b>
<b>Parameters</b>	<p><b>ip</b> — Configures SAP anti-spoof filtering to use only the source IP address in its lookup. If a static host exists on the SAP without an IP address specified, the <b>anti-spoof ip</b> command will fail.</p> <p><b>mac</b> — Configures SAP anti-spoof filtering to use only the source MAC address in its lookup. If a static host exists on the SAP without a specified MAC address, the <b>anti-spoof mac</b> command will fail.</p> <p><b>ip-mac</b> — Configures SAP anti-spoof filtering to use both the source IP address and the source MAC address in its lookup. If a static host exists on the SAP without both the IP address and MAC address specified, the <b>anti-spoof ip-mac</b> command will fail.</p>

### arp-host

<b>Syntax</b>	<b>arp-host</b>
<b>Context</b>	config>service>vpls>sap
<b>Description</b>	This command enables the context to configure ARP host parameters.

### diameter-auth-policy

<b>Syntax</b>	<b>diameter-auth-policy <i>name</i></b> <b>no diameter-auth-policy</b>
<b>Context</b>	config>service>vpls>sap
<b>Description</b>	This command is used to configure the Diameter NASREQ application policy to use for authentication.

**Parameters**     *name* — Specifies the name of the Diameter NASREQ application policy to use for authentication.

### host-limit

**Syntax**     **host-limit** *max-num-hosts*  
              **no host-limit**

**Context**     config>service>vpls>sap>arp-host

**Description**     This command configures the maximum number of ARP hosts.  
                      The **no** form of the command returns the value to the default.

**Default**       1

**Parameters**     *max-num-hosts* — specifies the maximum number of ARP hosts allowed on this SAP.

**Values**        1 — 32767

### min-auth-interval

**Syntax**     **min-auth-interval** *min-auth-interval*  
              **no min-auth-interval**

**Context**     config>service>vpls>sap>arp-host

**Description**     This command configures the minimum authentication interval.  
                      The **no** form of the command returns the value to the default.

**Default**       15

**Parameters**     *min-auth-interval* — Specifies the minimum authentic interval, in minutes.

**Values**        1 — 6000

### arp-reply-agent

**Syntax**     **arp-reply-agent** [*sub-ident*]  
              **no arp-reply-agent**

**Context**     config>service>vpls>sap

**Description**     This command enables a special ARP response mechanism in the system for ARP requests destined to static or dynamic hosts associated with the SAP. The system responds to each ARP request using the hosts MAC address as the both the source MAC address in the Ethernet header and the target hardware address in the ARP header.

                      ARP replies and requests received on a SAP with **arp-reply-agent** enabled will be evaluated by the system against the anti-spoof filter entries associated with the ingress SAP (if the SAP has anti-spoof filtering enabled). ARPs from unknown hosts on the SAP will be discarded when anti-spoof filtering is enabled.

The ARP reply agent only responds if the ARP request enters an interface (SAP, spoke-SDP or mesh-SDP) associated with the VPLS instance of the SAP.

A received ARP request that is not in the ARP reply agent table is flooded to all forwarding interfaces of the VPLS capable of broadcast except the ingress interface while honoring split-horizon constraints.

Static hosts can be defined on the SAP using the **host** command. Dynamic hosts are enabled on the system by enabling the **lease-populate** command in the SAP's **dhcp** context. In the event that both a static host and a dynamic host share the same IP and MAC address, the VPLS ARP reply agent will retain the host information until both the static and dynamic information are removed. In the event that both a static and dynamic host share the same IP address, but different MAC addresses, the VPLS ARP reply agent is populated with the static host information.

The **arp-reply-agent** command will fail if an existing static host on the SAP does not have both MAC and IP addresses specified. Once the ARP reply agent is enabled, creating a static host on the SAP without both an IP address and MAC address will fail.

The ARP-reply-agent may only be enabled on SAPs supporting Ethernet encapsulation.

The **no** form of the command disables ARP-reply-agent functions for static and dynamic hosts on the SAP.

**Default** not enabled

**Parameters** **sub-ident** — configures the arp-reply-agent to discard ARP requests received on the SAP that are targeted for a known host on the same SAP with the same subscriber identification.

Hosts are identified by their subscriber information. For DHCP subscriber hosts, the subscriber hosts, the subscriber information is configured using the optional subscriber parameter string.

When arp-reply-agent is enabled with **sub-ident**:

- If the subscriber information for the destination host exactly matches the subscriber information for the originating host and the destination host is known on the same SAP as the source, the ARP request is silently discarded.
- If the subscriber information for the destination host or originating host is unknown or undefined, the source and destination hosts are not considered to be the same subscriber. The ARP request is forwarded outside the SAP's Split Horizon Group.
- When **sub-ident** is not configured, the arp-reply-agent does not attempt to identify the subscriber information for the destination or originating host and will not discard an ARP request based on subscriber information.

## host

**Syntax** **host** {[**ip** *ip-address*] [**mac** *mac-address*]} [**subscriber-sap-id** | **subscriber** *sub-ident-string*] [**sub-profile** *sub-profile-name*] [**sla-profile** *sla-profile-name*] [**ancp-string** *ancp-string*] [**inter-dest-id** *intermediate-destination-id*] [**mac-linking** *ip-address*][**rip-policy** *rip-policy-name*]

**no host** {[**ip** *ip-address*] [**mac** *ieee-address*]}

**no host all** [**force**]

**Context** config>service>vpls>sap  
config>service>ies>if>sap

**Description** This command creates a static subscriber host for the SAP. Static subscriber hosts may be used by the system for various purposes. Applications within the system that make use of static host entries include anti-spoof, ARP reply agent and source MAC population into the VPLS forwarding database.

Multiple static hosts may be defined on the SAP. Each host is identified by either a source IP address, a source MAC address or both a source IP and source MAC address. Every static host definition must have at least one address defined, IP or MAC.

Static hosts can exist on the SAP even with anti-spoof and ARP reply agent features disabled. When enabled, each feature has different requirements for static hosts.

Use the **no** form of the command to remove a static entry from the system. The specified *ip-address* and *mac-address* must match the host's exact IP and MAC addresses as defined when it was created. When a static host is removed from the SAP, the corresponding anti-spoof filter entry and/or FDB entry is also removed.

**Default** none

**Parameters** **ip** *ip-address* — Specify this optional parameter when defining a static host. The IP address must be specified for **anti-spoof ip**, **anti-spoof ip-mac** and **arp-reply-agent**. Only one static host may be configured on the SAP with a given IP address.

**mac** *mac-address* — Specify this optional parameter when defining a static host. The MAC address must be specified for **anti-spoof mac**, **anti-spoof ip-mac** and **arp-reply-agent**. Multiple static hosts may be configured with the same MAC address given that each definition is distinguished by a unique IP address.

If the MAC address is not specified for a static host, the system will learn respective MAC address dynamically from ARP packets (arp-request or gratuitous-arp) generated by the host with the specified IP address. On a VPLS service, this can occur if arp-reply-agent function is enabled on a given SAP. On Layer 3 services, such as IES) the ARP packets are always examined so no further conditions are applicable.

The learned MAC address will be handled as a MAC address of static host with explicitly defined *mac-address*. Meaning:

- The MAC address will not be aged by the mac-aging or arp-aging timers.
- The MAC address will not be moved to another SAP as a consequence of re-learning event (= event when learning request for the same MAC address comes from another SAP)
- The MAC address will not be flushed from FDB due to SAP failure or STP flush messages.

Every time the given static-host uses different MAC address in its ARP request, the dynamic mac-learning process will be performed. The old MAC address will be overwritten by a new MAC address.

The learned MAC address will not be made persistent (a static host is not a part of the persistency file). A service discontinuity of such a host could be proportional to its arp-cache timeout.

The following interactions are described:

- **Antispoof (all services)** — In case a static IP-only host is configured on a given SAP, both anti-spoof types, IP and IP MAC are supported. Static hosts for which MAC address is not known will not have any antispoof entry. This will be added only after the corresponding MAC has been learned. As a consequence, all traffic generated by the host before sending any arp packets will be most likely dropped.

- Enhanced subscriber management (all services) — ESM is supported in a combination with a static ip-only host. It is assumed that ip-mac antispoofing is enabled. The resources (queues, etc.) are allocated at the time such a host is configured, although they will be effectively used only after antispoof entry has been installed.
- Dual-homing (all services) — It is assumed that static host is configured on both chassis. The dynamic mac-address learning event will be then synchronized (also, if the members are on two different nodes) and corresponding anti-spoof entries will be installed on both chassis.
- MAC-pinning (for VPLS services only) — The dynamically learned MAC address of the static-host will be considered as a static-mac and is not affected by the **no mac-pinning** command.
- ARP-reply-agent (VPLS services only) — It is possible to enable arp-reply-agent on a SAP where static host with ip-only configured. Besides the regular arp-reply-agent functionality (reply to all arp-requests targeting the given host's IP address) learning of the host's MAC address will be performed. As long as no MAC address have been learned no ARP replies on behalf of such host should be expected. Enabling of arp-reply-agent is optional for SAP with ip-only static hosts.

Every static host definition must have at least one address defined, IP or MAC.

**subscriber-sap-id** — Specifies to use the sap-id as the subscriber-id.

**subscriber** *sub-ident-string* — This optional parameter is used to specify an existing subscriber identification profile to be associated with the static subscriber host. The subscriber identification profile is configured in the **config>subscr-mgmt>sub-ident-policy** context. The subscriber information is used by the VPLS SAP arp-reply-agent to determine the proper handling of received ARP requests from subscribers.

- For VPLS SAPs with **arp-reply-agent** enabled with the optional *sub-ident* parameter, the static subscriber host's sub-ident-string is used to determine whether an ARP request received on the SAP is sourced from a host belonging to the same subscriber as the destination host. When both the destination and source hosts from the ARP request are known on the SAP and the subscriber identifications do not match, the ARP request may be forwarded to the rest of the VPLS destinations.

If the static subscriber host's *sub-ident* string is not defined, the host is not considered to belong to the same subscriber as another host on the SAP.

If source or destination host is unknown, the hosts are not considered to belong to the same subscriber. (ARP messages from unknown hosts are subject to anti-spoof filtering rules applied at the SAP.)

If *sub-ident* is not enabled on the SAP arp-reply-agent, subscriber identification matching is not performed on ARP requests received on the SAP.

ARP requests are never forwarded back to the same SAP or within the receiving SAPs Split Horizon Group.

**sub-profile** *sub-profile-name* — This optional parameter is used to specify an existing subscriber profile name to be associated with the static subscriber host. The subscriber profile is configured in the **config>subscr-mgmt>sub-profile** context.

**sla-profile** *sla-profile-name* — This optional parameter is used to specify an existing SLA profile name to be associated with the static subscriber host. The SLA profile is configured in the **config>subscr-mgmt>sla-profile** context.

**ancp-string** *ancp-string* — Specifies the ANCP string associated to this SAP host.

**inter-dest-id** *intermediate-destination-id* — Specifies to which intermediate destination (for example a DSLAM) this host belongs.

---

## Filter and QoS Policy Commands

### egress

<b>Syntax</b>	<b>egress</b>
<b>Context</b>	config>service>vpls>sap config>service>vpls>mesh-sdp config>service>ies>if>sap
<b>Description</b>	This command enables the context to configure egress Quality of Service (QoS) policies and filter policies.  If no QoS policy is defined, the system default QoS policy is used for egress processing. If no egress filter is defined, no filtering is performed.

### ingress

<b>Syntax</b>	<b>ingress</b>
<b>Context</b>	config>service>vpls>sap>egress config>service>vpls>mesh-sdp config>service>vpls>spoke-sdp config>service>ies>if>sap
<b>Description</b>	This command enables the context to configure ingress Quality of Service (QoS) policies and filter policies.  If no QoS policy is defined, the system default sap-ingress QoS policy is used for ingress processing. If no ingress filter is defined, no filtering is performed.

### filter

<b>Syntax</b>	<b>filter ip</b> <i>ip-filter-id</i> <b>filter mac</b> <i>mac-filter-id</i> <b>no filter</b>
<b>Context</b>	config>service>vpls>sap>egress config>service>vpls>sap>ingress config>service>vpls>mesh-sdp config>service>vpls>spoke-sdp config>service>ies>if>spoke-sdp>egress config>service>ies>if>spoke-sdp>ingress
<b>Description</b>	This command associates an IP filter policy or MAC filter policy with an ingress or egress Service Access Point (SAP) or IP interface.

Filter policies control the forwarding and dropping of packets based on IP or MAC matching criteria. There are two types of filter policies: IP and MAC. Only one type may be applied to a SAP at a time.

The **filter** command is used to associate a filter policy with a specified filter ID with an ingress or egress SAP. The filter ID must already be defined before the **filter** command is executed. If the filter policy does not exist, the operation will fail and an error message returned.

In general, filters applied to SAPs (ingress or egress) apply to all packets on the SAP. One exception is non-IP packets are not applied to IP match criteria, so the default action in the filter policy applies to these packets.

The **no** form of this command removes any configured filter ID association with the SAP or IP interface. The filter ID itself is not removed from the system unless the scope of the created filter is set to local. To avoid deletion of the filter ID and only break the association with the service object, use **scope** command within the filter definition to change the scope to **local** or **global**. The default scope of a filter is **local**.

<b>Special Cases</b>	<b>VPLS</b> — Both MAC and IP filters are supported on a VPLS service SAP.
<b>Parameters</b>	<p><b>ip</b> <i>ip-filter-id</i> — Specifies IP filter policy. The filter ID must already exist within the created IP filters.</p> <p><b>Values</b> 1 — 65535</p> <p><b>mac</b> <i>mac-filter-id</i> — Specifies the MAC filter policy. The specified filter ID must already exist within the created MAC filters. The filter policy must already exist within the created MAC filters.</p> <p><b>Values</b> 1 — 65535</p>

## agg-rate

<b>Syntax</b>	<b>[no] agg-rate</b>
<b>Context</b>	<pre>config&gt;service&gt;vprn&gt;subscriber-interface&gt;group-if&gt;sap&gt;egress config&gt;service&gt;ies&gt;subscriber-interface&gt;group-if&gt;sap&gt;egress</pre>
<b>Description</b>	This command is used to control an HQoS aggregate rate limit. It is used in conjunction with the following parameter commands: <b>rate</b> , <b>limit-unused-bandwidth</b> , and <b>queue-frame-based-accounting</b> .

## rate

<b>Syntax</b>	<p><b>rate {max   rate}</b></p> <p><b>no rate</b></p>
<b>Context</b>	<pre>config&gt;service&gt;vprn&gt;subscriber-interface&gt;group-if&gt;sap&gt;egress&gt;agg-rate config&gt;service&gt;ies&gt;subscriber-interface&gt;group-if&gt;sap&gt;egress&gt;agg-rate</pre>
<b>Description</b>	This command defines the enforced aggregate rate for all queues associated with the agg-rate context. A rate must be specified for the agg-rate context to be considered to be active on the context's object



(SAP, subscriber, VPORT etc.).

## limit-unused-bandwidth

<b>Syntax</b>	<b>[no] limit-unused-bandwidth</b>
<b>Context</b>	config>service>vprn>subscriber-interface>group-if>sap>egress>agg-rate config>service>ies>subscriber-interface>group-if>sap>egress>agg-rate
<b>Description</b>	This command is used to enable (or disable) aggregate rate overrun protection on the agg-rate context.

## queue-frame-based-accounting

<b>Syntax</b>	<b>[no] queue-frame-based-accounting</b>
<b>Context</b>	config>service>vprn>subscriber-interface>group-if>sap>egress>agg-rate
<b>Description</b>	This command is used to enabled (or disable) frame based accounting on all queues associated with the agg-rate context. Only supported on Ethernet ports. Not supported on HSMDB Ethernet ports.

## filter

<b>Syntax</b>	<b>filter ip <i>ip-filter-id</i></b> <b>no filter [<i>ip ip-filter-id</i>]</b>
<b>Context</b>	config>service>ies>if>sap>egress config>service>ies>if>sap>ingress
<b>Description</b>	<p>This command associates an IP filter policy with an ingress or egress Service Access Point (SAP). Filter policies control the forwarding and dropping of packets based on the matching criteria. MAC filters are only allowed on Epipe and Virtual Private LAN Service (VPLS) SAPs.</p> <p>The <b>filter</b> command is used to associate a filter policy with a specified filter ID with an ingress or egress SAP. The filter policy must already be defined before the <b>filter</b> command is executed. If the filter policy does not exist, the operation will fail and an error message returned.</p> <p>In general, filters applied to SAPs (ingress or egress) apply to all packets on the SAP. One exception is non-IP packets are not applied to the match criteria, so the default action in the filter policy applies to these packets.</p> <p>The <b>no</b> form of this command removes any configured filter ID association with the SAP. The filter ID itself is not removed from the system unless the scope of the created filter is set to local. To avoid deletion of the filter ID and only break the association with the service object, use <b>scope</b> command within the filter definition to change the scope to <b>local</b> or <b>global</b>. The default scope of a filter is <b>local</b>.</p>
<b>Special Cases</b>	<b>IES —</b> Only IP filters are supported on an IES IP interface, and the filters only apply to routed traffic.

<b>Parameters</b>	<b>ip</b> <i>ip-filter-id</i> — Specifies IP filter policy. The filter ID must already exist within the created IP filters.
<b>Values</b>	1 — 65535

## multicast-group

<b>Syntax</b>	<b>multicast-group</b> <i>group-name</i> <b>no multicast-group</b>
<b>Context</b>	config>service>vpls>sap>egress
<b>Description</b>	<p>This command places a VPLS Ethernet SAP into an egress multicast group. The SAP must comply with the egress multicast group's common requirements for member SAPs. If the SAP does not comply, the command will fail and the SAP will not be a member of the group. Common requirements for an egress multicast group are listed below:</p> <ul style="list-style-type: none"> <li>• If an egress-filter is specified on the egress multicast group, the SAP must have the same egress filter applied.</li> <li>• If an egress-filter is not defined on the egress multicast group, the SAP cannot have an egress filter applied.</li> <li>• If the egress multicast group has an encap-type set to null, the SAP must be defined on a port with the port encapsulation type set to null.</li> <li>• If the egress multicast group has an encap-type set to dot1q, the SAP must be defined on a port with the port encapsulation type set to dot1q and the port's dot1q-etype must match the dot1q-etype defined on the egress multicast group.</li> <li>• The access port the SAP is created on cannot currently be an egress mirror source.</li> </ul> <p>Once a SAP is a member of an egress multicast group, the following rules apply:</p> <ul style="list-style-type: none"> <li>• The egress filter defined on the SAP cannot be removed or modified. Egress filtering is managed at the egress multicast group for member SAPs.</li> <li>• If the encapsulation type for the access port the SAP is created on is set to dot1q, the port's dot1q-etype value cannot be changed.</li> <li>• Attempting to define an access port with a SAP that is currently defined in an egress multicast group as an egress mirror source will fail.</li> </ul>

Once a SAP is included in an egress multicast group, it is then eligible for efficient multicast replication if the egress forwarding plane performing replication for the SAP is capable. If the SAP is defined as a Link Aggregation Group (LAG) SAP, it is possible that some links in the LAG are on forwarding planes that support efficient multicast replication while others are not. The fact that some or all the forwarding planes associated with the SAP cannot perform efficient multicast replication does not affect the ability to place the SAP into an egress multicast group.

A SAP may be a member of one and only one egress multicast group. If the multicast-group command is executed with another egress multicast group name, the system will attempt to move the SAP to the specified group. If the SAP is not placed into the new group, the SAP will remain a member of the previous egress multicast group. Moving a SAP into an egress multicast group may cause a momentary gap in replications to the SAP destination while the move is being processed.

The **no** form of the command removes the SAP from any egress multicast group in which it may currently have membership. The SAP will be removed from all efficient multicast replication chains

and normal replication will apply to the SAP. A momentary gap in replications to the SAP destination while it is being moved is possible. If the SAP is not currently a member in an egress multicast group, the command has no effect.

<b>Default</b>	no multicast-group
<b>Parameters</b>	<i>group-name</i> — The <i>group-name</i> is required when specifying egress multicast group membership on a SAP. An egress multicast group with the specified egress-multicast-group-name must exist and the SAP must pass all common requirements or the command will fail.
<b>Values</b>	Any valid egress multicast group name.
<b>Default</b>	None, an egress multicast group name must be explicitly specified.

## qos

<b>Syntax</b>	<b>qos</b> <i>policy-id</i> <b>no qos</b>
<b>Context</b>	config>service>vpls>sap>egress config>service>vpls>sap>ingress config>service>ies>if>sap>egress config>service>ies>if>sap>ingress
<b>Description</b>	<p>This command associates a Quality of Service (QoS) policy with an ingress or egress Service Access Point (SAP) or IP interface.</p> <p>QoS ingress and egress policies are important for the enforcement of SLA agreements. The policy ID must be defined prior to associating the policy with a SAP or IP interface. If the <i>policy-id</i> does not exist, an error will be returned.</p> <p>The <b>qos</b> command is used to associate both ingress and egress QoS policies. The <b>qos</b> command only allows ingress policies to be associated on SAP or IP interface ingress and egress policies on SAP or IP interface egress. Attempts to associate a QoS policy of the wrong type returns an error.</p> <p>Only one ingress and one egress QoS policy can be associated with a SAP or IP interface at one time. Attempts to associate a second QoS policy of a given type will return an error.</p> <p>By default, no specific QoS policy is associated with the SAP or IP interface for ingress or egress, so the default QoS policy is used.</p> <p>The <b>no</b> form of this command removes the QoS policy association from the SAP or IP interface, and the QoS policy reverts to the default.</p> <p><i>policy-id</i> — The ingress/egress policy ID to associate with SAP or IP interface on ingress/egress. The policy ID must already exist.</p> <p><b>Values</b>      1 — 65535</p> <p><b>shared-queuing</b> — Specify the ingress shared queue policy used by this SAP. When the value of this object is null, the SAP will use individual ingress QoS queues, instead of the shared ones.</p>

## queue-override

<b>Syntax</b>	<b>[no] queue-override</b>
<b>Context</b>	config>service>ies>if>sap>egress
<b>Description</b>	This command enables the context to configure override values for the specified SAP egress QoS queue. These values override the corresponding ones specified in the associated SAP egress QoS policy.

## queue

<b>Syntax</b>	<b>[no] queue <i>queue-id</i></b>
<b>Context</b>	config>service>ies>if>sap>egress>queue-override
<b>Description</b>	This command specifies the ID of the queue whose parameters are to be overridden.
<b>Parameters</b>	<i>queue-id</i> — The queue ID whose parameters are to be overridden. <b>Values</b> 1 — 32

## adaptation-rule

<b>Syntax</b>	<b>adaptation-rule [pir {max   min   closest}] [cir {max   min   closest}] no adaptation-rule</b>
<b>Context</b>	config>service>ies>if>sap>egress>queue-override>queue
<b>Description</b>	<p>This command can be used to override specific attributes of the specified queue's adaptation rule parameters. The adaptation rule controls the method used by the system to derive the operational CIR and PIR settings when the queue is provisioned in hardware. For the CIR and PIR parameters individually, the system attempts to find the best operational rate depending on the defined constraint.</p> <p>The <b>no</b> form of the command removes any explicitly defined constraints used to derive the operational CIR and PIR created by the application of the policy. When a specific <b>adaptation-rule</b> is removed, the default constraints for <b>rate</b> and <b>cir</b> apply.</p>
<b>Default</b>	no adaptation-rule
<b>Parameters</b>	<p><b>pir</b> — The <b>pir</b> parameter defines the constraints enforced when adapting the PIR rate defined within the <b>queue <i>queue-id</i> rate</b> command. The <b>pir</b> parameter requires a qualifier that defines the constraint used when deriving the operational PIR for the queue. When the <b>rate</b> command is not specified, the default applies.</p> <p><b>max</b> — The <b>max</b> (maximum) option is mutually exclusive with the <b>min</b> and <b>closest</b> options. When <b>max</b> is defined, the operational PIR for the queue will be equal to or less than the administrative rate specified using the <b>rate</b> command.</p> <p><b>min</b> — The <b>min</b> (minimum) option is mutually exclusive with the <b>max</b> and <b>closest</b> options. When <b>min</b> is defined, the operational PIR for the queue will be equal to or greater than the administrative rate specified using the <b>rate</b> command.</p>

**closest** — The **closest** parameter is mutually exclusive with the **min** and **max** parameter. When **closest** is defined, the operational PIR for the queue will be the rate closest to the rate specified using the **rate** command.

**cir** — The **cir** parameter defines the constraints enforced when adapting the CIR rate defined within the **queue queue-id rate** command. The **cir** parameter requires a qualifier that defines the constraint used when deriving the operational CIR for the queue. When the **cir** parameter is not specified, the default constraint applies.

## avg-frame-overhead

<b>Syntax</b>	<b>avg-frame-overhead percent</b> <b>no avg-frame-overhead</b>
<b>Context</b>	config>service>ies>if>sap>egress>queue-override
<b>Description</b>	This command configures the average frame overhead to define the average percentage that the offered load to a queue will expand during the frame encapsulation process before sending traffic on-the-wire. While the avg-frame-overhead value may be defined on any queue, it is only used by the system for queues that egress a SONET or SDH port or channel. Queues operating on egress Ethernet ports automatically calculate the frame encapsulation overhead based on a 20 byte per packet rule (8 bytes for preamble and 12 bytes for Inter-Frame Gap).

When calculating the frame encapsulation overhead for port scheduling purposes, the system determines the following values:

- **Offered-load** — The offered-load of a queue is calculated by starting with the queue depth in octets, adding the received octets at the queue and subtracting queue discard octets. The result is the number of octets the queue has available to transmit. This is the packet based offered-load.
- **Frame encapsulation overhead** — Using the avg-frame-overhead parameter, the frame encapsulation overhead is simply the queue's current offered-load (how much has been received by the queue) multiplied by the avg-frame-overhead. If a queue had an offered load of 10000 octets and the avg-frame-overhead equals 10%, the frame encapsulation overhead would be 10000 x 0.1 or 1000 octets.

For egress Ethernet queues, the frame encapsulation overhead is calculated by multiplying the number of offered-packets for the queue by 20 bytes. If a queue was offered 50 packets then the frame encapsulation overhead would be 50 x 20 or 1000 octets.

- **Frame based offered-load** — The frame based offered-load is calculated by adding the offered-load to the frame encapsulation overhead. If the offered-load is 10000 octets and the encapsulation overhead was 1000 octets, the frame based offered-load would equal 11000 octets.
- **Packet to frame factor** — The packet to frame factor is calculated by dividing the frame encapsulation overhead by the queue's offered-load (packet based). If the frame encapsulation overhead is 1000 octets and the offered-load is 10000 octets then the packet to frame factor would be 1000 / 10000 or 0.1. When in use, the avg-frame-overhead will be the same as the packet to frame factor making this calculation unnecessary.
- **Frame based CIR** — The frame based CIR is calculated by multiplying the packet to frame factor with the queue's configured CIR and then adding that result to that CIR. If the queue CIR is set at 500 octets and the packet to frame factor equals 0.1, the frame based CIR would be 500 x 1.1 or 550 octets.

- **Frame based within-cir offered-load** — The frame based within-cir offered-load is the portion of the frame based offered-load considered to be within the frame-based CIR. The frame based within-cir offered-load is the lesser of the frame based offered-load and the frame based CIR. If the frame based offered-load equaled 11000 octets and the frame based CIR equaled 550 octets, the frame based within-cir offered-load would be limited to 550 octets. If the frame based offered-load equaled 450 octets and the frame based CIR equaled 550 octets, the frame based within-cir offered-load would equal 450 octets (or the entire frame based offered-load).

As a special case, when a queue or associated intermediate scheduler is configured with a CIR-weight equal to 0, the system automatically sets the queue's frame based within-cir offered-load to 0, preventing it from receiving bandwidth during the port scheduler's within-cir pass.

- **Frame based PIR** — The frame based PIR is calculated by multiplying the packet to frame factor with the queue's configured PIR and then adding the result to that PIR. If the queue PIR is set to 7500 octets and the packet to frame factor equals 0.1, the frame based PIR would be  $7500 \times 1.1$  or 8250 octets.
- **Frame based within-pir offered-load** — The frame based within-pir offered-load is the portion of the frame based offered-load considered to be within the frame based PIR. The frame based within-pir offered-load is the lesser of the frame based offered-load and the frame based PIR. If the frame based offered-load equaled 11000 octets and the frame based PIR equaled 8250 octets, the frame based within-pir offered-load would be limited to 8250 octets. If the frame based offered-load equaled 7000 octets and the frame based PIR equaled 8250 octets, the frame based within-pir offered load would equal 7000 octets.

**Port scheduler operation using frame transformed rates** — The port scheduler uses the frame based rates to figure the maximum rates that each queue may receive during the within-cir and above-cir bandwidth allocation passes. During the within-cir pass, a queue may receive up to its frame based within-cir offered-load. The maximum it may receive during the above-cir pass is the difference between the frame based within-pir offered load and the amount of actual bandwidth allocated during the within-cir pass.

**SAP and subscriber SLA-profile average frame overhead override** — The average frame overhead parameter on a sap-egress may be overridden at an individual egress queue basis. On each SAP and within the sla-profile policy used by subscribers an avg-frame-overhead command may be defined under the queue-override context for each queue. When overridden, the queue instance will use its local value for the average frame overhead instead of the sap-egress defined overhead.

The **no** form of this command restores the average frame overhead parameter for the queue to the default value of 0 percent. When set to 0, the system uses the packet based queue statistics for calculating port scheduler priority bandwidth allocation. If the no avg-frame-overhead command is executed in a queue-override queue id context, the avg-frame-overhead setting for the queue within the sap-egress QoS policy takes effect.

<b>Default</b>	0
<b>Parameters</b>	<i>percent</i> — This parameter sets the average amount of packet-to-frame encapsulation overhead expected for the queue. This value is not used by the system for egress Ethernet queues.
<b>Values</b>	0 — 100

## cbs

<b>Syntax</b>	<b>cbs</b> <i>size-in-kbytes</i> <b>no cbs</b>
<b>Context</b>	config>service>ies>if>sap>egress>queue-override>queue
<b>Description</b>	<p>This command can be used to override specific attributes of the specified queue's CBS parameters.</p> <p>It is permissible, and possibly desirable, to oversubscribe the total CBS reserved buffers for a given access port egress buffer pool. Oversubscription may be desirable due to the potential large number of service queues and the economy of statistical multiplexing the individual queue's CBS settings into the defined reserved total.</p> <p>When oversubscribing the reserved total, it is possible for a queue depth to be lower than its CBS setting and still not receive a buffer from the buffer pool for an ingress frame. As more queues are using their CBS buffers and the total in use exceeds the defined reserved total, essentially the buffers are being removed from the shared portion of the pool without the shared in use average and total counts being decremented. This can affect the operation of the high and low priority RED slopes on the pool, causing them to miscalculate when to start randomly drop packets.</p> <p>If the CBS value is larger than the MBS value, an error will occur, preventing the CBS change.</p> <p>The <b>no</b> form of this command returns the CBS size to the default value.</p>
<b>Default</b>	no cbs
<b>Parameters</b>	<p><i>size-in-kbytes</i> — The size parameter is an integer expression of the number of kilobytes reserved for the queue. If a value of 10KBytes is desired, enter the value 10. A value of 0 specifies that no reserved buffers are required by the queue (a minimal reserved size can still be applied for scheduling purposes).</p> <p><b>Values</b>      0 — 131072 or default</p>

## high-prio-only

<b>Syntax</b>	<b>high-prio-only</b> <i>percent</i> <b>no high-prio-only</b>
<b>Context</b>	config>service>ies>if>sap>egress>queue-override>queue
<b>Description</b>	<p>This command can be used to override specific attributes of the specified queue's high-prio-only parameters. The <b>high-prio-only</b> command configures the percentage of buffer space for the queue, used exclusively by high priority packets.</p> <p>The priority of a packet can only be set in the SAP ingress QoS policy and is only applicable on the ingress queues for a SAP. The <b>high-prio-only</b> parameter is used to override the default value derived from the <b>network-queue</b> command.</p> <p>The defined <b>high-prio-only</b> value cannot be greater than the MBS size of the queue. Attempting to change the MBS to a value smaller than the high priority reserve will generate an error and fail execution. Attempting to set the <b>high-prio-only</b> value larger than the current MBS size will also result in an error and fail execution.</p> <p>The <b>no</b> form of this command restores the default high priority reserved size.</p>

**Parameters** *percent* — The *percent* parameter is the percentage reserved for high priority traffic on the queue. If a value of 10KBytes is desired, enter the value 10. A value of 0 specifies that none of the MBS of the queue will be reserved for high priority traffic. This does not affect RED slope operation for packets attempting to be queued.

**Values** 0 — 100 | default

## mbs

**Syntax** **mbs** {*size-in-kbytes* | **default**}  
**no mbs**

**Context** config>service>ies>if>sap>egress>queue-override>queue

**Description** This command can be used to override specific attributes of the specified queue's MBS parameters. The MBS is a mechanism to override the default maximum size for the queue.

The sum of the MBS for all queues on an egress access port can oversubscribe the total amount of buffering available. When congestion occurs and buffers become scarce, access to buffers is controlled by the RED slope a packet is associated with. A queue that has not exceeded its MBS size is not guaranteed that a buffer will be available when needed or that the packets RED slope will not force the discard of the packet. Setting proper CBS parameters and controlling CBS oversubscription is one major safeguard to queue starvation (when a queue does not receive its fair share of buffers). Another is properly setting the RED slope parameters for the needs of services on this port or channel.

If the CBS value is larger than the MBS value, an error will occur, preventing the MBS change.

The **no** form of this command returns the MBS size assigned to the queue.

**Default** default

**Parameters** *size-in-kbytes* — The size parameter is an integer expression of the maximum number of kilobytes of buffering allowed for the queue. For a value of 100 kbps, enter the value 100. A value of 0 causes the queue to discard all packets.

**Values** 0 — 131072 or default

## rate

**Syntax** **rate** *pir-rate* [*cir cir-rate*]  
**no rate**

**Context** config>service>ies>if>sap>egress>queue-override>queue

**Description** This command can be used to override specific attributes of the specified queue's Peak Information Rate (PIR) and the Committed Information Rate (CIR) parameters.

The PIR defines the maximum rate that the queue can transmit packets out an egress interface (for SAP egress queues). Defining a PIR does not necessarily guarantee that the queue can transmit at the intended rate. The actual rate sustained by the queue can be limited by oversubscription factors or available egress bandwidth.



The CIR defines the rate at which the system prioritizes the queue over other queues competing for the same bandwidth. In-profile packets are preferentially queued by the system at egress and at subsequent next hop nodes where the packet can traverse. To be properly handled as in- or out-of-profile throughout the network, the packets must be marked accordingly for profiling at each hop.

The CIR can be used by the queue's parent commands *cir-level* and *cir-weight* parameters to define the amount of bandwidth considered to be committed for the child queue during bandwidth allocation by the parent scheduler.

The **rate** command can be executed at any time, altering the PIR and CIR rates for all queues created through the association of the SAP egress QoS policy with the *queue-id*.

The **no** form of the command returns all queues created with the *queue-id* by association with the QoS policy to the default PIR and CIR parameters (**max**, 0).

<b>Default</b>	<b>rate max cir 0</b> — The <b>max</b> default specifies the amount of bandwidth in kilobits per second (thousand bits per second). The <b>max</b> value is mutually exclusive to the <b>pir-rate</b> value.
<b>Parameters</b>	<p><i>pir-rate</i> — Defines the administrative PIR rate, in kilobits, for the queue. When the <b>rate</b> command is executed, a valid PIR setting must be explicitly defined. When the <b>rate</b> command has not been executed, the default PIR of <b>max</b> is assumed.</p> <p>Fractional values are not allowed and must be given as a positive integer.</p> <p>The actual PIR rate is dependent on the queue's <b>adaptation-rule</b> parameters and the actual hardware where the queue is provisioned.</p> <p><b>Values</b>      1 — 100000000</p> <p><b>Default</b>      <b>max</b></p> <p><b>cir</b> <i>cir-rate</i> — The <b>cir</b> parameter overrides the default administrative CIR used by the queue. When the <b>rate</b> command is executed, a CIR setting is optional. When the <b>rate</b> command has not been executed or the <b>cir</b> parameter is not explicitly specified, the default CIR (0) is assumed.</p> <p>Fractional values are not allowed and must be given as a positive integer. The <b>sum</b> keyword specifies that the CIR be used as the summed CIR values of the children schedulers or queues.</p> <p><b>Values</b>      0 — 100000000, <b>max</b>, <b>sum</b></p> <p><b>Default</b>      0</p>

## scheduler-policy

<b>Syntax</b>	<b>scheduler-policy</b> <i>scheduler-policy-name</i> <b>no scheduler-policy</b>
<b>Context</b>	config>service>vpls>sap>ingress config>service>vpls>sap>egress config>service>ies>if>sap>egress config>service>ies>if>sap>ingress
<b>Description</b>	This command applies an existing scheduler policy to an ingress or egress scheduler used by SAP queues associated with this multi-service customer site. The schedulers defined in the scheduler policy can only be created once the customer site has been appropriately assigned to a chassis port, channel or slot. Scheduler policies are defined in the <b>config&gt;qos&gt;scheduler-policy</b> <i>scheduler-policy-name</i> context.

The **no** form of this command removes the configured ingress or egress scheduler policy from the multi-service customer site. When the policy is removed, the schedulers created due to the policy are removed also making them unavailable for the ingress SAP queues associated with the customer site. Queues that lose their parent scheduler association are deemed to be orphaned and are no longer subject to a virtual scheduler. The SAPs that have ingress queues reliant on the removed schedulers enter into an operational state depicting the orphaned status of one or more queues. When the **no scheduler-policy** command is executed, the customer site ingress or egress node will not contain an applied scheduler policy.

*scheduler-policy-name* — The *scheduler-policy-name* parameter applies an existing scheduler policy that was created in the **config>qos>scheduler-policy scheduler-policy-name** context to create the hierarchy of ingress or egress virtual schedulers. The scheduler names defined within the policy are created and made available to any ingress or egress queues created on associated SAPs.

**Values** Any existing valid scheduler policy name.

## block-on-mesh-failure

<b>Syntax</b>	<b>[no] block-on-mesh-failure</b>
<b>Context</b>	config>service>vpls>spoke-sdp
<b>Description</b>	This command enables blocking (bring the spoke SDP to an operationally down state) after all configured mesh SDPs are in operationally down state. This event is signalled to corresponding T-LDP peer by withdrawing service label (status-bit-signaling non-capable peer) or by setting "PW not forwarding" status bit in T-LDP message (status-bit-signaling capable peer).
<b>Default</b>	disabled

## max-nbr-mac-addr

<b>Syntax</b>	<b>max-nbr-mac-addr table-size</b> <b>no max-nbr-mac-addr</b>
<b>Context</b>	config>service>vpls>sap config>service>vpls>spoke-sdp
<b>Description</b>	<p>This command specifies the maximum number of FDB entries for both learned and static MAC addresses for this SAP or spoke SDP.</p> <p>When the configured limit has been reached, and discard-unknown-source has been enabled for this SAP or spoke SDP (see the <b>discard-unknown-source</b> command), packets with unknown source MAC addresses will be discarded.</p> <p>The <b>no</b> form of the command restores the global MAC learning limitations for the SAP or spoke SDP.</p>
<b>Default</b>	no max-nbr-mac-addr
<b>Parameters</b>	<i>table-size</i> — The maximum number of MAC entries in the FDB from this SAP.
<b>Values</b>	1 — 131071

## multi-service-site

<b>Syntax</b>	<b>multi-service-site</b> <i>customer-site-name</i> <b>no multi-service-site</b>
<b>Context</b>	config>service>vpls>sap config>service>ies>sap
<b>Description</b>	<p>This command associates the SAP with a <i>customer-site-name</i>. If the specified <i>customer-site-name</i> does not exist in the context of the service customer ID an error occurs and the command will not execute. If <i>customer-site-name</i> exists, the current and future defined queues on the SAP (ingress and egress) will attempt to use the scheduler hierarchies created within <i>customer-site-name</i> as parent schedulers.</p> <p>This command is mutually exclusive with the SAP ingress and egress <b>scheduler-policy</b> commands. If a <b>scheduler-policy</b> has been applied to either the ingress or egress nodes on the SAP, the <b>multi-service-site</b> command will fail without executing. The locally applied scheduler policies must be removed prior to executing the <b>multi-service-site</b> command.</p> <p>The <b>no</b> form of the command removes the SAP from any multi-service customer site the SAP belongs to. Removing the site can cause existing or future queues to enter an orphaned state.</p>
<b>Default</b>	None
	<p><i>customer-site-name</i> — The customer-site-name must exist in the context of the customer-id defined as the service owner. If customer-site-name exists and local scheduler policies have not been applied to the SAP, the current and future queues defined on the SAP will look for their parent schedulers within the scheduler hierarchies defined on customer-site-name.</p> <p><b>Values</b> Any valid customer-site-name created within the context of the customer-id.</p>

## retail-service-id

<b>Syntax</b>	<b>retail-svc-id</b> <i>service-id</i> <b>no retail-svc-id</b>
<b>Context</b>	config>service>vprn>sub-if>grp-if>sap>static-host
<b>Description</b>	This command specifies the service id of the retailer IES/VP RN service to which the static IPv6 host belongs. A corresponding retailer subscriber interface must exist in the specified service.
<b>Default</b>	no retail-svc-id
<b>Parameters</b>	<i>service-id</i> — Specifies the retailer service id.
	<b>Values</b> 1— 2148007978 or service name up to 64 characters in length

## static-host

<b>Syntax</b>	<b>static-host ip</b> <i>ip-prefix[/prefix-length]</i> [ <b>mac</b> <i>ieee-address</i> ] [ <b>create</b> ] <b>no static-host ip</b> <i>ip-prefix[/prefix-length]</i> <b>mac</b> <i>ieee-address</i> <b>no static-host all</b> [ <b>force</b> ]
---------------	---

**no static-host ip** *ip-prefix[/prefix-length]*

<b>Context</b>	config>service>ies>if>sap config>service>vpls>sap		
<b>Description</b>	This command configures a static host on this SAP.		
<b>Syntax</b>	<b>ip</b> <i>ip-prefix[/prefix-length]</i> — Specifies the IPv4 address, IPv6 address or the IPv6 prefix.		
<b>Values</b>	ip-prefix[/prefix* :	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:d.d.d.d
			x - [0..FFFF]H
			d - [0..255]D
	ipv6-prefix-le - [0..128]		
	<b>mac</b> <i>ieee-address</i> — Specify this optional parameter when defining a static host. Every static host definition must have at least one address defined, IP or MAC.		
	<b>force</b> — Specifies the forced removal of the static host addresses.		
	<b>sla-profile</b> <i>sla-profile-name</i> — This optional parameter is used to specify an existing SLA profile name to be associated with the static subscriber host. The SLA profile is configured in the <b>config&gt;subscr-mgmt&gt;sla-profile</b> context.		

## ancp-string

<b>Syntax</b>	<b>ancp-string</b> <i>ancp-string</i> <b>no ancp-string</b>
<b>Context</b>	config>service>ies>if>sap>static-host config>service>vpls>sap>static-host
<b>Description</b>	This command specifies the ANCP string associated to this SAP host.
<b>Parameters</b>	<i>ancp-string</i> — Specifies the ANCP string up to 63 characters in length.

## app-profile

<b>Syntax</b>	<b>app-profile</b> <i>app-profile-name</i> <b>no app-profile</b>
<b>Context</b>	config>service>ies>if>sap>static-host config>service>vpls>sap>static-host
<b>Description</b>	This command specifies an application profile name.
<b>Parameters</b>	<i>app-profile-name</i> — Specifies the application profile name up to 32 characters in length.

## inter-dest-id

<b>Syntax</b>	<b>inter-dest-id</b> <i>intermediate-destination-id</i> <b>no inter-dest-id</b>
<b>Context</b>	config>service>ies>if>sap>static-host config>service>vpls>sap>static-host
<b>Description</b>	This command specifies to which intermediate destination (for example a DSLAM) this host belongs.
<b>Parameters</b>	<i>intermediate-destination-id</i> — Specifies the intermediate destination ID.

## sla-profile

<b>Syntax</b>	<b>sla-profile</b> <i>sla-profile-name</i> <b>no sla-profile</b>
<b>Context</b>	config>service>ies>if>sap>static-host config>service>vpls>sap>static-host
<b>Description</b>	This command specifies an existing SLA profile name to be associated with the static subscriber host. The SLA profile is configured in the <b>config&gt;subscr-mgmt&gt;sla-profile</b> context.
<b>Parameters</b>	<i>sla-profile-name</i> — Specifies the SLA profile name.

## sub-profile

<b>Syntax</b>	<b>sub-profile</b> <i>sub-profile-name</i> <b>no sub-profile</b>
<b>Context</b>	config>service>ies>if>sap>static-host config>service>vpls>sap>static-host
<b>Description</b>	This command specifies an existing subscriber profile name to be associated with the static subscriber host.
<b>Parameters</b>	<i>sub-profile-name</i> — Specifies the sub-profile name.

## subscriber

<b>Syntax</b>	<b>subscriber</b> <i>sub-ident</i> <b>no subscriber</b>
<b>Context</b>	config>service>ies>if>sap>static-host config>service>vpls>sap>static-host
<b>Description</b>	This command specifies an existing subscriber identification profile to be associated with the static subscriber host.

**Parameters**     *sub-ident* — Specifies the subscriber identification.

### subscriber-sap-id

**Syntax**     **[no] subscriber-sap-id**

**Context**     config>service>ies>if>sap>static-host  
                config>service>vpls>sap>static-host

**Description**     This command enables using the SAP ID as subscriber id.

**Parameters**     **subscriber-sap-id** — Specifies to use the sap-id as the subscriber-id.

### static-mac

**Syntax**     **[no] static-mac** *ieee-mac-address*

**Context**     config>service>vpls>sap  
                config>service>vpls>mesh-sdp  
                config>service>vpls>spoke-sdp

**Description**     This command creates a remote static MAC entry in the Virtual Private LAN Service (VPLS) forwarding database (FDB) associated with the Service Distribution Point (SDP).

In a VPLS service, MAC addresses are associated with a Service Access Point (SAP) or with a Service Distribution Point (SDP). MACs associated with a SAP are classified as local MACs, and MACs associated with an SDP are remote MACs.

Remote static MAC entries create a permanent MAC address to SDP association in the forwarding database for the VPLS instance so that MAC address will not be learned on the edge device.

Note that static MAC definitions on one edge device are not propagated to other edge devices participating in the VPLS instance, that is, each edge device has an independent forwarding database for the VPLS.

Only one static MAC entry (local or remote) can be defined per MAC address per VPLS instance.

By default, no static MAC address entries are defined for the SDP.

The **no** form of this command deletes the static MAC entry with the specified MAC address associated with the SDP from the VPLS forwarding database.

*ieee-mac-address* — Specifies the 48-bit MAC address for the static ARP in the form *aa:bb:cc:dd:ee:ff* or *aa-bb-cc-dd-ee-ff* where *aa*, *bb*, *cc*, *dd*, *ee* and *ff* are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

## VPLS and IES SDP and SAP Commands

### mesh-sdp

**mesh-sdp** *sdp-id*[:*vc-id*] [**vc-type** {**ether** | **vlan**}]  
**no mesh-sdp** *sdp-id*[:*vc-id*]

**Context** config>service>vpls  
 config>service>vpls>mesh-sdp

**Description** This command binds a VPLS service to an existing Service Distribution Point (SDP).  
 Mesh SDPs bound to a service are logically treated like a single bridge “port” for flooded traffic where flooded traffic received on any mesh SDP on the service is replicated to other “ports” (spoke SDPs and SAPs) and not transmitted on any mesh SDPs.  
 Note that this command creates a binding between a service and an SDP. The SDP has an operational state which determines the operational state of the SDP within the service. For example, if the SDP is administratively or operationally down, the SDP for the service will be down.  
 The SDP must already be defined in the **config>service>sdp** context in order to associate the SDP with an Epipe or VPLS service. If the **sdp sdp-id** is not already configured, an error message is generated. If the *sdp-id* does exist, a binding between that *sdp-id* and the service is created.  
 SDPs must be explicitly associated and bound to a service. If an SDP is not bound to a service, no far-end router devices can participate in the service.  
 The **no** form of this command removes the SDP binding from the service. The SDP configuration is not affected; only the binding of the SDP to a service. Once removed, no packets are forwarded to the far-end router.

**Default** No *sdp-id* is bound to a service.

**Special Cases** **VPLS** — Several SDPs can be bound to a VPLS. Each SDP must be destined to a different router. If two *sdp-id* bindings terminate on the same router, an error occurs and the second SDP binding is rejected.

**Parameters** *sdp-id* — The SDP identifier.  
**Values** 1 — 17407  
*vc-id* — The virtual circuit identifier. This value is used to validate the VC ID portion of each mesh SDP binding defined in the service. The default value of this object is equal to the service ID.  
**Values** 1 — 4294967295  
**vc-type** — This command overrides the default VC type signaled for the spoke or mesh binding to the far end of the SDP. The VC type is a 15 bit-quantity containing a value which represents the type of VC. The actual signaling of the VC type depends on the signaling parameter defined for the SDP. If signaling is disabled, the **vc-type** command can still be used to define the dot1q value expected by the far-end provider equipment. A change of the binding’s VC type causes the binding to signal the new VC type to the far end when signaling is enabled.  
 VC types are derived according to IETF *draft-martini-l2circuit-trans-mpls*.

- The VC type value for Ethernet is 0x0005.
- The VC type value for an Ethernet VLAN is 0x0004.

**ether** — Defines the VC type as Ethernet. The **vlan** keyword is mutually exclusive. When the VC type is not defined then the default is Ethernet for spoke SDP bindings. Defining Ethernet is the same as executing **no vc-type** and restores the default VC type for the spoke SDP binding. (hex 5)

**vlan** — Defines the VC type as VLAN. The **ethernet** keyword is mutually exclusive. When the VC type is not defined then the default is Ethernet for spoke SDP bindings.

## spoke-sdp

<b>Syntax</b>	<b>spoke-sdp</b> <i>sdp-id[:vc-id]</i> [ <b>vc-type</b> { <b>ether</b>   <b>vlan</b> }] [ <b>split-horizon-group</b> <i>group-name</i> ] <b>no spoke-sdp</b> <i>sdp-id[:vc-id]</i>
<b>Context</b>	config>service>vpls
<b>Description</b>	<p>This command binds a service to an existing Service Distribution Point (SDP).</p> <p>A spoke SDP is treated like the equivalent of a traditional bridge “port” where flooded traffic received on the spoke SDP is replicated on all other “ports” (other spoke and mesh SDPs or SAPs) and not transmitted on the port it was received.</p> <p>The SDP has an operational state which determines the operational state of the SDP within the service. For example, if the SDP is administratively or operationally down, the SDP for the service will be down.</p> <p>The SDP must already be defined in the <b>config&gt;service&gt;sdp</b> context in order to associate an SDP with a VPLS service. If the <b>sdp sdp-id</b> is not already configured, an error message is generated. If the <i>sdp-id</i> does exist, a binding between that <i>sdp-id</i> and the service is created.</p> <p>SDPs must be explicitly associated and bound to a service. If an SDP is not bound to a service, no far-end ESS-Series devices can participate in the service.</p> <p>The <b>no</b> form of this command removes the SDP binding from the service. The SDP configuration is not affected; only the binding of the SDP to a service. Once removed, no packets are forwarded to the far-end router.</p>
<b>Default</b>	No <i>sdp-id</i> is bound to a service.
<b>Special Cases</b>	<b>VPLS</b> — Several SDPs can be bound to a VPLS service. Each SDP must use unique <i>vc-ids</i> . An error message is generated if two SDP bindings with identical <i>vc-ids</i> terminate on the same router. Split horizon groups can only be created in the scope of a VPLS service.
<b>Parameters</b>	<p><i>sdp-id</i> — The SDP identifier.</p> <p><b>Values</b> 1 — 17407</p> <p><i>vc-id</i> — The virtual circuit identifier.</p> <p><b>Values</b> 1 — 4294967295</p> <p><b>vc-type</b> — This command overrides the default VC type signaled for the spoke or mesh binding to the far end of the SDP. The VC type is a 15 bit-quantity containing a value which represents the type of VC. The actual signaling of the VC type depends on the signaling parameter defined for</p>



the SDP. If signaling is disabled, the **vc-type** command can still be used to define the dot1q value expected by the far-end provider equipment. A change of the binding's VC type causes the binding to signal the new VC type to the far end when signaling is enabled.

VC types are derived according to IETF *draft-martini-l2circuit-trans-mpls*.

- The VC type value for Ethernet is 0x0005.
- The VC type value for an Ethernet VLAN is 0x0004.

**Values** ether, vlan

**ether** — Defines the VC type as Ethernet. The **ethernet**, **vlan**, and **vpls** keywords are mutually exclusive. When the VC type is not defined then the default is Ethernet for spoke SDP bindings. Defining Ethernet is the same as executing **no vc-type** and restores the default VC type for the spoke SDP binding. (hex 5)

**vlan** — Defines the VC type as VLAN. The **ethernet**, **vlan**, and **vpls** keywords are mutually exclusive. When the VC type is not defined then the default is Ethernet for spoke SDP bindings. The VLAN VC-type requires at least one dot1Q tag within each encapsulated Ethernet packet transmitted to the far end.

**split-horizon-group** *group-name* — Specifies the name of the split horizon group to which the SDP belongs.

## egress

<b>Syntax</b>	<b>egress</b>
<b>Context</b>	config>service>vpls>mesh-sdp config>service>vpls>spoke-sdp
<b>Description</b>	This command configures the egress SDP context.

## ingress

<b>Syntax</b>	<b>ingress</b>
<b>Context</b>	config>service>vpls>mesh-sdp config>service>vpls>spoke-sdp
<b>Description</b>	This command configures the ingress SDP context.

## vc-label

<b>Syntax</b>	<b>[no] vc-label</b> <i>egress-vc-label</i>
<b>Context</b>	config>service>vpls>mesh-sdp>egress config>service>vpls>spoke-sdp>egress
<b>Description</b>	This command configures the egress VC label.

**Parameters** *vc-label* — A VC egress value that indicates a specific connection.  
**Values** 16 — 1048575

## vc-label

**Syntax** **[no] vc-label** *ingress-vc-label*  
**Context** config>service>vpls>mesh-sdp>ingress  
 config>service>vpls>spoke-sdp>ingress  
**Description** This command configures the ingress VC label.  
**Parameters** *vc-label* — A VC ingress value that indicates a specific connection.  
**Values** 2048 — 18431

## vlan-vc-tag

**Syntax** **vlan-vc-tag** *0..4094*  
**no vlan-vc-tag** [*0..4094*]  
**Context** config>service>vpls>spoke-sdp  
 config>service>vpls>mesh-sdp  
**Description** This command specifies an explicit dot1q value used when encapsulating to the SDP far end. When signaling is enabled between the near and far end, the configured dot1q tag can be overridden by a received TLV specifying the dot1q value expected by the far end. This signaled value must be stored as the remote signaled dot1q value for the binding. The provisioned local dot1q tag must be stored as the administrative dot1q value for the binding.  
 When the dot1q tag is not defined, the default value of zero is stored as the administrative dot1q value. Setting the value to zero is equivalent to not specifying the value.  
 The **no** form of this command disables the command  
**Default** no vlan-vc-tag  
**Parameters** *0..4094* — Specifies a valid VLAN identifier to bind an 802.1Q VLAN tag ID.

## avg-frame-overhead

**Syntax** **avg-frame-overhead** *percent*  
**no avg-frame-overhead**  
**Context** config>service>vpls>sap>egress>queue-override>queue  
**Description** This command configures the average frame overhead to define the average percentage that the offered load to a queue will expand during the frame encapsulation process before sending traffic on-the-wire. While the avg-frame-overhead value may be defined on any queue, it is only used by the system for queues that egress a Sonet or SDH port or channel. Queues operating on egress Ethernet

ports automatically calculate the frame encapsulation overhead based on a 20 byte per packet rule (8 bytes for preamble and 12 bytes for Inter-Frame Gap).

When calculating the frame encapsulation overhead for port scheduling purposes, the system determines the following values:

- **Offered-load** — The offered-load of a queue is calculated by starting with the queue depth in octets, adding the received octets at the queue and subtracting queue discard octets. The result is the number of octets the queue has available to transmit. This is the packet based offered-load.
- **Frame encapsulation overhead** — Using the avg-frame-overhead parameter, the frame encapsulation overhead is simply the queues current offered-load (how much has been received by the queue) multiplied by the avg-frame-overhead. If a queue had an offered load of 10000 octets and the avg-frame-overhead equals 10%, the frame encapsulation overhead would be  $10000 \times 0.1$  or 1000 octets.

For egress Ethernet queues, the frame encapsulation overhead is calculated by multiplying the number of offered-packets for the queue by 20 bytes. If a queue was offered 50 packets then the frame encapsulation overhead would be  $50 \times 20$  or 1000 octets.

- **Frame based offered-load** — The frame based offered-load is calculated by adding the offered-load to the frame encapsulation overhead. If the offered-load is 10000 octets and the encapsulation overhead was 1000 octets, the frame based offered-load would equal 11000 octets.
- **Packet to frame factor** — The packet to frame factor is calculated by dividing the frame encapsulation overhead by the queues offered-load (packet based). If the frame encapsulation overhead is 1000 octets and the offered-load is 10000 octets then the packet to frame factor would be  $1000 / 10000$  or 0.1. When in use, the avg-frame-overhead will be the same as the packet to frame factor making this calculation unnecessary.
- **Frame based CIR** — The frame based CIR is calculated by multiplying the packet to frame factor with the queues configured CIR and then adding that result to that CIR. If the queue CIR is set at 500 octets and the packet to frame factor equals 0.1, the frame based CIR would be  $500 \times 1.1$  or 550 octets.
- **Frame based within-cir offered-load** — The frame based within-cir offered-load is the portion of the frame based offered-load considered to be within the frame-based CIR. The frame based within-cir offered-load is the lesser of the frame based offered-load and the frame based CIR. If the frame based offered-load equaled 11000 octets and the frame based CIR equaled 550 octets, the frame based within-cir offered-load would be limited to 550 octets. If the frame based offered-load equaled 450 octets and the frame based CIR equaled 550 octets, the frame based within-cir offered-load would equal 450 octets (or the entire frame based offered-load).

As a special case, when a queue or associated intermediate scheduler is configured with a CIR-weight equal to 0, the system automatically sets the queues frame based within-cir offered-load to 0, preventing it from receiving bandwidth during the port schedulers within-cir pass.

- **Frame based PIR** — The frame based PIR is calculated by multiplying the packet to frame factor with the queues configured PIR and then adding the result to that PIR. If the queue PIR is set to 7500 octets and the packet to frame factor equals 0.1, the frame based PIR would be  $7500 \times 1.1$  or 8250 octets.
- **Frame based within-pir offered-load** — The frame based within-pir offered-load is the portion of the frame based offered-load considered to be within the frame based PIR. The frame based within-pir offered-load is the lesser of the frame based offered-load and the frame based PIR. If the frame based offered-load equaled 11000 octets and the frame based PIR equaled 8250 octets, the frame based within-pir offered-load would be limited to 8250 octets. If the frame based

offered-load equaled 7000 octets and the frame based PIR equaled 8250 octets, the frame based within-pir offered load would equal 7000 octets.

Port scheduler operation using frame transformed rates — The port scheduler uses the frame based rates to figure the maximum rates that each queue may receive during the within-cir and above-cir bandwidth allocation passes. During the within-cir pass, a queue may receive up to its frame based within-cir offered-load. The maximum it may receive during the above-cir pass is the difference between the frame based within-pir offered load and the amount of actual bandwidth allocated during the within-cir pass.

SAP and subscriber SLA-profile average frame overhead override — The average frame overhead parameter on a sap-egress may be overridden at an individual egress queue basis. On each SAP and within the sla-profile policy used by subscribers an avg-frame-overhead command may be defined under the queue-override context for each queue. When overridden, the queue instance will use its local value for the average frame overhead instead of the sap-egress defined overhead.

The **no** form of this command restores the average frame overhead parameter for the queue to the default value of 0 percent. When set to 0, the system uses the packet based queue statistics for calculating port scheduler priority bandwidth allocation. If the no avg-frame-overhead command is executed in a queue-override queue id context, the avg-frame-overhead setting for the queue within the sap-egress QoS policy takes effect.

<b>Default</b>	0
<b>Parameters</b>	<i>percent</i> — This parameter sets the average amount of packet-to-frame encapsulation overhead expected for the queue. This value is not used by the system for egress Ethernet queues.
<b>Values</b>	0 — 100

## queue-override

<b>Syntax</b>	<b>[no] queue-override</b>
<b>Context</b>	config>service>vpls>sap>egress config>service>vpls>sap>ingress
<b>Description</b>	This command enables the context to configure override values for the specified SAP egress or ingress QoS queue. These values override the corresponding ones specified in the associated SAP egress or ingress QoS policy.

## queue

<b>Syntax</b>	<b>[no] queue <i>queue-id</i></b>
<b>Context</b>	config>service>vpls>sap>egress>queue-override config>service>vpls>sap>ingress>queue-override
<b>Description</b>	This command specifies the ID of the queue whose parameters are to be overridden.
<b>Parameters</b>	<i>queue-id</i> — The queue ID whose parameters are to be overridden.
<b>Values</b>	1 — 32

## adaptation-rule

<b>Syntax</b>	<b>adaptation-rule</b> [ <b>pir</b> { <b>max</b>   <b>min</b>   <b>closest</b> }] [ <b>cir</b> { <b>max</b>   <b>min</b>   <b>closest</b> }] <b>no adaptation-rule</b>
<b>Context</b>	config>service>vpls>sap>egress>queue-override>queue config>service>vpls>sap>ingress>queue-override>queue
<b>Description</b>	<p>This command can be used to override specific attributes of the specified queue's adaptation rule parameters. The adaptation rule controls the method used by the system to derive the operational CIR and PIR settings when the queue is provisioned in hardware. For the CIR and PIR parameters individually, the system attempts to find the best operational rate depending on the defined constraint.</p> <p>The <b>no</b> form of the command removes any explicitly defined constraints used to derive the operational CIR and PIR created by the application of the policy. When a specific <b>adaptation-rule</b> is removed, the default constraints for <b>rate</b> and <b>cir</b> apply.</p>
<b>Default</b>	no adaptation-rule
<b>Parameters</b>	<p><b>pir</b> — The <b>pir</b> parameter defines the constraints enforced when adapting the PIR rate defined within the <b>queue queue-id rate</b> command. The <b>pir</b> parameter requires a qualifier that defines the constraint used when deriving the operational PIR for the queue. When the <b>rate</b> command is not specified, the default applies.</p> <p><b>max</b> — The <b>max</b> (maximum) option is mutually exclusive with the <b>min</b> and <b>closest</b> options. When <b>max</b> is defined, the operational PIR for the queue will be equal to or less than the administrative rate specified using the <b>rate</b> command.</p> <p><b>min</b> — The <b>min</b> (minimum) option is mutually exclusive with the <b>max</b> and <b>closest</b> options. When <b>min</b> is defined, the operational PIR for the queue will be equal to or greater than the administrative rate specified using the <b>rate</b> command.</p> <p><b>closest</b> — The <b>closest</b> parameter is mutually exclusive with the <b>min</b> and <b>max</b> parameter. When <b>closest</b> is defined, the operational PIR for the queue will be the rate closest to the rate specified using the <b>rate</b> command.</p> <p><b>cir</b> — The <b>cir</b> parameter defines the constraints enforced when adapting the CIR rate defined within the <b>queue queue-id rate</b> command. The <b>cir</b> parameter requires a qualifier that defines the constraint used when deriving the operational CIR for the queue. When the <b>cir</b> parameter is not specified, the default constraint applies.</p>

## cbs

<b>Syntax</b>	<b>cbs</b> <i>size-in-kbytes</i> <b>no cbs</b>
<b>Context</b>	config>service>vpls>sap>egress>queue-override>queue config>service>vpls>sap>ingress>queue-override>queue
<b>Description</b>	<p>This command can be used to override specific attributes of the specified queue's CBS parameters.</p> <p>It is permissible, and possibly desirable, to oversubscribe the total CBS reserved buffers for a given access port egress buffer pool. Oversubscription may be desirable due to the potential large number of</p>

service queues and the economy of statistical multiplexing the individual queue's CBS settings into the defined reserved total.

When oversubscribing the reserved total, it is possible for a queue depth to be lower than its CBS setting and still not receive a buffer from the buffer pool for an ingress frame. As more queues are using their CBS buffers and the total in use exceeds the defined reserved total, essentially the buffers are being removed from the shared portion of the pool without the shared in use average and total counts being decremented. This can affect the operation of the high and low priority RED slopes on the pool, causing them to miscalculate when to start randomly drop packets.

If the CBS value is larger than the MBS value, an error will occur, preventing the CBS change.

The **no** form of this command returns the CBS size to the default value.

<b>Default</b>	no cbs
<b>Parameters</b>	<i>size-in-kbytes</i> — The size parameter is an integer expression of the number of kilobytes reserved for the queue. If a value of 10KBytes is desired, enter the value 10. A value of 0 specifies that no reserved buffers are required by the queue (a minimal reserved size can still be applied for scheduling purposes).
<b>Values</b>	0 — 131072 or default

## high-prio-only

<b>Syntax</b>	<b>high-prio-only</b> <i>percent</i> <b>no high-prio-only</b>
<b>Context</b>	config>service>vpls>sap>egress>queue-override>queue config>service>vpls>sap>ingress>queue-override>queue
<b>Description</b>	<p>This command can be used to override specific attributes of the specified queue's high-prio-only parameters. The <b>high-prio-only</b> command configures the percentage of buffer space for the queue, used exclusively by high priority packets.</p> <p>The priority of a packet can only be set in the SAP ingress QoS policy and is only applicable on the ingress queues for a SAP. The <b>high-prio-only</b> parameter is used to override the default value derived from the <b>network-queue</b> command.</p> <p>The defined <b>high-prio-only</b> value cannot be greater than the MBS size of the queue. Attempting to change the MBS to a value smaller than the high priority reserve will generate an error and fail execution. Attempting to set the <b>high-prio-only</b> value larger than the current MBS size will also result in an error and fail execution.</p> <p>The <b>no</b> form of this command restores the default high priority reserved size.</p>
<b>Parameters</b>	<p><i>percent</i> — The <i>percent</i> parameter is the percentage reserved for high priority traffic on the queue. If a value of 10KBytes is desired, enter the value 10. A value of 0 specifies that none of the MBS of the queue will be reserved for high priority traffic. This does not affect RED slope operation for packets attempting to be queued.</p> <p><b>Values</b> 0 — 100   default</p>

## mbs

<b>Syntax</b>	<b>mbs</b> { <i>size-in-kbytes</i>   <b>default</b> } <b>no mbs</b>
<b>Context</b>	config>service>vpls>sap>egress>queue-override>queue
<b>Description</b>	<p>This command can be used to override specific attributes of the specified queue's MBS parameters. The MBS is a mechanism to override the default maximum size for the queue.</p> <p>The sum of the MBS for all queues on an egress access port can oversubscribe the total amount of buffering available. When congestion occurs and buffers become scarce, access to buffers is controlled by the RED slope a packet is associated with. A queue that has not exceeded its MBS size is not guaranteed that a buffer will be available when needed or that the packets RED slope will not force the discard of the packet. Setting proper CBS parameters and controlling CBS oversubscription is one major safeguard to queue starvation (when a queue does not receive its fair share of buffers). Another is properly setting the RED slope parameters for the needs of services on this port or channel.</p> <p>If the CBS value is larger than the MBS value, an error will occur, preventing the MBS change.</p> <p>The <b>no</b> form of this command returns the MBS size assigned to the queue.</p>
<b>Default</b>	default
<b>Parameters</b>	<p><i>size-in-kbytes</i> — The size parameter is an integer expression of the maximum number of kilobytes of buffering allowed for the queue. For a value of 100 kbps, enter the value 100. A value of 0 causes the queue to discard all packets.</p> <p><b>Values</b>      0 — 131072 or default</p>

## mbs

<b>Syntax</b>	<b>mbs</b> { <i>size-in-kbytes</i>   <b>default</b> } <b>no mbs</b>
<b>Context</b>	config>service>vpls>sap>ingress>queue-override>queue
<b>Description</b>	<p>This command can be used to override specific attributes of the specified queue's MBS parameters. The MBS value is used by a queue to determine whether it has exhausted all of its buffers while enqueueing packets. Once the queue has exceeded the amount of buffers allowed by MBS, all packets are discarded until packets have been drained from the queue.</p> <p>The sum of the MBS for all queues on an ingress access port can oversubscribe the total amount of buffering available. When congestion occurs and buffers become scarce, access to buffers is controlled by the RED slope a packet is associated with. A queue that has not exceeded its MBS size is not guaranteed that a buffer will be available when needed or that the packets RED slope will not force the discard of the packet. Setting proper CBS parameters and controlling CBS oversubscription is one major safeguard to queue starvation (when a queue does not receive its fair share of buffers). Another is properly setting the RED slope parameters for the needs of services on this port or channel.</p> <p>If the CBS value is larger than the MBS value, an error will occur, preventing the MBS change.</p> <p>The defined high-prio-only value cannot be greater than the MBS size of the queue. Attempting to change the MBS to a value smaller than the high priority reserve will generate an error and fail execu-</p>

tion. Attempting to set the high-prio-only value larger than the current MBS size will also result in an error and fail execution.

The **no** form of this command returns the MBS size assigned to the queue to the value.

**Default** default

**Parameters** *size-in-kbytes* — The size parameter is an integer expression of the maximum number of kilobytes of buffering allowed for the queue. For a value of 100 kbps, enter the value 100. A value of 0 causes the queue to discard all packets.

**Values** 0 — 131072 or default

## rate

**Syntax** **rate** *pir-rate* [**cir** *cir-rate*]  
**no rate**

**Context** config>service>vpls>sap>egress>queue-override>queue  
config>service>vpls>sap>ingress>queue-override>queue

**Description** This command can be used to override specific attributes of the specified queue's Peak Information Rate (PIR) and the Committed Information Rate (CIR) parameters.

The PIR defines the maximum rate that the queue can transmit packets out an egress interface (for SAP egress queues). Defining a PIR does not necessarily guarantee that the queue can transmit at the intended rate. The actual rate sustained by the queue can be limited by oversubscription factors or available egress bandwidth.

The CIR defines the rate at which the system prioritizes the queue over other queues competing for the same bandwidth. In-profile packets are preferentially queued by the system at egress and at subsequent next hop nodes where the packet can traverse. To be properly handled as in- or out-of-profile throughout the network, the packets must be marked accordingly for profiling at each hop.

The CIR can be used by the queue's parent commands *cir-level* and *cir-weight* parameters to define the amount of bandwidth considered to be committed for the child queue during bandwidth allocation by the parent scheduler.

The **rate** command can be executed at anytime, altering the PIR and CIR rates for all queues created through the association of the SAP egress QoS policy with the *queue-id*.

The **no** form of the command returns all queues created with the *queue-id* by association with the QoS policy to the default PIR and CIR parameters (**max**, 0).

**Default** **rate max cir 0** — The **max** default specifies the amount of bandwidth in kilobits per second (thousand bits per second). The **max** value is mutually exclusive to the **pir-rate** value.

**Parameters** *pir-rate* — Defines the administrative PIR rate, in kilobits, for the queue. When the **rate** command is executed, a valid PIR setting must be explicitly defined. When the **rate** command has not been executed, the default PIR of **max** is assumed.

Fractional values are not allowed and must be given as a positive integer.



The actual PIR rate is dependent on the queue's **adaptation-rule** parameters and the actual hardware where the queue is provisioned.

**Values** 1 — 100000000

**Default** max

*cir-rate* — The **cir** parameter overrides the default administrative CIR used by the queue. When the **rate** command is executed, a CIR setting is optional. When the **rate** command has not been executed or the **cir** parameter is not explicitly specified, the default CIR (0) is assumed. Fractional values are not allowed and must be given as a positive integer.

**Values** 0 — 100000000, max

**Default** 0

## scheduler-override

**Syntax** [no] **scheduler-override**

**Context** config>service>vpls>sap>egress  
config>service>vpls>sap>ingress

**Description** This command specifies the set of attributes whose values have been overridden via management on this virtual scheduler. Clearing a given flag will return the corresponding overridden attribute to the value defined on the SAP's ingress scheduler policy.

## scheduler

**Syntax** **scheduler** *scheduler-name*  
**no scheduler** *scheduler-name*

**Context** config>service>vpls>sap>egress>sched-override

**Description** This command can be used to override specific attributes of the specified scheduler name.

A scheduler defines a bandwidth control that limits each child (other schedulers and queues) associated with the scheduler. Scheduler objects are created within the hierarchical tiers of the policy. It is assumed that each scheduler created will have queues or other schedulers defined as child associations. The scheduler can be a child (take bandwidth from a scheduler in a higher tier, except for schedulers created in tier 1). A total of 32 schedulers can be created within a single scheduler policy with no restriction on the distribution between the tiers.

Each scheduler must have a unique name within the context of the scheduler policy; however the same name can be reused in multiple scheduler policies. If *scheduler-name* already exists within the policy tier level (regardless of the inclusion of the keyword create), the context changes to that scheduler name for the purpose of editing the scheduler parameters. Modifications made to an existing scheduler are executed on all instantiated schedulers created through association with the policy of the edited scheduler. This can cause queues or schedulers to become orphaned (invalid parent association) and adversely affect the ability of the system to enforce service level agreements (SLAs).

If the *scheduler-name* exists within the policy on a different tier (regardless of the inclusion of the keyword **create**), an error occurs and the current CLI context will not change.

If the *scheduler-name* does not exist in this or another tier within the scheduler policy, it is assumed that an attempt is being made to create a scheduler of that name. The success of the command execution is dependent on the following:

1. The maximum number of schedulers has not been configured.
2. The provided *scheduler-name* is valid.
3. The **create** keyword is entered with the command if the system is configured to require it (enabled in the **environment create** command).

When the maximum number of schedulers has been exceeded on the policy, a configuration error occurs and the command will not execute, nor will the CLI context change.

If the provided scheduler-name is invalid according to the criteria below, a name syntax error will occur, the command will not execute, and the CLI context will not change.

**Parameters** *scheduler-name* — The name of the scheduler.

**Values** Valid names consist of any string up to 32 characters long composed of printable, 7-bit ASCII characters excluding double quotes. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

**Default** **None.** Each scheduler must be explicitly created.

*create* — This optional keyword explicitly specifies that it is acceptable to create a scheduler with the given *scheduler-name*. If the **create** keyword is omitted, **scheduler-name** is not created when the system environment variable **create** is set to true. This safeguard is meant to avoid accidental creation of system objects (such as schedulers) while attempting to edit an object with a mistyped name or ID. The keyword has no effect when the object already exists.

## rate

<b>Syntax</b>	<b>rate</b> <i>pir-rate</i> [ <i>cir cir-rate</i> ] <b>no rate</b>
<b>Context</b>	config>service>vpls>sap>egress>sched-override>scheduler
<b>Description</b>	<p>This command can be used to override specific attributes of the specified scheduler rate. The <b>rate</b> command defines the maximum bandwidth that the scheduler can offer its child queues or schedulers. The maximum rate is limited to the amount of bandwidth the scheduler can receive from its parent scheduler. If the scheduler has no parent, the maximum rate is assumed to be the amount available to the scheduler. When a parent is associated with the scheduler, the CIR parameter provides the amount of bandwidth to be considered during the parent scheduler's 'within CIR' distribution phase.</p> <p>The actual operating rate of the scheduler is limited by bandwidth constraints other than its maximum rate. The scheduler's parent scheduler may not have the available bandwidth to meet the scheduler's needs or the bandwidth available to the parent scheduler could be allocated to other child schedulers or child queues on the parent based on higher priority. The children of the scheduler may not need the maximum rate available to the scheduler due to insufficient offered load or limits to their own maximum rates.</p>

When a scheduler is defined without specifying a rate, the default rate is **max**. If the scheduler is a root scheduler (no parent defined), the default maximum rate must be changed to an explicit value. Without this explicit value, the scheduler will assume that an infinite amount of bandwidth is available and allow all child queues and schedulers to operate at their maximum rates.

The **no** form of this command returns all queues created with this *queue-id* by association with the QoS policy to the default PIR and CIR parameters.

#### Parameters

*pir-rate* — The **pir** parameter accepts a step multiplier value that specifies the multiplier used to determine the PIR rate at which the queue will operate. A value of 0 to 100000000 or the keyword **max** is accepted. Any other value will result in an error without modifying the current PIR rate.

To calculate the actual PIR rate, the rate described by the queue's **rate** is multiplied by the *pir-rate*.

The SAP ingress context for PIR is independent of the defined forwarding class (fc) for the queue. The default **pir** and definable range is identical for each class. The PIR in effect for a queue defines the maximum rate at which the queue will be allowed to forward packets in a given second, thus shaping the queue's output.

The PIR parameter for SAP ingress queues do not have a negate (**no**) function. To return the queue's PIR rate to the default value, that value must be specified as the PIR value.

**Values** 1 — 100000000, **max**

**Default** **max**

*cir cir-rate* — The **cir** parameter accepts a step-multiplier value that specifies the multiplier used to determine the CIR rate at which the queue will operate. A value of 0 to 250 or the keyword **max** is accepted. Any other value will result in an error without modifying the current CIR rate.

To calculate the actual CIR rate, the rate described by the **rate pir pir-rate** is multiplied by the *cir cir-rate*. If the **cir** is set to **max**, then the CIR rate is set to infinity.

The SAP ingress context for CIR is dependent on the defined forwarding class (fc) for the queue. The default CIR and definable range is different for each class. The CIR in effect for a queue defines both its profile (in or out) marking level as well as the relative importance compared to other queues for scheduling purposes during congestion periods.

**Values** 0 — 100000000, **max**, **sum**

**Default** **sum**

## match-qinq-dot1p

<b>Syntax</b>	<b>match-qinq-dot1p {top   bottom}</b> <b>no match-qinq-dot1p</b>
<b>Context</b>	config>service>vpls>sap>ingress config>service>ies>if>sap>ingress
<b>Description</b>	This command specifies which Dot1Q tag position Dot1P bits in a QinQ encapsulated packet should be used to evaluate Dot1P QoS classification.

The **match-qinq-dot1p** command allows the top or bottom PBits to be used when evaluating the applied sap-ingress QoS policy's Dot1P entries. The **top** and **bottom** keywords specify which position should be evaluated for QinQ encapsulated packets.

The **no** form of the command restores the default dot1p evaluation behavior for the SAP.

By default, the bottom-most service delineating Dot1Q tag's Dot1P bits are used. [Table 5](#) defines the default behavior for Dot1P evaluation when the **match-qinq-dot1p** command is not executed.

**Table 5: Default QinQ and TopQ SAP Dot1P Evaluation**

Port / SAP Type	Existing Packet Tags	PBits Used for Match
Null	None	None
Null	Dot1P (VLAN-ID 0)	Dot1P PBits
Null	Dot1Q	Dot1Q PBits
Null	TopQ BottomQ	TopQ PBits
Null	TopQ (No BottomQ)	TopQ PBits
Dot1Q	None (Default SAP)	None
Dot1Q	Dot1P (Default SAP VLAN-ID 0)	Dot1P PBits
Dot1Q	Dot1Q	Dot1Q PBits
QinQ / TopQ	TopQ	TopQ PBits
QinQ / TopQ	TopQ BottomQ	TopQ PBits
QinQ / QinQ	TopQ BottomQ	BottomQ PBits

**Default** no match-qinq-dot1p (no filtering based on p-bits)  
(top or bottom must be specified to override the default QinQ dot1p behavior)

**Parameters** **top** — The top parameter is mutually exclusive to the bottom parameter. When the top parameter is specified, the top most PBits are used (if existing) to match any dot1p dot1p-value entries. [Table 6](#) defines the dot1p evaluation behavior when the top parameter is specified.

**Table 6: Top Position QinQ and TopQ SAP Dot1P Evaluation**

Port / SAP Type	Existing Packet Tags	PBits Used for Match
Null	None	None
Null	Dot1P (VLAN-ID 0)	Dot1P PBits
Null	Dot1Q	Dot1Q PBits
Null	TopQ BottomQ	TopQ PBits
Null	TopQ (No BottomQ)	TopQ PBits
Dot1Q	None (Default SAP)	None

**Table 6: Top Position QinQ and TopQ SAP Dot1P Evaluation (Continued)**

Port / SAP Type	Existing Packet Tags	PBits Used for Match
Dot1Q	Dot1P (Default SAP VLAN-ID 0)	Dot1P PBits
Dot1Q	Dot1Q	Dot1Q PBits
QinQ / TopQ	TopQ	TopQ PBits
QinQ / TopQ	TopQ BottomQ	TopQ PBits
QinQ / QinQ	TopQ BottomQ	TopQ PBits

**bottom** — The bottom parameter is mutually exclusive to the top parameter. When the bottom parameter is specified, the bottom most PBits are used (if existing) to match any dot1p dot1p-value entries. [Table 7](#) defines the dot1p evaluation behavior when the bottom parameter is specified.

**Table 7: Bottom Position QinQ and TopQ SAP Dot1P Evaluation**

Port / SAP Type	Existing Packet Tags	PBits Used for Match
Null	None	None
Null	Dot1P (VLAN-ID 0)	Dot1P PBits
Null	Dot1Q	Dot1Q PBits
Null	TopQ BottomQ	BottomQ PBits
Null	TopQ (No BottomQ)	TopQ PBits
Dot1Q	None (Default SAP)	None
Dot1Q	Dot1P (Default SAP VLAN-ID 0)	Dot1P PBits
Dot1Q	Dot1Q	Dot1Q PBits
QinQ / TopQ	TopQ	TopQ PBits
QinQ / TopQ	TopQ BottomQ	BottomQ PBits
QinQ / QinQ	TopQ BottomQ	BottomQ PBits

## discard-unknown-source

<b>Syntax</b>	<b>[no] discard-unknown-source</b>
<b>Context</b>	config>service>vpls>sap config>service>vpls>spoke-sdp
<b>Description</b>	When this command is enabled, packets received on a SAP or on a spoke SDP with an unknown source MAC address will be dropped only if the maximum number of MAC addresses for that SAP or spoke SDP (see <a href="#">max-nbr-mac-addr</a> on <a href="#">page 126</a> ) has been reached. If max-nbr-mac-addr has not been set for the SAP or spoke SDP, enabling discard-unknown-source has no effect.

When disabled, the packets are forwarded based on the destination MAC addresses.

The **no** form of this command causes packets with an unknown source MAC addresses to be forwarded by destination MAC addresses in VPLS.

**Default**    **no discard-unknown**

*qtag1, qtag2* — Specifies the encapsulation value used to identify the SAP on the port or sub-port. If this parameter is not specifically defined, the default value is 0.

**Values**      qtag1:            0 — 4094  
                   qtag2 :            \*, 0 — 4094

The values depends on the encapsulation type configured for the interface. The following table describes the allowed values for the port and encapsulation types..

Port Type	Encap-Type	Allowed Values	Comments
Ethernet	Null	0	The SAP is identified by the port.
Ethernet	Dot1q	0 - 4094	The SAP is identified by the 802.1Q tag on the port. Note that a 0 qtag1 value also accepts untagged packets on the dot1q port.
Ethernet	QinQ	qtag1: 0 - 4094 qtag2: 0 - 4094	The SAP is identified by two 802.1Q tags on the port. Note that a 0 qtag1 value also accepts untagged packets on the dot1q port.
SONET/SDH	IPCP	-	The SAP is identified by the channel. No BCP is deployed and all traffic is IP.
SONET/SDH	BCP-Null	0	The SAP is identified with a single service on the channel. Tags are assumed to be part of the customer packet and not a service delimiter.
SONET/SDH	BCP-Dot1q	0 - 4094	The SAP is identified by the 802.1Q tag on the channel.

**create** — Keyword used to create a SAP instance.

**split-horizon-group** *group-name* — Specifies the name of the split horizon group to which the SAP belongs.

## bpdu-translation

<b>Syntax</b>	<b>bpdu-translation {auto   pvst   stp}</b> <b>no bpdu-translation</b>
<b>Context</b>	config>service>vpls>sap config>service>vpls>spoke-sdp
<b>Description</b>	This command enables the translation of BPDUs to a given format, meaning that all BPDUs transmitted on a given SAP or spoke SDP will have a specified format.  The <b>no</b> form of this command reverts to the default setting.
<b>Default</b>	no bpdu-translation
<b>Parameters</b>	<b>auto</b> — Specifies that appropriate format will be detected automatically, based on type of bpdus received on such port.  <b>pvst</b> — Specifies the BPDU-format as PVST. Note that the correct VLAN tag is included in the payload (depending on encapsulation value of outgoing SAP).  <b>stp</b> — Specifies the BPDU-format as STP.

## l2pt-termination

<b>Syntax</b>	<b>[no] l2pt-termination</b>
<b>Context</b>	config>service>vpls>sap config>service>vpls>spoke-sdp
<b>Description</b>	This commands enables L2PT termination on a given SAP or spoke SDP. L2PT termination will be supported only for STP BPDUs. PDUs of other protocols will be discarded.  This feature can be enabled only if STP is disabled in the context of the given VPLS service.
<b>Default</b>	no l2pt-termination

## def-mesh-vc-id

<b>Syntax</b>	<b>[no] def-mesh-vc-id <i>vc-id</i></b>
<b>Context</b>	config>service>vpls
<b>Description</b>	This command configures the value used by each end of a tunnel to identify the VC. If this command is not configured, then the service ID value is used as the VC-ID.  This VC-ID is used instead of a label to identify a virtual circuit. The VC-ID is significant between peer ESS on the same hierarchical level. The value of a VC-ID is conceptually independent from the value of the label or any other datalink specific information of the VC.  The <b>no</b> form of this command disables the VC-ID.
<b>Default</b>	none

**Values** 1 — 4294967295

## mac-move

<b>Syntax</b>	<b>[no] mac-move</b>
<b>Context</b>	config>service>vpls
<b>Description</b>	<p>This command enables the context to configure MAC move attributes. A sustained high re-learn rate can be a sign of a loop somewhere in the VPLS topology. Typically, STP detects loops in the topology, but for those networks that do not run STP, the mac-move feature is an alternative way to protect your network against loops.</p> <p>When enabled in a VPLS, <b>mac-move</b> monitors the re-learn rate of each MAC. If the rate exceeds the configured maximum allowed limit, it disables the SAP where the source MAC was last seen. The SAP can be disabled permanently (until a <b>shutdown/no shutdown</b> command is executed) or for a length of time that grows linearly with the number of times the given SAP was disabled. You have the option of marking a SAP as non-blockable in the <b>config&gt;service&gt;vpls&gt;sap&gt;limit-mac-move</b> or <b>config&gt;service&gt;vpls&gt;spoke-sdp&gt;limit-mac-move</b> contexts, see <a href="#">page 92</a>. This means that when the re-learn rate has exceeded the limit, another (blockable) SAP will be disabled instead.</p> <p>The <b>mac-move</b> command enables the feature at the service level for SAPs and spoke SDPs, as only those objects can be blocked by this feature. Mesh SDPs are never blocked, but their re-learn rates (sap-to-mesh/spoke-to-mesh or vice versa) are still measured.</p> <p>The operation of this feature is the same on the SAP and spoke SDP. For example, if a MAC address moves from SAP to SAP, from SAP to spoke SDP, or between spoke SDPs, one will be blocked to prevent thrashing. If the MAC address moves between a SAP and mesh SDP or spoke SDP and mesh SDP combinations, the respective SAP or spoke SDP will be blocked.</p> <p>The re-learn rate is computed as the number of times a MAC moves in a 5 second interval. Therefore, the fastest a loop can be detected and broken is 5 seconds.</p> <p>The <b>no</b> form of this command disables MAC move.</p>
<b>Default</b>	not enabled

## mac-protect

<b>Syntax</b>	<b>mac-protect</b>
<b>Context</b>	config>service>vpls
<b>Description</b>	<p>This command indicates whether or not this MAC is protected. When enabled, the agent will protect the MAC from being learned or re-learned on a SAP that has restricted learning enabled.</p>
<b>Default</b>	disabled



## mac

<b>Syntax</b>	<b>[no] mac</b> <i>ieee-address</i>
<b>Context</b>	config>service>vpls>mac-protect
<b>Description</b>	This command specifies the 48-bit IEEE 802.3 MAC address.
<b>Parameters</b>	<i>ieee-address</i> — Specifies the 48-bit MAC address in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee, and ff are hexadecimal numbers.

## mac-subnet-length

<b>Syntax</b>	<b>mac-subnet-length</b> <i>subnet-length</i> <b>no mac-subnet-length</b>
<b>Context</b>	config>service>vpls
<b>Description</b>	<p>This command specifies the number of bits to be considered when performing MAC learning (MAC source) and MAC switching (MAC destination). Specifically, this value identifies how many bits, starting from the beginning of the MAC address are used. For example, if the mask-value of 28 is used, MAC learning will only do a lookup for the first 28 bits of the source MAC address when comparing with existing FIB entries. Then, it will install the first 28 bits in the FIB while zeroing out the last 20 bits of the MAC address. When performing switching in the reverse direction, only the first 28 bits of the destination MAC address will be used to perform a FIB lookup to determine the next hop.</p> <p>The <b>no</b> form of this command switches back to full MAC lookup.</p>
<b>Parameters</b>	<i>subnet-length</i> — Specifies the number of bits to be considered when performing MAC learning or MAC switching.
<b>Values</b>	24 — 48

## mcr-default-gtw

<b>Syntax</b>	<b>mcr-default-gtw</b>
<b>Context</b>	config>service>vpls
<b>Description</b>	This command enables the context to configure the default gateway information when using Dual Homing in L2-TPSDA. The IP and MAC address of the default gateway used for subscribers on an L2 MC-Ring are configured in this context. After a ring heals or fails, the system will send out a gratuitous ARP on an active ring SAP in order to attract traffic from subscribers on the ring with connectivity to that SAP.

## ip

<b>Syntax</b>	<b>ip</b> <i>address</i> <b>no ip</b>
<b>Context</b>	config>service>vpls>mcr-default-gtw
<b>Description</b>	This command relates to a system configured for Dual Homing in L2-TPSDA. It defines the IP address used when the system sends out a gratuitous ARP on an active SAP after a ring heals or fails in order to attract traffic from subscribers on the ring with connectivity to that SAP.
<b>Default</b>	no ip
<b>Parameters</b>	<i>address</i> — Specifies the IP address in a.b.c.d. format.

## mac

<b>Syntax</b>	<b>mac</b> <i>ieee-address</i> <b>no mac</b>
<b>Context</b>	config>service>vpls>mcr-default-gtw
<b>Description</b>	This command relates to a system configured for Dual Homing in L2-TPSDA. It defines the MAC address used when the system sends out a gratuitous ARP on an active SAP after a ring heals or fails in order to attract traffic from subscribers on the ring with connectivity to that SAP.
<b>Default</b>	no mac
<b>Parameters</b>	<i>ieee-address</i> — Specifies the address in xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx format (cannot be all zeros).

## move-frequency

<b>Syntax</b>	<b>move-frequency</b> <i>frequency</i> <b>no move-frequency</b>
<b>Context</b>	config>service>vpls>mac-move
<b>Description</b>	<p>This object indicates the maximum rate at which MAC's can be re-learned in the VPLS service, before the SAP where the moving MAC was last seen is automatically disabled in order to protect the system against undetected loops or duplicate MAC's.</p> <p>The rate is computed as the maximum number of re-learns allowed in a 5 second interval. For example, the default rate of 10 relearns per second corresponds to 50 relearns in a 5 second period.</p> <p>The <b>no</b> form of the command reverts to the default value.</p>
<b>Default</b>	2 (when mac-move is enabled)
<b>Parameters</b>	<i>frequency</i> — Specifies the rate, in 5-second intervals for the maximum number of relearns.
<b>Values</b>	1 — 100

## retry-timeout

<b>Syntax</b>	<b>retry-timeout</b> <i>timeout</i> <b>no retry-timeout</b>
<b>Context</b>	config>service>vpls>mac-move
<b>Description</b>	<p>This objects indicates the time in seconds to wait before a SAP that has been disabled after exceeding the maximum relearn rate is reenabled.</p> <p>A zero value indicates that the SAP will not be automatically re-enabled after being disabled. If, after the SAP is reenabled it is disabled again, the effective retry timeout is doubled in order to avoid thrashing.</p> <p>The <b>no</b> form of the command reverts to the default value.</p>
<b>Default</b>	10 (when mac-move is enabled)
<b>Parameters</b>	<p><i>timeout</i> — Specifies the time, in seconds, to wait before a SAP that has been disabled after exceeding the maximum relearn rate is reenabled.</p> <p><b>Values</b>      0 — 120</p>

## mfib-table-high-wmark

<b>Syntax</b>	<b>[no] mfib-table-high-wmark</b> <i>high-water-mark</i>
<b>Context</b>	config>service>vpls
<b>Description</b>	This command specifies the multicast FIB high watermark. When the percentage filling level of the multicast FIB exceeds the configured value, a trap is generated and/or a log entry is added.
<b>Parameters</b>	<p><i>high-water-mark</i> — Specifies the multicast FIB high watermark as a percentage.</p> <p><b>Values</b>      1 — 100</p> <p><b>Default</b>      95%</p>

## mfib-table-low-wmark

<b>Syntax</b>	<b>[no] mfib-table-low-wmark</b> <i>low-water-mark</i>
<b>Context</b>	config>service>vpls
<b>Description</b>	This command specifies the multicast FIB low watermark. When the percentage filling level of the Multicast FIB drops below the configured value, the corresponding trap is cleared and/or a log entry is added.
<b>Parameters</b>	<p><i>low-water-mark</i> — Specifies the multicast FIB low watermark as a percentage.</p> <p><b>Values</b>      1 — 100</p> <p><b>Default</b>      90%</p>

## mfib-table-size

<b>Syntax</b>	<b>mfib-table-size</b> <i>size</i> <b>no mfib-table-size</b>
<b>Context</b>	config>service>vpls
<b>Description</b>	<p>This command specifies the maximum number of (s,g) entries in the multicast forwarding database (MFIB) for this VPLS instance.</p> <p>The <i>mfib-table-size</i> parameter specifies the maximum number of multicast database entries for both learned and static multicast addresses for the VPLS instance.</p> <p>When a table-size limit is set on the mfib of a service which is lower than the current number of dynamic entries present in the mfib then the number of entries remains above the limit.</p> <p>The <b>no</b> form of this command removes the configured maximum MFIB table size.</p>
<b>Default</b>	none
<b>Parameters</b>	<p><i>size</i> — The maximum number of (s,g) entries allowed in the Multicast FIB.</p> <p><b>Values</b> 1 — 16383</p>

## send-flush-on-failure

<b>Syntax</b>	<b>[no] send-flush-on-failure</b>
<b>Context</b>	config>service>vpls
<b>Description</b>	<p>This command enables sending out “flush-all-from-ME” messages to all LDP peers included in affected VPLS, in the event of physical port failures or “oper-down” events of individual SAPs. This feature provides an LDP-based mechanism for recovering a physical link failure in a dual-homed connection to a VPLS service. This method provides an alternative to RSTP solutions where dual homing redundancy and recovery, in the case of link failure, is resolved by RSTP running between a PE router and CE devices.</p> <p>This feature cannot be enabled on management VPLS.</p>
<b>Default</b>	no send-flush-on-failure

## restrict-protected-src

<b>Syntax</b>	<b>[no] restrict-protected-src</b>
<b>Context</b>	config>service>vpls>split-horizon-group config>service>vpls>sap
<b>Description</b>	<p>This command indicates how the agent will handle relearn requests for protected MAC addresses. When enabled, requests to relearn a protected MAC address will be ignored, and the SAP where the protected source MAC was seen will be brought operationally down.</p>
<b>Default</b>	no restrict-protected-src

## restrict-unprotected-dst

<b>Syntax</b>	<b>[no] restrict-unprotected-dst</b>
<b>Context</b>	config>service>vpls>split-horizon-group config>service>vpls>sap
<b>Description</b>	This command indicates how the system will forward packets destined to an unprotected MAC address. When enabled, packets destined to an unprotected MAC address will be dropped.
<b>Default</b>	no restrict-unprotected-dst

---

## DHCP Commands

### dhcp

<b>Syntax</b>	<b>dhcp</b>
<b>Context</b>	<pre>config&gt;service&gt;vpls&gt;sap config&gt;service&gt;vpls&gt;spoke-sdp config&gt;service&gt;vpls&gt;mesh-sdp config&gt;service&gt;ies&gt;interface config&gt;service&gt;vpn&gt;interface config&gt;service&gt;ies&gt;subscriber-interface&gt;group-interface config&gt;service&gt;ies&gt;subscriber-interface</pre>
<b>Description</b>	This command enables the context to configure DHCP parameters.

### gi-address

<b>Syntax</b>	<b>gi-address</b> <i>ip-address</i> [ <i>src-ip-addr</i> ] <b>no gi-address</b>
<b>Context</b>	<pre>config&gt;service&gt;ies&gt;if&gt;dhcp config&gt;service&gt;vpn&gt;interface&gt;dhcp config&gt;service&gt;ies&gt;subscriber-interface&gt;grp-if&gt;dhcp config&gt;service&gt;ies&gt;subscriber-interface&gt;dhcp</pre>
<b>Description</b>	<p>This command configures the gateway interface address for the DHCP relay. A subscriber interface can include multiple group interfaces with multiple SAPs. The GI address is needed, when the router functions as a DHCP relay, to distinguish between the different subscriber interfaces and potentially between the group interfaces defined.</p> <p>By default, the GI address used in the relayed DHCP packet is the primary IP address of a normal IES interface. Specifying the GI address allows the user to choose a secondary address. For group interfaces a GI address must be specified under the group interface DHCP context or subscriber-interface DHCP context in order for DHCP to function.</p>
<b>Default</b>	no gi-address
<b>Parameters</b>	<p><i>ip-address</i> — Specifies the host IP address to be used for DHCP relay packets.</p> <p><i>src-ip-address</i> — Specifies that this GI address is to be the source IP address for DHCP relay packets.</p>

## lease-populate

<b>Syntax</b>	<b>lease-populate</b> [ <i>nbr-of-leases</i> ] <b>lease-populate</b> [ <i>nbr-of-leases</i> ] <b>I2-header</b> [ <i>mac ieee-address</i> ] <b>no lease-populate</b>
<b>Context</b>	config>subscr-mgmt>msap-policy>vpls-only>dhcp config>service>vpls>sap>dhcp config>service>ies>interface>dhcp config>service>vprn>interface>dhcp config>service>ies>sub-if>grp-if>dhcp config>service>vprn>sub-if>grp-if>dhcp config>service>vprn>sub-if>dhcp
<b>Description</b>	<p>This command enables and disables dynamic host DHCPv4 lease state management for SAPs.</p> <p>For VPLS, DHCP snooping must be explicitly enabled (using the <b>snoop</b> command) at all points where DHCP messages requiring snooping enter the VPLS instance (both from the DHCP server and from the subscribers). Lease state information is extracted from snooped DHCP ACK messages to populate lease state table entries for the SAP.</p> <p>The optional number-of-entries parameter defines the number lease state table entries allowed.</p> <ul style="list-style-type: none"> <li>• for this SAP in case of a VPLS service</li> <li>• for this interface in case of an IES or VPRN interface</li> <li>• for each SAP in case of an IES or VPRN group-interface</li> <li>• for this interface in case of an IES or VPRN retail subscriber-interface</li> </ul> <p>If number-of-entries is omitted, only a single entry is allowed. Once the maximum number of entries has been reached, subsequent lease state entries are not allowed and subsequent DHCP ACK messages are discarded.</p> <p>The retained lease state information representing dynamic hosts may be used to:</p> <ul style="list-style-type: none"> <li>• Populate a SAP based anti-spoof filter table to provide dynamic anti-spoof filtering. If the system is unable to populate the dynamic host information in the anti-spoof filter table on the SAP, the DHCP ACK message must be discarded without adding new lease state entry or updating an existing lease state entry.</li> <li>• Populate the system's ARP cache based on the arp-populate configuration. Applicable to IES and VPRN interfaces or group-interfaces.</li> <li>• Populate managed entries into a VPLS forwarding database. VPLS forwarding database population is an implicit feature that automatically places the dynamic host's MAC address into the VPLS FDB. When a dynamic host's MAC address is placed in the lease state table, it will automatically be populated into the VPLS forwarding database associated with the SAP on which the host is learned. The dynamic host MAC address will override any static MAC entries using the same MAC and prevent dynamic learning of the MAC on another interface. Existing static MAC entries with the same MAC address as the dynamic host are marked as inactive but not deleted. If all entries in the lease state table associated with the MAC address are removed, the static MAC may be populated. New static MAC definitions for the VPLS instance may be created while a dynamic host exists associated with the static MAC address</li> <li>• Generate dynamic ARP replies if <b>arp-reply-agent</b> is enabled. Applicable to VPLS service SAPs</li> </ul>

<b>Default</b>	no lease-populate
<b>Parameters</b>	<i>nbr-of-leases</i> — Specifies the number of DHCPv4 leases allowed.
<b>Values</b>	1 — 32767 1 — 65535 (chassis-mode d, SF/CPM-4 or later) 1 — 262143 (chassis-mode d, SF/CPM-4 or later, retail subscriber interfaces only)
<b>i2-header</b>	— Indicates a mode of operation where anti-spoof entry associated with the given DHCP state is created based on the MAC address from the Layer 2 header. The Layer 2 header flag is not set by default. This parameter is only applicable for group-interfaces.
<b>mac</b>	— Specifies that the provisioned ieee-address will be used in the anti-spoofing entries for this SAP. The parameter may be changed mid-session. Existing sessions will not be re-programmed unless a tools perform command is issues for the lease. This parameter is only applicable for group-interfaces.

## match-circuit-id

<b>Syntax</b>	<b>[no] match-circuit-id</b>
<b>Context</b>	config>service>ies>sub-if>grp-if>dhcp
<b>Description</b>	<p>This command enables Option 82 circuit ID on relayed DHCP packet matching.</p> <p>For Routed CO, the group interface DHCP relay process is stateful. When packets are relayed to the server the virtual router ID, transaction ID, SAP ID, and client hardware MAC address of the relayed packet are tracked. When get a response back from the server the virtual router ID, transaction ID, and client HW MAC address must be matched to determine the SAP on which to send the packet out. In some cases, the virtual router ID, transaction ID, and client HW MAC address are not guaranteed to be unique.</p> <p>When the <b>match-circuit-id</b> command is enabled this part of the key is used to guarantee correctness in our lookup. This is really only needed when dealing with an IP aware DSLAM that proxies the client HW mac address.</p>
<b>Default</b>	no match-circuit-id

## option

<b>Syntax</b>	<b>[no] option</b>
<b>Context</b>	config>service>vpls>sap>dhcp config>service>ies>if>dhcp config>service>vprn>if>dhcp config>service>ies>subscriber-interface>grp-if>dhcp
<b>Description</b>	<p>This command enables DHCP Option 82 (Relay Agent Information Option) parameters processing and enters the context for configuring Option 82 sub-options.</p> <p>The <b>no</b> form of this command returns the system to the default.</p>
<b>Default</b>	no option



## action

<b>Syntax</b>	<b>action {replace   drop   keep}</b> <b>no action</b>
<b>Context</b>	config>service>vpls>sap>dhcp>option config>service>ies>interface>dhcp>option config>service>vprn>interface>dhcp>option config>service>ies>subscriber-interface>grp-if>dhcp
<b>Description</b>	This command configures the Relay Agent Information Option (Option 82) processing. The <b>no</b> form of this command returns the system to the default value.
<b>Default</b>	The default is to keep the existing information intact.
<b>Parameters</b>	<b>replace</b> — In the upstream direction (from the user), the Option 82 field from the router is inserted in the packet (overwriting any existing Option 82 field). In the downstream direction (towards the user) the Option 82 field is stripped (in accordance with RFC 3046). <b>drop</b> — The DHCP packet is dropped if an Option 82 field is present, and a counter is incremented. <b>keep</b> — The existing information is kept in the packet and the router does not add any additional information. In the downstream direction the Option 82 field is not stripped and is forwarded towards the client.

## circuit-id

<b>Syntax</b>	<b>circuit-id [ascii-tuple   vlan-ascii-tuple]</b> <b>no circuit-id</b>
<b>Context</b>	config>service>vpls>sap>dhcp>option
<b>Description</b>	When enabled, the router sends an ASCII-encoded tuple in the <b>circuit-id</b> sub-option of the DHCP packet. This ASCII-tuple consists of the access-node-identifier, service-id, and SAP-ID, separated by “ ”. If disabled, the <b>circuit-id</b> sub-option of the DHCP packet will be left empty. The <b>no</b> form of this command returns the system to the default.
<b>Default</b>	circuit-id
<b>Parameters</b>	<b>ascii-tuple</b> — Specifies that the ASCII-encoded concatenated tuple consisting of the access-node-identifier, service-id, and interface-name is used. <b>vlan-ascii-tuple</b> — Specifies that the format will include VLAN-id and dot1p bits in addition to what is included in ascii-tuple already. The format is supported on dot1q and qinq ports only. Thus, when the option 82 bits are stripped, dot1p bits will be copied to the Ethernet header of an outgoing packet.

## circuit-id

<b>Syntax</b>	<b>circuit-id</b> [ascii-tuple   ifindex   sap-id   vlan-ascii-tuple] <b>no circuit-id</b>
<b>Context</b>	config>service>ies>if>dhcp>option
<b>Description</b>	<p>When enabled, the router sends an ASCII-encoded tuple in the <b>circuit-id</b> sub-option of the DHCP packet. This ASCII-tuple consists of the access-node-identifier, service-id, and SAP-ID, separated by “ ”.</p> <p>If disabled, the <b>circuit-id</b> sub-option of the DHCP packet will be left empty.</p> <p>The <b>no</b> form of this command returns the system to the default.</p>
<b>Default</b>	circuit-id
	<p><b>ascii-tuple</b> — Specifies that the ASCII-encoded concatenated tuple will be used which consists of the access-node-identifier, service-id, and interface-name, separated by “ ”.</p> <p><b>ifindex</b> — Specifies that the interface index will be used. (The If Index of a router interface can be displayed using the command <code>show&gt;router&gt;interface&gt;detail</code>)</p> <p><b>sap-id</b> — Specifies that the SAP identifier will be used.</p> <p><b>vlan-ascii-tuple</b> — Specifies that the format will include VLAN-id and dot1p bits in addition to what is included in ascii-tuple already. The format is supported on dot1q-encapsulated ports only. Thus, when the option 82 bits are stripped, dot1p bits will be copied to the Ethernet header of an outgoing packet.</p>

## remote-id

<b>Syntax</b>	<b>remote-id</b> [mac   string <i>string</i> ] <b>no remote-id</b>
<b>Context</b>	config>service>vpls>sap>dhcp>option config>service>ies>interface>dhcp>option config>subscr-mgmt>msap-policy>vpls-only>dhcp>option
<b>Description</b>	<p>This command specifies what information goes into the remote-id sub-option in the DHCP Relay packet.</p> <p>If disabled, the <b>remote-id</b> sub-option of the DHCP packet will be left empty.</p> <p>The <b>no</b> form of this command returns the system to the default.</p>
<b>Default</b>	remote-id
<b>Parameters</b>	<p><b>mac</b> — This keyword specifies the MAC address of the remote end is encoded in the sub-option.</p> <p><b>string</b> <i>string</i> — Specifies the remote-id.</p>

## snoop

<b>Syntax</b>	<b>[no] snoop</b>
<b>Context</b>	config>service>vpls>sap>dhcp - config>service>vpls>spoke-sdp>dhcp config>service>vpls>mesh-sdp>dhcp
<b>Description</b>	<p>This command enables DHCP snooping of DHCP messages on the SAP or SDP. Enabling DHCP snooping on interfaces (SAPs and SDP bindings) is required where DHCP messages important to lease state table population are received, or where Option 82 information is to be inserted. This includes interfaces that are in the path to receive messages from either DHCP servers or from subscribers.</p> <p>Use the <b>no</b> form of the command to disable DHCP snooping on the specified VPLS SAP or SDP binding.</p>
<b>Default</b>	no snoop

## server

<b>Syntax</b>	<b>server server1 [server2...(up to 8 max)]</b>
<b>Context</b>	config>service>ies>if>dhcp config>service>vprn>if>dhcp config>service>ies>subscriber-interface>grp-if>dhcp
<b>Description</b>	<p>This command specifies a list of servers where requests will be forwarded. The list of servers can be entered as either IP addresses or fully qualified domain names. There must be at least one server specified for DHCP relay to work. If there are multiple servers then the request is forwarded to all of the servers in the list.</p> <p>There can be a maximum of 8 DHCP servers configured.</p>
<b>Default</b>	no server
<b>Parameters</b>	<i>server</i> — Specify the DHCP server IP address.

## trusted

<b>Syntax</b>	<b>[no] trusted</b>
<b>Context</b>	config>service>ies>if>dhcp
<b>Description</b>	<p>According to RFC 3046, <i>DHCP Relay Agent Information Option</i>, a DHCP request where the giaddr is 0.0.0.0 and which contains a Option 82 field in the packet, should be discarded, unless it arrives on a "trusted" circuit. If trusted mode is enabled on an IP interface, the Relay Agent (the router) will modify the request's giaddr to be equal to the ingress interface and forward the request.</p> <p>Note that this behavior only applies when the action in the Relay Agent Information Option is "keep". In the case where the option 82 field is being replaced by the Relay Agent (action = "replace"), the</p>

## Service Commands

original Option 82 information is lost anyway, and there is thus no reason for enabling the trusted option.

The **no** form of this command returns the system to the default.

**Default** not enabled

## Egress Multicast Group Commands

### egress-multicast-group

<b>Syntax</b>	<b>egress-multicast-group</b> <i>egress-multicast-group-name</i> <b>no egress-multicast-group</b> <i>group-name</i>
<b>Context</b>	config>service
<b>Description</b>	<p>This command creates an egress multicast group (EMG) context. An EMG is created as an object used to group VPLS SAPs that are allowed to participate in efficient multicast replication (EMR). EMR is a method to increase the performance of egress multipoint forwarding by sacrificing some destination-based features. Eliminating the requirement to perform unique features for each destination allows the egress forwarding plane to chain together multiple destinations into a batch replication process. In order to perform this batch replication function, similar characteristics are required on each SAP within the EMG.</p> <p>Only SAPs defined on Ethernet access ports are allowed into an egress-multicast-group.</p> <p>In order to understand the purpose of an egress-multicast-group, an understanding of the system's use of flooding lists is required. A flooding list is maintained at the egress forwarding plane to define a set of destinations to which a packet must be replicated. Multipoint services make use of flooding lists to enable forwarding a single packet to many destinations. Examples of multipoint services that use flooding lists are VPLS, IGMP snooping and IP multicast routing. Currently, the egress forwarding plane will only use efficient multicast replication for VPLS and IGMP snooping flooding lists.</p> <p>In VPLS services, a unique flooding list is created for each VPLS context. The flooding list is used when a packet has a broadcast, multicast or unknown destination MAC address. From a system perspective, proper VPLS handling requires that a broadcast, multicast or unknown destined packet be sent to all destinations that are in the forwarding state. The ingress forwarding plane ensures the packet gets to all egress forwarding planes that include a destination in the VPLS context. It is the egress forwarding plane's job to replicate the packet to the subset of the destinations that are reached through its interfaces and each of these destinations are included in the VPLS context's flooding list.</p> <p>For IGMP snooping, a unique flooding list is created for each IP multicast (s,g) record. This (s,g) record is associated with an ingress VPLS context and may be associated with VPLS destinations in the source VPLS instance or other VPLS instances (in the case of MVR). Again, the ingress forwarding plane ensures that an ingress IP multicast packet matching the (s,g) record gets to all egress forwarding planes that have a VPLS destination associated with the (s,g) record. The egress forwarding plane uses the flooding list owned by the (s,g) record to replicate the packet to all VPLS destinations in the flooding list. The IGMP Snooping function identifies which VPLS destinations should be associated with the (s,g) record.</p> <p>With normal multicast replication, the egress forwarding plane examines which features are enabled for each destination. This includes ACL filtering, mirroring, encapsulation and queuing. The resources used to perform this per destination multicast processing are very expensive to the egress forwarding plane when high replication bandwidth is required. If destinations with similar egress functions can be grouped together, the egress forwarding plane can process them in a more efficient manner and maximize replication bandwidth.</p> <p>The egress-multicast-group object is designed to allow the identification of SAPs with similar egress characteristics. When a SAP is successfully provisioned into an egress-multicast-group, the system is</p>

ensured that it may be batched together with other SAPs in the same group at the egress forwarding plane for efficient multicast replication. A SAP that does not meet the common requirements is not allowed into the egress-multicast-group.

At the forwarding plane level, a VPLS flooding list is categorized into chainable and non-chainable destinations. Currently, the only chainable destinations are SAPs within an egress-multicast-group. The chainable destinations are further separated by egress-multicast-group association. Chains are then created following the rules below:

- A replication batch chain may only contain SAPs from the same egress-multicast-group
- A replication batch chain length may not exceed the dest-chain-limit of the egress-multicast-group to which the SAPs are members

Further subcategories are created for an IGMP (s,g) flooding list. A Layer 2 (s,g) record is created in a specific VPLS instance (the instance the (s,g) flow ingresses). SAPs within that VPLS context that join the (s,g) record are considered native SAPs within the flooding list. SAPs that join the (s,g) flooding list through the multicast VPLS registration process (MVR) from another VPLS context using the **from-vpls** command are considered alien SAPs. The distinction between native and alien in the list is maintained to allow the forwarding plane to enforce or suspend split-horizon-group (SHG) squelching. When the source of the (s,g) matching packet is in the same SHG as a native SAP, the packet must not be replicated to that SAP. For a SAP in another VPLS context, the source SHG of the packet has no meaning and the forwarding plane must disregard SHG matching between the native source of the packet and the alien destination. Because the SHG squelch decision is done for the whole chain based on the first SAP in the chain, all SAPs in the chain must be all native or all alien SAPs. Chains for IGMP (s,g) flooding lists are created using the following rules:

1. A replication batch chain may only contain SAPs from the same egress-multicast-group.
2. A replication batch chain may only contain all alien or all native SAPs.
3. A replication batch chain length may not exceed the dest-chain-limit of the egress-multicast-group to which the SAPs are members

When a packet associated with a flooding list is received by the egress forwarding plane, it processes the packet by evaluating each destination on the list sequentially in a replication context. If the current entry being processed in the list is a non-chained destination, the forwarding plane processes the packet for that destination and then moves on to process other packets currently in the forwarding plane before returning to process the next destination in the list. If the current entry being processed is a chained destination, the forwarding plane remains in the replication context until it has forwarded to each entry in that chain. Once the replication context finishes with the last entry in the chain, it moves on to process other packets waiting for egress processing before returning to the replication context. Processing continues in this manner until the packet has been forwarded to all destinations in the list.

Batch chain processing of a chain of SAPs improves replication efficiency by bypassing the functions that perform egress mirroring decisions on SAPs within the chain and making a single ACL filtering decision for the whole chain. Each destination in the chain may have a unique egress QoS policy and per destination queuing is still performed for each destination in the chain. Also, while each SAP in the chain must be on access ports with the same encap-type, if the encap-type is dot1q, each SAP may have a unique dot1q tag.

One caveat to each SAP having a unique egress QoS policy in the chain is that only the Dot1P marking decisions for the first SAP in the list is enforced. If the first SAP's QoS policy forwarding class action states that the packet should not be remarked, none of the replicated packets in the chain will have the dot1P bits remarked. If the first SAP's QoS policy forwarding class action states that the packet should be remarked with a specific dot1P value, all the replicated packets for the remaining SAPs in the chain will have the same dot1P marking.

While the system supports 32 egress multicast groups, a single group would usually suffice. An instance where multiple groups would be needed is when all the SAPs requiring efficient multicast replication cannot share the same common requirements. In this case, an egress multicast group would be created for each set of common requirements. An egress multicast group may contain SAPs from many different VPLS instances. It should be understood that an egress multicast group is not equivalent to an egress forwarding plane flooding list. An egress multicast group only identifies which SAPs may participate in efficient multicast replication. As stated above, entries in a flooding list are populated due to VPLS destination creation or IGMP snooping events.

The **no** form of the command removes a specific egress multicast group. Deleting an egress multicast group will only succeed when the group has no SAP members. To remove SAP members, use the **no multicast-group group-name** command under each SAP's egress context.

**Note:** Efficient multicast replication will only be performed on IOMs that support chassis mode b. If an IOM does not support mode b operation, egress-multicast-group membership is ignored on that IOM's egress forwarding planes. The chassis need not be placed into mode b for efficient multicast replication to be performed on the capable IOMs.

<b>Parameters</b>	<i>group-name</i> — Multiple egress multicast groups may be created on the system. Each must have a unique name. The egress-multicast-group-name is an ASCII string up to 16 characters in length and follows all the naming rules as other named policies in the system. The group's name is used throughout the system to uniquely identify the Egress Multicast Group and is used to provision a SAP into the group.
<b>Default</b>	None, each egress multicast group must be explicitly configured.
<b>Values</b>	Up to 32 egress multicast groups may be created on the system.

## description

<b>Syntax</b>	<b>description</b> <i>description-string</i> <b>no description</b>
<b>Context</b>	config>service>egress-multicast-group
<b>Description</b>	This command defines an ASCII string associated with egress-multicast-group-name. The <b>no</b> form of the command removes an existing description string from egress-multicast-group.
<b>Default</b>	none
<b>Parameters</b>	<i>description-string</i> — The description command accepts a description-string parameter. The description-string parameter is an ASCII string of up to 80 characters in length. Only printable 127 bit ASCII characters are allowed. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.
<b>Values</b>	An ASCII string up to 80 characters in length.

dest-chain-limit

<b>Syntax</b>	<b>dest-chain-limit</b> <i>destinations per pass</i> <b>no dest-chain-limit</b>
<b>Context</b>	config>service>egress-multicast-group
<b>Description</b>	<p>This command defines the maximum length of an egress forwarding plane efficient multicast replication chain for an egress-multicast-group. Varying the maximum length of chains created for an egress multicast group has the effect of efficient multicast batched chain replication on other packets flowing through the egress forwarding plane. While replicating for the SAPs within a replication chain, other packets are waiting for the forwarding plane to finish. As the chain length increases, forwarding latency for the other waiting packets may increase. When the chain length decreases, a loss of efficiency in the replication process will be observed.</p> <p>The <b>no</b> form of the command restores the default value of 10 to the dest-chain-limit parameter for the egress-multicast-group.</p>
<b>Default</b>	no dest-chain-limit
<b>Parameters</b>	<p><i>destinations per pass</i> — This parameter must be specified when executing the <b>dest-chain-limit</b> command. When executed, the command will use the number-of-destinations parameter to reorganize all efficient multicast SAP chains that contain members from the egress-multicast-group.</p> <p>The <i>destinations per pass</i> parameter can be modified at any time. Be aware that when changing the maximum chain length, the system will rebuild the chains according to the new limit. When this happens, it is possible that packets will not be replicated to a destination while it is being reorganized in the flooding lists' chains. Only the chains associated with the egress-multicast-group context the command is executed in will be affected by changing the parameter.</p> <p>It is expected that the optimal replication chain length will be between 10 and 16. Since so many variables affect efficient multicast (i.e. ingress packet rate, number of chains, size of replicated packets), only proper testing in the environment that replication will be performed will identify the best dest-chain-limit value for each Egress Multicast Group.</p> <p>Setting the <i>destinations per pass</i> parameter to a value of 0 has the effect of removing from all egress forwarding planes all chains with members from the egress-multicast-group. Replication to each destination SAP from the group is performed using the normal method (non-efficient replication). The value 0 is not considered a normal value for dest-chain-limit and is provided for debugging purposes only. Setting the value to 0 is persistent between reboots of the system.</p> <p>Setting the <i>destinations per pass</i> parameter to a value of 1 has the effect of placing each egress-multicast-group member SAP into a chain with a single SAP. The value 1 is not considered a normal value for the <b>dest-chain-limit</b> and is provided for debugging purposes only. Setting the value to 1 is persistent between reboots of the system.</p> <p><b>Values</b>        0 — 255</p> <p><b>Default</b>       10</p>



## sap-common-requirements

<b>Syntax</b>	<b>sap-common-requirements</b>
<b>Context</b>	config>service>egress-multicast-group
<b>Description</b>	This command configures the common SAP parameter requirements. The SAP common requirements are used to evaluate each SAP for group membership. If a SAP does not meet the specified requirements, the SAP is not allowed into the egress-multicast-group. Once a SAP is a member of the group, attempting to change the parameters on the SAP will fail.

## egress-filter

<b>Syntax</b>	<b>egress-filter</b> [ <b>ip</b> <i>ip-filter-id</i> ] <b>egress-filter</b> [ <b>mac</b> <i>mac-filter-id</i> ] <b>no egress-filter</b> [ <b>ip</b> <i>ip-filter-id</i> ] [ <b>mac</b> <i>mac-filter-id</i> ]
<b>Context</b>	config>service>egress-multicast-group>sap-common-requirements
<b>Description</b>	<p>This command identifies the type of filter and actual filter ID that must be provisioned on the SAP prior to the SAP being made a member of the egress-multicast-group. If the SAP does not have the specified filter applied, the SAP cannot be provisioned into the group. It is important that the egress filter applied to each SAP within the egress-multicast-group be the same since the batch replication process on an efficient multicast replication chain will apply the first SAP's ACL decision to all other SAPs on the chain.</p> <p>Once the SAP is made a member of the egress-multicast-group, the SAP's egress filter cannot be changed on the SAP.</p> <p>Changing the <b>egress-filter</b> parameters within the <b>sap-common-requirements</b> node automatically changes the egress filter applied to each member SAP. If the filter cannot be changed on the SAP due to resource constraints, the modification will fail.</p> <p>The specified egress-filter does not contain an entry that is defined as an egress mirror-source. Once the filter is associated with the egress-multicast-group, attempting to define one of its entries as an egress mirror source will fail.</p> <p>The <b>no</b> form of the command removes the egress-filter from each member SAP. The <b>no egress-filter</b> command specifies that an egress filter (IP or MAC) is not applied to a new member SAP within the egress-multicast-group.</p>
<b>Default</b>	<b>no filter</b> . The egress filter ID must be defined with the associated <b>ip</b> or <b>mac</b> keyword. If an egress-filter is not specified or the no egress-filter command is executed in the sap-common-requirements node, a new member SAP does not have an egress IP or MAC filter defined.
<b>Parameters</b>	<p><b>ip</b> <i>ip-filter-id</i> — Specifies IP filter policy. The filter ID must already exist within the created IP filters.</p> <p><b>Values</b> 1 — 65535</p> <p><b>mac</b> <i>mac-filter-id</i> — Specifies the MAC filter policy. The specified filter ID must already exist within the created MAC filters. The filter policy must already exist within the created MAC filters.</p> <p><b>Values</b> 1 — 65535</p>

## encap-type

<b>Syntax</b>	<b>encap-type {dot1q   null}</b> <b>no encap-type</b>
<b>Context</b>	config>service>egress-multicast-group>sap-common-requirements
<b>Description</b>	<p>This command specifies the encapsulation type that must exist on the SAP's access port to allow the SAP membership within the egress-multicast-group. The <b>config&gt;port&gt;ethernet&gt;access&gt;encap-type</b> command is used to define the encapsulation type for the Ethernet port. The allowed encapsulation type values are dot1q and null. If the SAP does not exist on a port with the specified encap-type, it will not be allowed into the egress-multicast-group.</p> <p>If at least one SAP is currently a member of the efficient-multicast-group, the <b>encap-type</b> cannot be changed within the sap-common-requirements node. If the efficient-multicast-group does not contain any member SAPs, the <b>encap-type</b> may be changed at anytime.</p> <p>There is no interaction between an efficient-multicast-group and the corresponding access ports associated with its members since all SAPs must be deleted from a port before its encap-type can be changed. When the SAPs are deleted from the port, they are also automatically deleted from the efficient-multicast-group.</p> <p>The <b>no</b> form of the command returns the egress-multicast-group required encapsulation type for SAPs to dot1q. If the current encap-type is set to null, the command cannot be executed when SAPs exist within the egress-multicast-group.</p>
<b>Default</b>	<p><b>dot1q</b> — For an egress-multicast-group.</p> <p><b>null</b> — If member SAPs are on a null encapsulated access port.</p>
<b>Parameters</b>	<p><b>null</b> — The <b>null</b> keyword is mutually exclusive with the <b>dot1q</b> keyword. When the encap-type within the sap-common-requirements is specified to be null, the encapsulation type for the access ports associated with all SAPs within the egress-multicast-group must be set to null.</p> <p><b>dot1q</b> — The <b>dot1q</b> keyword is mutually exclusive with the <b>null</b> keyword. When the encap-type within the sap-common-requirements is specified to be dot1q, the encapsulation type for the access ports associated with all SAPs within the egress-multicast-group must be set to dot1q.</p>

## dot1q-etype

<b>Syntax</b>	<b>dot1q-etype [0x0600..0xffff]</b> <b>no dot1q-etype</b>
<b>Context</b>	config>service>egress-multicast-group>sap-common-requirements
<b>Description</b>	<p>This command specifies the dot1q EtherType that must exist on the SAP's access port to allow the SAP membership within the egress-multicast-group. The <b>config&gt;port&gt;ethernet&gt;access&gt;dot1q-etype</b> command is used to define the EtherType used when encapsulating a packet with a dot1q tag on the Ethernet port. Any valid EtherType is allowed on the port.</p> <p>If the current encap-type for the egress-multicast-group is set to null, the dot1q-etype EtherType is ignored when evaluating SAP membership in the group. If the encap-type is set to dot1q (the default), a member SAP's access port must be configured with the same dot1q-etype EtherType as the egress-multicast-group.</p>

If at least one SAP is currently a member of the efficient-multicast-group, the dot1q-etype value cannot be changed within the sap-common-requirements node. If the efficient-multicast-group does not contain any member SAPs, the dot1q-etype value may be changed at anytime.

If an access port currently has SAPs associated with it that are defined within an egress-multicast-group and the port is currently set to encap-type dot1q, the dot1q-etype value defined on the port cannot be changed.

The **no** form of the command returns the egress-multicast-group dot1q EtherType to the default value of 0x8100. If the current encap-type is set to a value other than 0x8100, the command cannot be executed when SAPs exist within the egress-multicast-group.

<b>Default</b>	The default dot1q-etype is 0x8100 for an egress-multicast-group.
<b>Parameters</b>	<i>ethertype</i> — Defines the dot1q EtherType that must be associated with a SAP's access port when the encap-type is set to dot1q. Any valid EtherType may be specified.
<b>Values</b>	[0x0600 — 0xffff]: [1536 — 65535] in decimal or hex
<b>Default</b>	0x8100

---

## Interface Commands

### interface

<b>Syntax</b>	<b>interface</b> <i>ip-int-name</i> <b>no interface</b> <i>ip-int-name</i>
<b>Context</b>	config>service>vprn
<b>Description</b>	<p>This command creates a logical IP routing interface for a Virtual Private Routed Network (VPRN). Once created, attributes like an IP address and service access point (SAP) can be associated with the IP interface.</p> <p>The <b>interface</b> command, under the context of services, is used to create and maintain IP routing interfaces within VPRN service IDs. The <b>interface</b> command can be executed in the context of an VPRN service ID. The IP interface created is associated with the service core network routing instance and default routing table. The typical use for IP interfaces created in this manner is for subscriber internet access.</p> <p>Interface names are case sensitive and must be unique within the group of defined IP interfaces defined for <b>config router interface</b> and <b>config service vprn interface</b>. Interface names must not be in the dotted decimal notation of an IP address. For example, the name “1.1.1.1” is not allowed, but “int-1.1.1.1” is allowed. Show commands for router interfaces use either interface names or the IP addresses. Use unique IP address values and IP address names to maintain clarity. It could be unclear to the user if the same IP address and IP address name values are used. Although not recommended, duplicate interface names can exist in different router instances.</p> <p>The available IP address space for local subnets and routes is controlled with the <b>config router service-prefix</b> command. The <b>service-prefix</b> command administers the allowed subnets that can be defined on service IP interfaces. It also controls the prefixes that may be learned or statically defined with the service IP interface as the egress interface. This allows segmenting the IP address space into <b>config router</b> and <b>config service</b> domains.</p> <p>When a new name is entered, a new logical router interface is created. When an existing interface name is entered, the user enters the router interface context for editing and configuration.</p> <p>By default, there are no default IP interface names defined within the system. All VPRN IP interfaces must be explicitly defined. Interfaces are created in an enabled state.</p> <p>The <b>no</b> form of this command removes IP the interface and all the associated configuration. The interface must be administratively shutdown before issuing the <b>no interface</b> command.</p> <p>For VPRN services, the IP interface must be shutdown before the SAP on that interface may be removed. VPRN services do not have the <b>shutdown</b> command in the SAP CLI context. VPRN service SAPs rely on the interface status to enable and disable them.</p>
<b>Parameters</b>	<i>ip-int-name</i> — Specifies the name of the IP interface. Interface names must be unique within the group of defined IP interfaces for <b>config router interface</b> and <b>config service vprn interface</b> commands. An interface name cannot be in the form of an IP address. Interface names can be from 1 to 32 alphanumeric characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

If *ip-int-name* already exists within the service ID, the context will be changed to maintain that IP interface. If *ip-int-name* already exists within another service ID or is an IP interface defined within the **config router** commands, an error will occur and context will not be changed to that IP interface. If *ip-int-name* does not exist, the interface is created and context is changed to that interface for further command processing.

## interface

<b>Syntax</b>	<b>interface</b> <i>ip-int-name</i> <b>no interface</b> <i>ip-int-name</i>
<b>Context</b>	config>service>ies
<b>Description</b>	<p>This command creates a logical IP routing interface for an Internet Ethernet Service (IES). Once created, attributes like an IP address and service access point (SAP) can be associated with the IP interface.</p> <p>The <b>interface</b> command, under the context of services, is used to create and maintain IP routing interfaces within IES service IDs. The <b>interface</b> command can be executed in the context of an IES, service ID. The IP interface created is associated with the service core network routing instance and default routing table. The typical use for IP interfaces created in this manner is for subscriber internet access.</p> <p>Interface names are case sensitive and must be unique within the group of defined IP interfaces defined for <b>config router interface</b> and <b>config service ies interface</b> (that is, the network core router instance). Interface names must not be in the dotted decimal notation of an IP address. For example, the name “1.1.1.1” is not allowed, but “int-1.1.1.1” is allowed. Show commands for router interfaces use either interface names or the IP addresses. Use unique IP address values and IP address names to maintain clarity. It could be unclear to the user if the same IP address and IP address name values are used. Although not recommended, duplicate interface names can exist in different router instances.</p> <p>The available IP address space for local subnets and routes is controlled with the <b>config router service-prefix</b> command. The <b>service-prefix</b> command administers the allowed subnets that can be defined on IES IP interfaces. It also controls the prefixes that may be learned or statically defined with the IES IP interface as the egress interface. This allows segmenting the IP address space into <b>config router</b> and <b>config service</b> domains.</p> <p>When a new name is entered, a new logical router interface is created. When an existing interface name is entered, the user enters the router interface context for editing and configuration. By default, there are no default IP interface names defined within the system. All IES IP interfaces must be explicitly defined. Interfaces are created in an enabled state.</p> <p>The <b>no</b> form of this command removes the IP interface and all the associated configuration. The interface must be administratively shutdown before issuing the <b>no interface</b> command.</p> <p>For IESServices, the IP interface must be shutdown before the SAP on that interface may be removed. IESServices do not have the <b>shutdown</b> command in the SAP CLI context. IESService SAPs rely on the interface status to enable and disable them.</p>
<b>Parameters</b>	<i>ip-int-name</i> — Specifies the name of the IP interface. Interface names must be unique within the group of defined IP interfaces for <b>config router interface</b> and <b>config service ies interface</b> commands. An interface name cannot be in the form of an IP address. Interface names can be from 1 to 32 alphanumeric characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

If *ip-int-name* already exists within the service ID, the context will be changed to maintain that IP interface. If *ip-int-name* already exists within another service ID or is an IP interface defined within the **config router** commands, an error will occur and context will not be changed to that IP interface. If *ip-int-name* does not exist, the interface is created and context is changed to that interface for further command processing.

address

- Syntax

**address** {*ip-address/mask*|*ip-address netmask*} [**broadcast all-ones**|**host-ones**] [**track-srrp srrp-instance**]  
**no address**
- Context

config>service>ies>interface  
config>service>vpn>interface
- Description

This command assigns an IP address, IP subnet, and broadcast address format to an IES IP router interface. Only one IP address can be associated with an IP interface.

An IP address must be assigned to each IES IP interface. An IP address and a mask are used together to create a local IP prefix. The defined IP prefix must be unique within the context of the routing instance. It cannot overlap with other existing IP prefixes defined as local subnets on other IP interfaces in the same routing context within the router.

The local subnet that the **address** command defines must be part of the services' address space within the routing context using the **config router service-prefix** command. The default is to disallow the complete address space to services. Once a portion of the address space is allocated as a service prefix, that portion can be made unavailable for IP interfaces defined within the **config router interface** CLI context for network core connectivity with the **exclude** option in the **config router service-prefix** command.

The IP address for the interface can be entered in either CIDR (Classless Inter-Domain Routing) or traditional dotted decimal notation. The show commands display CIDR notation and is stored in configuration files.

By default, no IP address or subnet association exists on an IP interface until it is explicitly created.

Use the **no** form of this command to remove the IP address assignment from the IP interface. When the **no address** command is entered, the interface becomes operationally down.

Address	Admin State	Oper State
No address	up	down
No address	down	down
1.1.1.1	up	up
1.1.1.1	down	down

The operational state is a read-only variable and the only controlling variables are the address and admin states. The address and admin states are independent and can be set independently. If an interface is in an administratively up state and an address is assigned, it becomes operationally up and the protocol interfaces and the MPLS LSPs associated with that IP interface will be reinitialized.

*ip-address* — The IP address of the IP interface. The *ip-address* portion of the **address** command specifies the IP host address that will be used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).

/ — The forward slash is a parameter delimiter and separates the *ip-address* portion of the IP address from the mask that defines the scope of the local subnet. No spaces are allowed between the *ip-address*, the “/” and the *mask-length* parameter. If a forward slash is not immediately following the *ip-address*, a dotted decimal mask must follow the prefix.

*mask-length* — The subnet mask length when the IP prefix is specified in CIDR notation. When the IP prefix is specified in CIDR notation, a forward slash (/) separates the *ip-address* from the *mask-length* parameter. The mask length parameter indicates the number of bits used for the network portion of the IP address; the remainder of the IP address is used to determine the host portion of the IP address. Allowed values are integers in the range 0 – 31. Note that a mask length of 32 is reserved for loopback addresses (includes system addresses).

*mask* — The subnet mask in dotted decimal notation. When the IP prefix is not specified in CIDR notation, a space separates the *ip-address* from a traditional dotted decimal mask. The *mask* parameter indicates the complete mask that will be used in a logical ‘AND’ function to derive the local subnet of the IP address. Allowed values are dotted decimal addresses in the range 128.0.0.0 – 255.255.255.252. Note that a mask of 255.255.255.255 is reserved for system IP addresses.

**broadcast** — The optional **broadcast** parameter overrides the default broadcast address used by the IP interface when sourcing IP broadcasts on the IP interface. If no broadcast format is specified for the IP address, the default value is **host-ones** which indicates a subnet broadcast address. Use this parameter to change the broadcast address to **all-ones** or revert back to a broadcast address of **host-ones**.

The broadcast format on an IP interface can be specified when the IP address is assigned or changed.

This parameter does not affect the type of broadcasts that can be received by the IP interface. A host sending either the local broadcast (**all-ones**) or the valid subnet broadcast address (**host-ones**) will be received by the IP interface.

**Default**      host-ones

**all-ones** — The **all-ones** keyword following the **broadcast** parameter specifies the broadcast address used by the IP interface for this IP address will be 255.255.255.255, also known as the local broadcast.

**host-ones** — The **host-ones** keyword following the **broadcast** parameter specifies that the broadcast address used by the IP interface for this IP address will be the subnet broadcast address. This is an IP address that corresponds to the local subnet described by the *ip-address* and the *mask-length* or *mask* with all the host bits set to binary one. This is the default broadcast address used by an IP interface.

The **broadcast** parameter within the **address** command does not have a negate feature, which is usually used to revert a parameter to the default value. To change the **broadcast** type to **host-ones** after being changed to **all-ones**, the **address** command must be executed with the **broadcast** parameter defined.

## allow-directed-broadcast

<b>Syntax</b>	<b>[no] allow-directed-broadcast</b>
<b>Context</b>	config>service>ies>interface
<b>Description</b>	<p>This command enables the forwarding of directed broadcasts out of the IP interface.</p> <p>A directed broadcast is a packet received on a local router interface destined for the subnet broadcast address on another IP interface. The <b>allow-directed-broadcasts</b> command on an IP interface enables or disables the transmission of packets destined to the subnet broadcast address of the egress IP interface.</p> <p>When enabled, a frame destined to the local subnet on this IP interface will be sent as a subnet broadcast out this interface. Care should be exercised when allowing directed broadcasts as it is a well-known mechanism used for denial-of-service attacks.</p> <p>When disabled, directed broadcast packets discarded at this egress IP interface will be counted in the normal discard counters for the egress SAP.</p> <p>By default, directed broadcasts are not allowed and will be discarded at this egress IP interface.</p> <p>The <b>no</b> form of this command disables the forwarding of directed broadcasts out of the IP interface.</p>
<b>Default</b>	<b>no allow-directed-broadcasts</b> — Directed broadcasts are dropped

## loopback

<b>Syntax</b>	<b>[no] loopback</b>
<b>Context</b>	config>service>ies>interface
<b>Description</b>	<p>This command specifies that the associated interface is a loopback interface that has no associated physical interface. As a result, the associated IES interface cannot be bound to a SAP.</p> <p>Note that you can configure an IES interface as a loopback interface by issuing the <b>loopback</b> command instead of the <b>sap sap-id</b> command. The loopback flag cannot be set on an interface where a SAP is already defined and a SAP cannot be defined on a loopback interface.</p>
<b>Default</b>	None

## mac

<b>Syntax</b>	<b>mac ieee-address</b> <b>no mac</b>
<b>Context</b>	config>service>ies>interface
<b>Description</b>	<p>This command assigns a specific MAC address to an IES IP interface.</p> <p>The <b>no</b> form of the command returns the MAC address of the IP interface to the default value.</p>
<b>Default</b>	The physical MAC address associated with the Ethernet interface that the SAP is configured on (the default MAC address assigned to the interface, assigned by the system).



**Parameters** *ieee-address* — Specifies the 48-bit MAC address for the static ARP in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee, and ff are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

## sap

**Syntax** **[no] sap sap-id [create]**

**Context** config>service>ies>interface

**Description** This command creates a Service Access Point (SAP) within a service. A SAP is a combination of port and encapsulation parameters which identifies the service access point on the interface and within the . Each SAP must be unique.

All SAPs must be explicitly created. If no SAPs are created within a service or on an IP interface, a SAP will not exist on that object.

Enter an existing SAP without the **create** keyword to edit SAP parameters. The SAP is owned by the service in which it was created.

A SAP can only be associated with a single service. A SAP can only be defined on a port that has been configured as an access port using the **config interface port-type port-id mode access** command.

If a port is shutdown, all SAPs on that port become operationally down. When a service is shutdown, SAPs for the service are not displayed as operationally down although all traffic traversing the service will be discarded. The operational state of a SAP is relative to the operational state of the port on which the SAP is defined.

Note that you can configure an IES interface as a loopback interface by issuing the **loopback** command instead of the **sap sap-id** command. The loopback flag cannot be set on an interface where a SAP is already defined and a SAP cannot be defined on a loopback interface.

The **no** form of this command deletes the SAP with the specified port. When a SAP is deleted, all configuration parameters for the SAP will also be deleted. For IES, the IP interface must be shutdown before the SAP on that interface may be removed.

**Default** No SAPs are defined.

**Special Cases** **IES** — IES SAP can be defined with Ethernet or SONET/SDH ports. A SAP is defined within the context of an IP routed interface. Each IP interface is limited to a single SAP definition. Attempts to create a second SAP on an IP interface will fail and generate an error; the original SAP will not be affected.

## secondary

<b>Syntax</b>	<b>secondary</b> { <i>ip-address/mask</i>   <i>ip-address netmask</i> } [ <b>broadcast all-ones</b>   <b>host-ones</b> ] [ <b>igp-inhibit</b> ] <b>no secondary</b> <i>ip-address</i>
<b>Context</b>	config>service>ies>interface  This command assigns a secondary IP address/IP subnet/broadcast address format to the interface.
<b>Default</b>	none
<b>Parameters</b>	<p><i>ip-address</i> — The IP address of the IP interface. The <i>ip-address</i> portion of the <b>address</b> command specifies the IP host address that will be used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).</p> <p><i>mask</i> — The subnet mask in dotted decimal notation. When the IP prefix is not specified in CIDR notation, a space separates the <i>ip-address</i> from a traditional dotted decimal mask. The <i>mask</i> parameter indicates the complete mask that will be used in a logical ‘AND’ function to derive the local subnet of the IP address. Allowed values are dotted decimal addresses in the range 128.0.0.0 – 255.255.255.252. Note that a mask of 255.255.255.255 is reserved for system IP addresses.</p> <p><i>netmask</i> — Specifies a string of 0s and 1s that mask or screen out the network part of an IP address so that only the host computer part of the address remains.</p> <p><b>broadcast</b> — The optional <b>broadcast</b> parameter overrides the default broadcast address used by the IP interface when sourcing IP broadcasts on the IP interface. If no broadcast format is specified for the IP address, the default value is <b>host-ones</b> which indicates a subnet broadcast address. Use this parameter to change the broadcast address to <b>all-ones</b> or revert back to a broadcast address of <b>host-ones</b>.</p> <p>The broadcast format on an IP interface can be specified when the IP address is assigned or changed.</p> <p>This parameter does not affect the type of broadcasts that can be received by the IP interface. A host sending either the local broadcast (<b>all-ones</b>) or the valid subnet broadcast address (<b>host-ones</b>) will be received by the IP interface. (Default: <i>host-ones</i>)</p> <p><b>all-ones</b> — The <b>all-ones</b> keyword following the <b>broadcast</b> parameter specifies the broadcast address used by the IP interface for this IP address will be 255.255.255.255, also known as the local broadcast.</p> <p><b>host-ones</b> — The <b>host-ones</b> keyword following the <b>broadcast</b> parameter specifies that the broadcast address used by the IP interface for this IP address will be the subnet broadcast address. This is an IP address that corresponds to the local subnet described by the <i>ip-address</i> and the <i>mask-length</i> or <i>mask</i> with all the host bits set to binary one. This is the default broadcast address used by an IP interface.</p> <p>The <b>broadcast</b> parameter within the <b>address</b> command does not have a negate feature, which is usually used to revert a parameter to the default value. To change the <b>broadcast</b> type to <b>host-ones</b> after being changed to <b>all-ones</b>, the <b>address</b> command must be executed with the <b>broadcast</b> parameter defined.</p> <p><b>igp-inhibit</b> — The optional <b>igp-inhibit</b> parameter signals that the given secondary IP interface should not be recognized as a local interface by the running IGP. For OSPF and IS-IS, this means</p>

that the specified secondary IP interfaces will not be injected and used as passive interfaces and will not be advertised as internal IP interfaces into the IGP's link state database. For RIP, this means that these secondary IP interfaces will not source RIP updates.

## tos-marking-state

<b>Syntax</b>	<b>tos-marking-state {trusted   untrusted}</b> <b>no tos-marking-state</b>
<b>Context</b>	config>service>ies>interface
<b>Description</b>	<p>This command is used to change the default trusted state to a non-trusted state. When unset or reverted to the trusted default, the ToS field will not be remarked by egress network IP interfaces unless the egress network IP interface has the remark-trusted state set, in which case the egress network interface treats all IES and network IP interfaces as untrusted.</p> <p>When the ingress interface is set to untrusted, all egress network IP interfaces will remark IP packets received on the network interface according to the egress marking definitions on each network interface. The egress network remarking rules also apply to the ToS field of IP packets routed using IGP shortcuts (tunneled to a remote next-hop). However, the tunnel QoS markings are always derived from the egress network QoS definitions.</p> <p>Egress marking and remarking is based on the internal forwarding class and profile state of the packet once it reaches the egress interface. The forwarding class is derived from ingress classification functions. The profile of a packet is either derived from ingress classification or ingress policing.</p> <p>The default marking state for network IP interfaces is trusted. This is equivalent to declaring no tos-marking-state on the network IP interface. When undefined or set to tos-marking-state trusted, the trusted state of the interface will not be displayed when using show config or show info unless the detail parameter is given. The <b>save config</b> command will not store the default tos-marking-state trusted state for network IP interfaces unless the detail parameter is also specified.</p> <p>The <b>no</b> tos-marking-state command is used to restore the trusted state to a network IP interface. This is equivalent to executing the tos-marking-state trusted command.</p>
<b>Default</b>	trusted
<b>Parameters</b>	<p><b>trusted</b> — The default prevents the ToS field to not be remarked by egress network IP interfaces unless the egress network IP interface has the remark-trusted state set.</p> <p><b>untrusted</b> — Specifies that all egress network IP interfaces will remark IP packets received on the network interface according to the egress marking definitions on each network interface.</p>

## arp-populate

<b>Syntax</b>	<b>[no] arp-populate</b>
<b>Context</b>	config>service>ies>interface

<b>Description</b>	<p>This command, when enabled, disables dynamic learning of ARP entries. Instead, the ARP table is populated with dynamic entries from the DHCP Lease State Table (enabled with <b>lease-populate</b>), and optionally with static entries entered with the <b>host</b> command.</p> <p>Enabling the <b>arp-populate</b> command will remove any dynamic ARP entries learned on this interface from the ARP cache.</p> <p>The <b>arp-populate</b> command will fail if an existing static ARP entry exists for this interface.</p> <p>The <b>arp-populate</b> command will fail if an existing static subscriber host on the SAP does not have both MAC and IP addresses specified.</p> <p>Once <b>arp-populate</b> is enabled, creating a static subscriber host on the SAP without both an IP address and MAC address will fail.</p> <p>When <b>arp-populate</b> is enabled, the system will not send out ARP requests for hosts that are not in the ARP cache. Only statically configured and DHCP learned hosts are reachable through an IP interface with <b>arp-populate</b> enabled. The <b>arp-populate</b> command can only be enabled on IES interfaces supporting Ethernet encapsulation.</p> <p>Use the <b>no</b> form of the command to disable ARP cache population functions for static and dynamic hosts on the interface. All static and dynamic host information for this interface will be removed from the system's ARP cache.</p>
<b>Default</b>	not enabled

## arp-timeout

<b>Syntax</b>	<b>arp-timeout</b> <i>seconds</i> <b>no arp-timeout</b>
<b>Context</b>	config>service>ies>interface
<b>Description</b>	<p>This command configures the minimum time in seconds an ARP entry learned on the IP interface will be stored in the ARP table. ARP entries are automatically refreshed when an ARP request or gratuitous ARP is seen from an IP host, otherwise, the ARP entry is aged from the ARP table. If <b>arp-timeout</b> is set to a value of zero seconds, ARP aging is disabled.</p> <p>When the <b>arp-populate</b> and <b>lease-populate</b> commands are enabled on an IES interface, the ARP table entries will no longer be dynamically learned, but instead by snooping DHCP ACK message from a DHCP server. In this case the configured <b>arp-timeout</b> value has no effect.</p> <p>The default value for <b>arp-timeout</b> is 14400 seconds (4 hours).</p> <p>The <b>no</b> form of this command restores <b>arp-timeout</b> to the default value.</p>
<b>Default</b>	14400 seconds
<b>Parameters</b>	<p><i>seconds</i> — The minimum number of seconds a learned ARP entry will be stored in the ARP table, expressed as a decimal integer. A value of zero specifies that the timer is inoperative and learned ARP entries will not be aged.</p> <p><b>Values</b>      0 — 65535</p>

## host-connectivity-verify

<b>Syntax</b>	<b>host-connectivity-verify</b> [ <b>source</b> { <b>vrrp</b>   <b>interface</b> }] [ <b>interval</b> <i>interval</i> ] [ <b>action</b> { <b>remove</b>   <b>alarm</b> }]
<b>Context</b>	config>service>ies>if
<b>Description</b>	<p>This command enables subscriber host connectivity verification for all hosts on this interface. This tool will periodically scan all known hosts (from dhcp-state) and perform UC ARP requests. The subscriber host connectivity verification will maintain state (connected vs. not-connected) for all hosts.</p>
<b>Default</b>	no host-connectivity-verify
<b>Parameters</b>	<p><b>source</b> {<b>vrrp</b>   <b>interface</b>} — Specifies the source to be used for generation of subscriber host connectivity verification packets. The <b>vrrp</b> keyword specifies that the VRRP state should be used to select proper IP and MAC (active uses VRID, back-up uses interface addresses). The <b>interface</b> keyword forces the use of the interface mac and ip addresses. Note that there are up to 16 possible subnets on a given interface, therefore the subscriber host connectivity verification tool will use always an address of the subnet to which the given host is pertaining. In case of group-interfaces, one of the parent subscriber-interface subnets (depending on host's address) will be used.</p> <p><b>interval</b> <i>interval</i> — The interval, expressed in minutes, which specifies when all known sources should be verified. The actual rate is then dependent on number of known hosts and interval.</p> <p><b>Values</b>      1 — 6000</p>

Note that a zero value can be used by the SNMP agent to disable host-connectivity-verification.

**action {remove | alarm}** — Defines the action taken on a subscriber host connectivity verification failure for a given host. The **remove** keyword raises an alarm and removes dhcp-state and releases all allocated resources (queues, table entries, etc.). The **alarm** keyword raises an alarm indicating that the host is disconnected.

### icmp

<b>Syntax</b>	<b>icmp</b>
<b>Context</b>	config>service>ies>interface
<b>Description</b>	This command enables the context to configure Internet Control Message Protocol (ICMP) parameters on a service.

### allow-multiple-wan-addresses

<b>Syntax</b>	<b>[no] allow-multiple-wan-addresses</b>
<b>Context</b>	config>service>ies>sub-if>grp-if>ipv6 config>service>vprn>sub-if>grp-if>ipv6
<b>Description</b>	This command enables host to have two WAN addresses, one from DHCP IA_NA and one from SLAAC assignment.
<b>Default</b>	no allow-multiple-wan-addresses

### nd

<b>Syntax</b>	<b>nd</b>
<b>Context</b>	config>service>vprn>sub-if>group-interface>ipv6 config>service>ies>sub-if>group-interface>ipv6
<b>Description</b>	This command enables the context to configure neighbor discovery parameters.

### dad-snooping

<b>Syntax</b>	<b>[no] dad-snooping</b>
<b>Context</b>	config>service>vprn>sub-if>group-interface>ipv6>nd config>service>ies>sub-if>group-interface>ipv6>nd
<b>Description</b>	This command allows the router to populate the neighbor discovery table through snooping subscribers' duplicate address detection messages.

**Default** no dad-snooping

## neighbor-limit

**Syntax** **neighbor-limit** [1..8]  
**no neighbor-limit**

**Context** config>service>vprn>sub-if>group-interface>ipv6>nd  
 config>service>ies>sub-if>group-interface>ipv6>nd

**Description** This command configures the maximum number of neighbors learned.

**Parameters** 1..8 — Specifies the maximum number of neighbors learned.

## router-advertisements

**Syntax** **[no] router-advertisements**

**Context** config>service>ies>sub-if>group-interface>ipv6  
 config>service>vprn>sub-if>ipv6  
 config>service>ies>sub-if>ipv6

**Description** This command configures IPv6 router advertisements for this group-interface.

## current-hop-limit

**Syntax** **[no] current-hop-limit** *limit*

**Context** config>service>ies>sub-if>group-interface>ipv6>rtr-adv  
 config>service>ies>sub-if>grp-if>ipv6>rtr-sol  
 config>service>ies>sub-if>grp-if>ipv6>rtr-sol

**Description** This command configures the hop-limit advertised for this group-interface.

**Default** 64

**Parameters** *limit* — Specifies the default value to be placed in the current hop limit field in router advertisements sent from this interface.

**Values** 0 — 255

## dns-options

**Syntax** **[no] dns-options**

**Context** config>service>ies>sub-if>grp-if>ipv6>rtr-adv  
 config>service>vprn>sub-if>grp-if>ipv6>rtr-adv  
 config>service>ies>sub-if>grp-if>ipv6>rtr-sol

config>service>ies>sub-if>grp-if>ipv6>rtr-sol

**Description** This command enables the context to configure IPv6 DNS options for SLAAC hosts

## include-dns

**Syntax** [no] include-dns

**Context** config>service>ies>sub-if>grp-if>ipv6>rtr-adv>dns-opt  
config>service>vprn>sub-if>grp-if>ipv6>rtr-adv>dns-opt  
config>service>vprn>sub-if>ipv6>rtr-adv>dns-opt  
config>service>ies>sub-if>ipv6>rtr-adv>dns-opt

**Description** This command specifies to include the Recursive DNS Server (RDNSS) Option as defined in RFC 6106 in IPv6 Router Advertisements for DNS name resolution of IPv6 SLAAC hosts

**Default** no include-dns

## rdnss-lifetime

**Syntax** rdnss-lifetime seconds  
rdnss-lifetime infinite  
no rdnss-lifetime

**Context** config>service>ies>sub-if>grp-if>ipv6>rtr-adv>dns-opt  
config>service>vprn>sub-if>grp-if>ipv6>rtr-adv>dns-opt  
config>service>vprn>sub-if>ipv6>rtr-adv>dns-opt  
config>service>ies>sub-if>ipv6>rtr-adv>dns-opt

**Description** Specify the maximum time in seconds that the RDNSS address may be used for name resolution.

**Default** rdnss-lifetime 3600

**Parameters** seconds — Specifies the time in seconds.

**Values** 900 — 3600

**infinite** — The RDNSS address can be used permanently.

## force-mcast

**Syntax** force-mcast [ip] [mac]  
no force-mcast

**Context** config>service>ies>sub-if>grp-if>ipv6>rtr-adv  
config>service>vprn>sub-if>grp-if>ipv6>rtr-adv  
config>service>ies>sub-if>grp-if>ipv6>rtr-sol  
config>service>ies>sub-if>grp-if>ipv6>rtr-sol

**Description** This command configures the multicast router advertisements on this interface, either IP or MAC.



## managed-configuration

<b>Syntax</b>	<b>[no] managed-configuration</b>
<b>Context</b>	config>service>ies>sub-if>grp-if>ipv6>rtr-adv config>service>vpn>sub-if>grp-if>ipv6>rtr-adv config>service>ies>sub-if>grp-if>ipv6>rtr-sol config>service>ies>sub-if>grp-if>ipv6>rtr-sol
<b>Description</b>	This command sets or resets managed address configuration flag for this group-interface.

## max-advertisement

<b>Syntax</b>	<b>max-advertisement <i>seconds</i></b> <b>no max-advertisement</b>
<b>Context</b>	config>service>ies>sub-if>grp-if>ipv6>rtr-adv config>service>vpn>sub-if>grp-if>ipv6>rtr-adv config>service>ies>sub-if>grp-if>ipv6>rtr-sol config>service>ies>sub-if>grp-if>ipv6>rtr-sol
<b>Description</b>	This command specifies the maximum time allowed between sending unsolicited router advertisements from this interface.
<b>Parameters</b>	<i>seconds</i> — Specifies the maximum advertisement interval in seconds for this group-interface. <b>Values</b> 900 — 11800

## min-advertisement

<b>Syntax</b>	<b>min-advertisement <i>seconds</i></b> <b>no min-advertisement</b>
<b>Context</b>	config>service>ies>sub-if>grp-if>ipv6>rtr-adv config>service>vpn>sub-if>grp-if>ipv6>rtr-adv config>service>ies>sub-if>grp-if>ipv6>rtr-sol config>service>ies>sub-if>grp-if>ipv6>rtr-sol
<b>Description</b>	This command specifies the minimum time allowed between sending unsolicited router advertisements from this interface.
<b>Parameters</b>	<i>seconds</i> — Specifies the minimum advertisement interval in seconds for this group-interface. <b>Values</b> 900 — 1350

## mtu

<b>Syntax</b>	<b>mtu</b> <i>bytes</i> <b>no mtu</b>
<b>Context</b>	config>service>ies>sub-if>grp-if>ipv6>rtr-adv config>service>vpn>sub-if>grp-if>ipv6>rtr-adv config>service>ies>sub-if>grp-if>ipv6>rtr-sol config>service>ies>sub-if>grp-if>ipv6>rtr-sol
<b>Description</b>	This command specifies the value to be placed in link MTU options sent by the router on this interface.
<b>Parameters</b>	<i>bytes</i> — Sets the advertised MTU value in bytes for this group-interface. <b>Values</b> 1280 — 9212

## other-stateful-configuration

<b>Syntax</b>	<b>[no] other-stateful-configuration</b>
<b>Context</b>	config>service>ies>sub-if>grp-if>ipv6>rtr-adv config>service>vpn>sub-if>grp-if>ipv6>rtr-adv config>service>ies>sub-if>grp-if>ipv6>rtr-sol config>service>ies>sub-if>grp-if>ipv6>rtr-sol
<b>Description</b>	This command sets and resets the other-stateful-configuration flag for this group-interface.

## prefix-options

<b>Syntax</b>	<b>[no] prefix-options</b>
<b>Context</b>	config>service>ies>sub-if>grp-if>ipv6>rtr-adv config>service>vpn>sub-if>grp-if>ipv6>rtr-adv config>service>ies>sub-if>grp-if>ipv6>rtr-sol config>service>ies>sub-if>grp-if>ipv6>rtr-sol
<b>Description</b>	This command enables the context to configure prefix options for this group-interface.

## autonomous

<b>Syntax</b>	<b>[no] autonomous</b>
<b>Context</b>	config>service>ies>sub-if>grp-if>ipv6>rtr-adv config>service>vpn>sub-if>grp-if>ipv6>rtr-adv
<b>Description</b>	This command enables or disables the option that determines whether or not the prefix can be used for stateless address autoconfiguration.

## on-link

<b>Syntax</b>	<b>[no] on-link</b>
<b>Context</b>	config>service>ies>sub-if>grp-if>ipv6>rtr-adv config>service>vpn>sub-if>grp-if>ipv6>rtr-adv
<b>Description</b>	This command specifies whether the prefix will be assigned to an interface on the specified link.

## preferred-lifetime

<b>Syntax</b>	<b>preferred-lifetime</b> <i>seconds</i> <b>preferred-lifetime</b> <i>infinite</i> <b>no preferred-lifetime</b>
---------------	---

## reachable-time

<b>Syntax</b>	<b>reachable-time</b> <i>milli-seconds</i> <b>no reachable-time</b>
<b>Context</b>	config>service>ies>sub-if>grp-if>ipv6>rtr-adv config>service>vpn>sub-if>grp-if>ipv6>rtr-adv config>service>ies>sub-if>grp-if>ipv6>rtr-sol config>service>ies>sub-if>grp-if>ipv6>rtr-sol
<b>Description</b>	This command configures the value to be placed in the reachable time field in router advertisement messages sent from this interface.
<b>Default</b>	0
<b>Parameters</b>	<i>milli-seconds</i> — Specifies the reachable time in milli-seconds for advertisements from this group-interface.
<b>Values</b>	0 — 3600000

## retransmit-time

<b>Syntax</b>	<b>retransmit-time</b> <i>milli-seconds</i> <b>no retransmit-time</b>
<b>Context</b>	config>service>ies>sub-if>grp-if>ipv6>rtr-adv config>service>vpn>sub-if>grp-if>ipv6>rtr-adv config>service>ies>sub-if>grp-if>ipv6>rtr-sol config>service>ies>sub-if>grp-if>ipv6>rtr-sol
<b>Description</b>	This command configures the value to be placed in the retransmit timer field in router advertisements sent from this interface.
<b>Default</b>	0

<b>Parameters</b>	<i>milli-seconds</i> — Specifies the retransmit time in milli-seconds for advertisement from this group-interface.
<b>Values</b>	0 — 1800000

### router-lifetime

<b>Syntax</b>	<b>router-lifetime</b> <i>seconds</i> <b>router-lifetime no-default-router</b> <b>no router-lifetime</b>
<b>Context</b>	config>service>ies>sub-if>grp-if>ipv6>rtr-adv config>service>vprn>sub-if>grp-if>ipv6>rtr-adv config>service>ies>sub-if>grp-if>ipv6>rtr-sol config>service>ies>sub-if>grp-if>ipv6>rtr-sol
<b>Description</b>	This command configures the value to be placed in the router lifetime field of router advertisements sent from this interface.
<b>Default</b>	4500
<b>Parameters</b>	<i>seconds</i> — Specifies the router lifetime in seconds for this group-interface. <b>Values</b> 2700 — 9000 <b>no-default-router</b> — Specifies that the router is not to be used as a default router.

### router-solicit

<b>Syntax</b>	<b>router-solicit</b>
<b>Context</b>	config>service>ies>sub-if>grp-if>ipv6 config>service>vprn>sub-if>grp-if>ipv6 config>service>vprn>sub-if>ipv6 config>service>ies>sub-if>ipv6
<b>Description</b>	This command enables the context to configure parameters used for router-solicit based authentication.

### inactivity-timer

<b>Syntax</b>	<b>inactivity-timer</b> [ <i>days days</i> ] [ <i>hrs hours</i> ] [ <i>min minutes</i> ] [ <b>sec</b> <i>seconds</i> ] <b>no inactivity-timer</b>
<b>Context</b>	config>service>ies>sub-if>grp-if>ipv6>rtr-sol config>service>ies>sub-if>grp-if>ipv6>rtr-sol
<b>Description</b>	This command specifies the time before an inactive host is removed.
<b>Default</b>	no interactive-timer

<b>Parameters</b>	<b>infinite</b> — An idle host is never removed. <b>days <i>days</i></b> — An idle host is removed if idle within the number of specified days. <b>hrs <i>hours</i></b> — An idle host is removed if idle within the number of specified hours. <b>min <i>minutes</i></b> — An idle host is removed if idle within the number of specified minutes. <b>sec <i>seconds</i></b> — An idle host is removed if idle within the number of specified seconds.
-------------------	---

## min-auth-interval

<b>Syntax</b>	<b>min-auth-interval [days <i>days</i>] [hrs <i>hours</i>] [min <i>minutes</i>] [sec <i>seconds</i>] no min-auth-interval</b>
<b>Context</b>	config>service>ies>sub-if>grp-if>ipv6>rtr-sol config>service>ies>sub-if>grp-if>ipv6>rtr-sol
<b>Description</b>	This command specify the minimum interval between two consecutive authentication attempts from the same host.
<b>Default</b>	no min-auth-interval
<b>Parameters</b>	<b>days <i>days</i></b> — The number of days that a user must wait for the next authentication attempt. <b>hrs <i>hours</i></b> — The number of hours that a user must wait for the next authentication attempt. <b>min <i>minutes</i></b> — The number of minutes that a user must wait for the next authentication attempt. <b>sec <i>seconds</i></b> — The number of seconds that a user must wait for the next authentication attempt.

## user-db

<b>Syntax</b>	<b>[no] user-db</b>
<b>Context</b>	config>service>ies>sub-if>grp-if>ipv6>rtr-sol config>service>ies>sub-if>grp-if>ipv6>rtr-sol
<b>Description</b>	This command enables the use of the local-user-database for authentication.
<b>Default</b>	no user-db
<b>Parameters</b>	<i>local-user-db-name</i> — Specifies the name of the local-user-database to authenticate the router-solicit. The local-user-database can also return a static prefix or a pool name for address assignment.

## shutdown

<b>Syntax</b>	<b>[no] shutdown</b>
<b>Context</b>	config>service>ies>sub-if>grp-if>ipv6>rtr-sol config>service>ies>sub-if>grp-if>ipv6>rtr-sol
<b>Description</b>	This command enables or disables SLAAC triggered host creation.

**Default** no shutdown

### local-proxy-arp

**Syntax** [no] local-proxy-arp

**Context** config>service>ies>interface

**Description** This command enables local proxy ARP. When local proxy ARP is enabled on an IP interface, the system responds to all ARP requests for IP addresses belonging to the subnet with its own MAC address, and thus will become the forwarding point for all traffic between hosts in that subnet.

When local-proxy-arp is enabled, ICMP redirects on the ports associated with the service are automatically blocked.

**Default** no local-proxy-arp

### mask-reply

**Syntax** [no] mask-reply

**Context** config>service>ies>if>icmp

**Description** This command enables responses to Internet Control Message Protocol (ICMP) mask requests on the router interface.

If a local node sends an ICMP mask request to the router interface, the **mask-reply** command configures the router interface to reply to the request.

By default, the router instance will reply to mask requests.

The **no** form of this command disables replies to ICMP mask requests on the router interface.

**Default** mask-reply

### proxy-arp-policy

**Syntax** [no] proxy-arp-policy *policy-name* [*policy-name*...(up to 5 max)]

**Context** config>service>ies>interface

**Description** This command configures a proxy ARP policy for the interface.

The **no** form of this command disables the proxy ARP capability.

**Default** no proxy-arp-policy

**Parameters** *policy-name* — The export route policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters excluding double quotes. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

The specified name(s) must already be defined.

## redirects

<b>Syntax</b>	<b>redirects</b> [ <i>number seconds</i> ] <b>no redirects</b>
<b>Context</b>	config>service>ies>if>icmp
<b>Description</b>	<p>This command configures the rate for Internet Control Message Protocol (ICMP) redirect messages issued on the router interface.</p> <p>When routes are not optimal on this router and another router on the same subnetwork has a better route, the router can issue an ICMP redirect to alert the sending node that a better route is available.</p> <p>The <b>redirects</b> command enables the generation of ICMP redirects on the router interface. The rate at which ICMP redirects is issued can be controlled with the optional <i>number</i> and <i>seconds</i> parameters by indicating the maximum number of redirect messages that can be issued on the interface for a given time interval.</p> <p>By default, generation of ICMP redirect messages is enabled at a maximum rate of 100 per 10 second time interval. (<i>Default: redirects 100 10</i>)</p> <p>The <b>no</b> form of this command disables the generation of icmp redirects on the router interface.</p>
<b>Default</b>	<b>redirects 100 10</b> — Maximum of 100 redirect messages in 10 seconds.
<b>Parameters</b>	<p><i>number</i> — The maximum number of ICMP redirect messages to send. This parameter must be specified with the <i>second</i> parameter.</p> <p><b>Values</b> 10 — 1000</p> <p><i>seconds</i> — The time frame in seconds used to limit the <i>number</i> of ICMP redirect messages that can be issued.</p> <p><b>Values</b> 1 — 60</p>

## remote-proxy-arp

<b>Syntax</b>	<b>[no] remote-proxy-arp</b>
<b>Context</b>	config>service>ies>interface
<b>Description</b>	<p>This command enables remote proxy ARP on the interface.</p> <p>Remote proxy ARP is similar to proxy ARP. It allows the router to answer an ARP request on an interface for a subnet that is not provisioned on that interface. This allows the router to forward to the other subnet on behalf of the requester. To distinguish remote proxy ARP from local proxy ARP, local proxy ARP performs a similar function but only when the requested IP is on the receiving interface.</p>
<b>Default</b>	no remote-proxy-arp

## static-arp

<b>Syntax</b>	<b>static-arp</b> <i>ip-address ieee-mac-address</i>
---------------	--

**no static-arp** *ip-address ieee-mac-address*

**Context** config>service>ies>interface

**Description** This command configures a static address resolution protocol (ARP) entry associating a subscriber IP address with a MAC address for the core router instance. This static ARP will appear in the core routing ARP table. A static ARP can only be configured if it exists on the network attached to the IP interface.

If an entry for a particular IP address already exists and a new MAC address is configured for the IP address, the existing MAC address will be replaced with the new MAC address.

The **no** form of this command removes a static ARP entry.

**Default** none

**Parameters** *ip-address* — Specifies the IP address for the static ARP in IP address dotted decimal notation.  
*ieee-mac-address* — Specifies the 48-bit MAC address for the static ARP in the form *aa:bb:cc:dd:ee:ff* or *aa-bb-cc-dd-ee-ff* where *aa*, *bb*, *cc*, *dd*, *ee* and *ff* are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

## ttl-expired

**Syntax** **ttl-expired** *number seconds*  
**no ttl-expired**

**Context** config>service>ies>if>icmp

**Description** This command configures the rate Internet Control Message Protocol (ICMP) TTL expired messages are issued by the IP interface.  
 By default, generation of ICMP TTL expired messages is enabled at a maximum rate of 100 per 10 second time interval.

The **no** form of this command disables the limiting the rate of TTL expired messages on the router interface.

**Default** ttl-expired 100 10

**Parameters** *number* — The maximum number of ICMP TTL expired messages to send, expressed as a decimal integer. This parameter must be specified with the *seconds* parameter.

**Values** 10 — 1000

*seconds* — The time frame in seconds used to limit the *number* of ICMP TTL expired messages that can be issued, expressed as a decimal integer.

**Values** 1 — 60

## unreachables

**Syntax** **unreachables** [*number seconds*]  
**no unreachables**



<b>Context</b>	config>service>ies>if>icmp
<b>Description</b>	<p>This command configures the rate for ICMP host and network destination unreachable messages issued on the router interface.</p> <p>The <b>unreachables</b> command enables the generation of ICMP destination unreachables on the router interface. The rate at which ICMP unreachables is issued can be controlled with the optional <i>number</i> and <i>seconds</i> parameters by indicating the maximum number of destination unreachable messages which can be issued on the interface for a given time interval.</p> <p>By default, generation of ICMP destination unreachables messages is enabled at a maximum rate of 100 per 10 second time interval.</p> <p>The <b>no</b> form of this command disables the generation of icmp destination unreachables on the router interface.</p>
<b>Default</b>	<b>unreachables 100 10</b>
<b>Parameters</b>	<p><i>number</i> — The maximum number of ICMP unreachable messages to send. This parameter must be specified with the <i>seconds</i> parameter.</p> <p><b>Values</b>      10 — 1000</p> <p><i>seconds</i> — The time frame in seconds used to limit the <i>number</i> of ICMP unreachable messages that can be issued.</p> <p><b>Values</b>      1 — 60</p>

---

## Interface IPv6 Commands

### ipv6

<b>Syntax</b>	<b>[no] ipv6</b>
<b>Context</b>	config>service>ies>interface config>service>ies>sub-if>grp-if
<b>Description</b>	This command enables the context to configure IPv6 for an interface.

### address

<b>Syntax</b>	<b>address</b> <i>ipv6-address/prefix-length</i> [ <b>eui-64</b> ] <b>no address</b> <i>ipv6-address/prefix-length</i>		
<b>Context</b>	config>service>ies>if>ipv6		
<b>Description</b>	This command assigns an IPv6 address to the IES interface.		
<b>Parameters</b>	<i>ipv6-address/prefix-length</i> — Specify the IPv6 address on the interface.		
<b>Values</b>	ipv6-address	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	
	prefix-length	1 — 128	
<b>eui-64</b> — When the <b>eui-64</b> keyword is specified, a complete IPv6 address from the supplied prefix and 64-bit interface identifier is formed. The 64-bit interface identifier is derived from MAC address on Ethernet interfaces. For interfaces without a MAC address, for example ATM interfaces, the Base MAC address of the chassis is used.			

### default-dns

<b>Syntax</b>	<b>default-dns</b> <i>ipv6-address</i> [ <b>secondary</b> <i>ipv6-address</i> ] <b>no default-dns</b>		
<b>Context</b>	config>service>ies>sub-if>ipv6 config>service>vprn>sub-if>ipv6		
<b>Description</b>	Configure last resort IPv6 DNS addresses that can be used for name resolution by IPEv6 hosts (IA_NA, IA_PD and SLAAC) and PPPoEv6 hosts (IA_NA, IA_PD and SLAAC)		
<b>Default</b>	no default-dns		
<b>Parameters</b>	<i>ipv6-address</i> — s - IPv6 address of the primary DNS server <b>secondary</b> <i>ipv6-address</i> — IPv6 address of the secondary DNS server (optional)		

## allow-unmatching-prefixes

<b>Syntax</b>	<b>[no] allow-unmatching-prefixes</b>
<b>Context</b>	configure>service>vpn>sub-if>ipv6 configure>service>ies>sub-if>ipv6
<b>Description</b>	<p>This command allows address assignments for IPoEv6 and PPPoEv6 hosts in cases where the subscriber host assigned IPv6 address or prefix falls outside of the subscriber-prefix range explicitly configured for the subscriber-interface ( <b>configure&gt;service&gt;vpn/ies&gt;sub-if&gt;ipv6&gt;</b> or the subscriber-prefix is not configured at all.</p> <p>SLAAC hosts will be installed in the FIB as /64 entries, the length of the installed DHCP-PD prefix will be dictated by the prefix-length and the DHCP-NA host will be installed as /128 entries.</p> <p>IPv4 subscriber hosts are unaffected by this command.</p>
<b>Default</b>	no allow-unmatching-prefixes

## delegated-prefix-length

<b>Syntax</b>	<b>delegated-prefix-length <i>bits</i></b> <b>delegated-prefix-length <i>variable</i></b> <b>no delegated-prefix-length</b>
<b>Context</b>	configure>service>vpn>sub-if>ipv6 configure>service>ies>sub-if>ipv6
<b>Description</b>	<p>This command configures the subscriber-interface level setting for delegated prefix length. The delegated prefix length for a subscriber- interface can be either set to a fixed value that is explicitly configured under the subscriber-interface CLI hierarchy or a variable value that can be obtained from various sources. This command can be changed only when no IPv6 prefixes are configured under the subscriber-interface.</p>
<b>Default</b>	no delegated-prefix-length This means that the delegated prefix length is 64.
<b>Parameters</b>	<p><i>bits</i> — The delegated prefix length in bits. This value will be applicable to the entire subscriber-interface. In case that the delegated prefix length is also supplied via other means (LUDB, RADIUS or DHCP Server), such supplied value must match the value configured under the subscriber-interface. Otherwise the prefix instantiation in 7x50 will fail.</p> <p><b>Values</b>      48 — 64</p> <p><b>variable</b> — The delegated prefix value can be of any length between 48..64. The value itself can vary between the prefixes and it will be provided at the time of prefix instantiation. The order of priority for the source of the delegated prefix length is:</p> <ul style="list-style-type: none"> <li>• LUDB</li> <li>• RADIUS</li> <li>• DHCPv6 server</li> </ul>

## subscriber-prefixes

<b>Syntax</b>	<b>[no] subscriber-prefixes</b>
<b>Context</b>	config>services>ies>sub-if>ipv6 config>services>vpn>sub-if>ipv6
<b>Description</b>	This command specifies aggregate off-link subscriber prefixes associated with this subscriber interface. Individual prefixes are specified under the prefix context list aggregate routes in which the next-hop is indirect via the subscriber interface.

## prefix

<b>Syntax</b>	<b>prefix</b> <i>ipv6-address/prefix-length</i> [ <b>pd</b> ] [ <b>wan-host</b> ] <b>no prefix</b> <i>ipv6-address/prefix-length</i>
<b>Context</b>	config>services>ies>sub-if>ipv6>sub-prefixes config>services>vpn>sub-if>ipv6>sub-prefixes
<b>Description</b>	This command allows a list of prefixes(using the prefix command multiple times) to be routed to hosts associated with this subscriber interface. Each prefix will be represented in the associated FIB with a reference to the subscriber interface. Prefixes are defined as being for prefix delegation (pd) or use on a WAN interface or host (wan-host).
<b>Parameters</b>	<i>ipv6-address</i> — Specifies the 128-bit IPv6 address. <div style="margin-left: 40px;"><b>Values</b> 128-bit hexadecimal IPv6 address in compressed form.</div> <i>prefix-length</i> — Specifies the length of any associated aggregate prefix. <div style="margin-left: 40px;"><b>Values</b> 32-63</div> <b>pd</b> — Specifies that this aggregate is used by IPv6 ESM hosts for DHCPv6 prefix-delegation. <b>wan-host</b> — Specifies that this aggregate is used by IPv6 ESM hosts for local addressing or by a routing gateway's WAN interface.

## dhcp6-relay

<b>Syntax</b>	<b>[no] dhcp6-relay</b>
<b>Context</b>	config>service>ies>if>ipv6
<b>Description</b>	This command enables the context to configure DHCPv6 relay parameters for the IES interface. The <b>no</b> form of the command disables DHCP6 relay.

## lease-populate

<b>Syntax</b>	<b>lease-populate</b> [ <i>nbr-of-leases</i> ] <b>lease-populate</b> [ <i>nbr-of-leases</i> ] <b>route-populate</b> [pd] na [ta] <b>lease-populate</b> [ <i>nbr-of-leases</i> ] <b>route-populate</b> pd [na] [ta] [exclude] <b>lease-populate</b> [ <i>nbr-of-leases</i> ] <b>route-populate</b> [pd] [na] ta <b>no lease-populate</b>
<b>Context</b>	config>service>ies>if>ipv6>dhcp6-relay config>service>vprn>if>ipv6>dhcp6-relay
<b>Description</b>	<p>This command specifies the maximum number of DHCPv6 lease states allocated by the DHCPv6 relay function, allowed on this interface.</p> <p>Optionally, by specifying “route-populate” parameter, system could:</p> <ul style="list-style-type: none"> <li>• Create routes based on the IA_PD/IA_NA/IA_TA prefix option in relay-reply message.</li> <li>• Create black hole routes based on OPTION_PD_EXCLUDE in IA_PD in relay-reply message.</li> </ul> <p>These routes could be redistributed into IGP/BGP by using route-policy, following protocol types that could be used in “from protocol”:</p> <ul style="list-style-type: none"> <li>• dhcpv6-pd</li> <li>• dhcpv6-na</li> <li>• dhcpv6-ta</li> <li>• dhcpv6-pd-excl</li> </ul> <p>The <b>no</b> form of the command disables dynamic host lease state management.</p>
<b>Default</b>	no lease-populate
<b>Parameters</b>	<p><i>nbr-of-leases</i> — Defines the number lease state table entries allowed for this interface. If this parameter is omitted, only a single entry is allowed. Once the maximum number of entries has been reached, subsequent lease state entries are not allowed and subsequent DHCP6 ACK messages are discarded.</p> <p><b>Values</b>      1 — 8000</p> <p><b>route-populate</b> — Create routes based on options in relay-reply messages.</p> <p><b>Values</b>      pd/na/ta — Create route based on specified option.  exclude — Create blackhole route based on OPTION_PD_EXCLUDE</p>

## neighbor-resolution

<b>Syntax</b>	<b>[no] neighbor-resolution</b>
<b>Context</b>	config>service>ies>if>ipv6>dhcp6-relay
<b>Description</b>	This command enables neighbor resolution with DHCPv6 relay. The <b>no</b> form of the command disables neighbor resolution.

## remote-id

<b>Syntax</b>	<b>[no] remote-id</b>
<b>Context</b>	config>service>ies>if>ipv6>dhcp6>option
<b>Description</b>	This command enables the sending of remote ID option in the DHCPv6 relay packet. The client DHCP Unique Identifier (DUID) is used as the remote ID.

## option

<b>Syntax</b>	<b>[no] option</b>
<b>Context</b>	config>service>ies>if>ipv6>dhcp6-relay
<b>Description</b>	This command enables the context to configure DHCPv6 relay information options. The <b>no</b> form of the command disables DHCPv6 relay information options.

## interface-id

<b>Syntax</b>	<b>interface-id</b> <b>interface-id ascii-tuple</b> <b>interface-id ifindex</b> <b>interface-id sap-id</b> <b>no interface-id</b>
<b>Context</b>	config>service>ies>if>ipv6>dhcp6>option
<b>Description</b>	This command enables the sending of interface ID options in the DHCPv6 relay packet. The <b>no</b> form of the command disables the sending of interface ID options in the DHCPv6 relay packet
<b>Parameters</b>	<b>ascii-tuple</b> — Specifies that the ASCII-encoded concatenated tuple will be used which consists of the access-node-identifier, service-id, and interface-name, separated by “ ”. <b>ifindex</b> — Specifies that the interface index will be used. (The If Index of a router interface can be displayed using the command show>router>interface>detail) <b>sap-id</b> — Specifies that the SAP identifier will be used.

The **no** form of the command disables the sending of remote ID option in the DHCPv6 relay packet.

## icmp6

<b>Syntax</b>	<b>icmp6</b>
<b>Context</b>	config>service>ies>if>ipv6
<b>Description</b>	This command configures ICMPv6 parameters for the interface.

## packet-too-big

<b>Syntax</b>	<b>packet-too-big</b> [ <i>number seconds</i> ] <b>no packet-too-big</b>								
<b>Context</b>	config>service>ies>if>ipv6>icmp6								
<b>Description</b>	This command specifies whether “packet-too-big” ICMP messages should be sent. When enabled, ICMPv6 “packet-too-big” messages are generated by this interface.  The <b>no</b> form of the command disables the sending of ICMPv6 “packet-too-big” messages.								
<b>Default</b>	100 10								
<b>Parameters</b>	<i>number</i> — Specifies the number of “packet-too-big” ICMP messages to send in the time frame specified by the <i>seconds</i> parameter.  <table> <tr> <td><b>Values</b></td><td>10 — 1000</td></tr> <tr> <td><b>Default</b></td><td>100</td></tr> </table> <i>seconds</i> — Specifies the time frame in seconds that is used to limit the number of “packet-too-big” ICMP messages issued.  <table> <tr> <td><b>Values</b></td><td>1 — 60</td></tr> <tr> <td><b>Default</b></td><td>10</td></tr> </table>	<b>Values</b>	10 — 1000	<b>Default</b>	100	<b>Values</b>	1 — 60	<b>Default</b>	10
<b>Values</b>	10 — 1000								
<b>Default</b>	100								
<b>Values</b>	1 — 60								
<b>Default</b>	10								

## param-problem

<b>Syntax</b>	<b>param-problem</b> [ <i>number seconds</i> ] <b>no packet-too-big</b>
<b>Context</b>	config>service>ies>if>ipv6>icmp6
<b>Description</b>	This command specifies whether “parameter-problem” ICMP messages should be sent. When enabled, “parameter-problem” ICMP messages are generated by this interface.  The <b>no</b> form of the command disables the sending of “parameter-problem” ICMP messages.
<b>Default</b>	100 10

*number* — Specifies the number of “parameter-problem” ICMP messages to send in the time frame specified by the *seconds* parameter.

**Values** 10 — 1000

**Default** 100

*seconds* — Specifies the time frame in seconds that is used to limit the number of “parameter-problem” ICMP messages issued.

**Values** 1 — 60

**Default** 10

## redirects

**Syntax** **redirects** [*number seconds*]  
**no redirects**

**Context** config>service>ies>if>ipv6>icmp6

**Description** This command configures ICMPv6 redirect messages. When enabled, ICMPv6 redirects are generated when routes are not optimal on this router and another router on the same subnetwork has a better route in order to alert that node that a better route is available.

When disabled, ICMPv6 redirects are not generated.

**Default** 100 10

*number* — Specifies the number of version 6 redirects are to be issued in the time frame specified by the *seconds* parameter.

**Values** 10 — 1000

**Default** 100

*seconds* — Specifies the time frame in seconds that is used to limit the number of version 6 redirects issued.

**Values** 1 — 60

**Default** 10

## time-exceeded

**Syntax** **time-exceeded** [*number seconds*]  
**no time-exceeded**

**Context** config>service>ies>if>ipv6>icmp6

**Description** This command specifies whether “time-exceeded” ICMP messages should be sent. When enabled, ICMPv6 “time-exceeded” messages are generated by this interface.

When disabled, ICMPv6 “time-exceeded” messages are not sent.

**Default** 100 10



*number* — Specifies the number of “time-exceeded” ICMP messages are to be issued in the time frame specified by the *seconds* parameter.

**Values** 10 — 1000

**Default** 100

*seconds* — Specifies the time frame in seconds that is used to limit the number of “time-exceeded” ICMP message to be issued.

**Values** 1 — 60

**Default** 10

## unreachables

**Syntax** **unreachables** [*number seconds*]  
**no unreachables**

**Context** config>service>ies>if>ipv6>icmp6

**Description** This command specifies that ICMPv6 host and network unreachable messages are generated by this interface.

When disabled, ICMPv6 host and network unreachable messages are not sent.

**Default** 100 10

*number* — Specifies the number of destination unreachable ICMPv6 messages are issued in the time frame specified by the *seconds* parameter.

**Values** 10 — 1000

**Default** 100

*seconds* — Specifies the time frame in seconds that is used to limit the number of destination unreachable ICMPv6 messages to be issued.

**Values** 1 — 60

**Default** 10

## local-proxy-nd

**Syntax** [no] **local-proxy-nd**

**Context** config>service>ies>if>ipv6

**Description** When this command is enabled, the interface will reply to neighbor solicitation requests when both the hosts are on the same subnet. In this case, ICMP redirects will be disabled. When this command is disabled, the interface will not reply to neighbor solicitation requests if both the hosts are on the same subnet.

**Default** disabled

## neighbor

<b>Syntax</b>	<b>neighbor</b> <i>ipv6-address mac-address</i> <b>no neighbor</b> <i>ipv6-address</i>
<b>Context</b>	config>service>ies>if>ipv6
<b>Description</b>	This command configures IPv6-to-MAC address mapping on the IES interface.
<b>Parameters</b>	<p><i>ipv6-address</i> — The IPv6 address of the interface for which to display information.</p> <p><b>Values</b></p> <p>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  prefix-length [1..128]</p> <p><i>mac-address</i> — Specifies the 48-bit MAC address for the IPv6-to-MAC address mapping in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee and ff are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.</p>

## proxy-nd-policy

<b>Syntax</b>	<b>proxy-nd-policy</b> <i>policy-name [policy-name...(up to 5 max)]</i> <b>no proxy-nd-policy</b>
<b>Context</b>	config>service>ies>if>ipv6
<b>Description</b>	<p>This command configures a proxy neighbor discovery policy for the interface. This policy determines networks and sources for which proxy ND will be attempted, when local proxy neighbor discovery is enabled.</p> <p>The <b>no</b> form of this command reverts to the default value.</p>
<b>Default</b>	no proxy-nd-policy
<b>Parameters</b>	<p><i>policy-name</i> — The export route policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters excluding double quotes. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.</p> <p>The specified name(s) must already be defined.</p>

## Show Commands

### egress-label

<b>Syntax</b>	<b>egress-label</b> <i>egress-label1</i> [ <i>egress-label2</i> ]
<b>Context</b>	show>service
<b>Description</b>	<p>Display services using the range of egress labels.</p> <p>If only the mandatory <i>egress-label1</i> parameter is specified, only services using the specified label are displayed.</p> <p>If both <i>egress-label1</i> and <i>egress-label2</i> parameters are specified, the services using the range of labels X where <i>egress-label1</i> &lt;= X &lt;= <i>egress-label2</i> are displayed.</p> <p>Use the <b>show router ldp bindings</b> command to display dynamic labels.</p>
<b>Parameters</b>	<p><i>egress-label1</i> — The starting egress label value for which to display services using the label range. If only <i>egress-label1</i> is specified, services only using <i>egress-label1</i> are displayed.</p> <p><b>Values</b> 0, 2049 — 131071</p> <p><i>egress-label2</i> — The ending egress label value for which to display services using the label range.</p> <p><b>Default</b> The <i>egress-label1</i> value.</p> <p><b>Values</b> 2049 — 131071</p>
<b>Output</b>	<b>Show Service Egress Command Output</b> — The following table describes show service egress label output fields.

Label	Description
Svc Id	The ID that identifies a service.
Sdp Id	The ID that identifies an SDP.
Type	Indicates whether the SDP binding is a spoke or a mesh.
I. Lbl	The VC label used by the far-end device to send packets to this device in this service by the SDP.
E. Lbl	The VC label used by this device to send packets to the far-end device in this service by the SDP.
Number of bindings found	The total number of SDP bindings that exist within the specified egress label range.

**Sample Output**

```

A:ALA-12# show service egress-label 0 10000
=====
Martini Service Labels
=====
Svc Id      Sdp Id      Type I.Lbl      E.Lbl
-----
1           10:1        Mesh 0          0
1           20:1        Mesh 0          0
1           30:1        Mesh 0          0
1           100:1       Mesh 0          0
...
1           107:1       Mesh 0          0
1           108:1       Mesh 0          0
1           300:1       Mesh 0          0
1           301:1       Mesh 0          0
1           302:1       Mesh 0          0
1           400:1       Mesh 0          0
1           500:2       Spok 131070     2001
1           501:1       Mesh 131069     2000
100         300:100     Spok 0          0
200         301:200     Spok 0          0
300         302:300     Spok 0          0
400         400:400     Spok 0          0
-----
Number of Bindings Found : 23
=====
A:ALA-12#

```

**fdb-info**

<b>Syntax</b>	<b>fdb-info</b>
<b>Context</b>	show>service
<b>Description</b>	This command displays global FDB usage information.
<b>Output</b>	<b>Show FDB-Info Command Output</b> — The following table describes show FDB-Info command output.

Label	Description
Service ID	The ID that identifies a service.
Mac Move	Indicates the administrative state of the MAC movement feature associated with the service.
Mac Move Rate	The maximum rate at which MACs can be re-learned in this service, before the SAP where the moving MAC was last seen is automatically disabled in order to protect the system against undetected loops or duplicate MACs. The rate is computed as the maximum number of re-learns allowed in a 5 second interval: for example, the default rate of 10 re-learns per second corresponds to 50 re-learns in a 5 second period.

Label	Description (Continued)
Mac Move Timeout	Indicates the time in seconds to wait before a SAP that has been disabled after exceeding the maximum re-learn rate is re-enabled. A value of zero indicates that the SAP will not be automatically re-enabled after being disabled. If after the SAP is re-enabled it is disabled again, the effective retry timeout is doubled in order to avoid thrashing.
Table Size	The maximum number of learned and static entries allowed in the FDB.
Total Count	The current number of entries (both learned and static) in the FDB.
Learned Count	The current number of learned entries in the FDB.
Static Count	The current number of static entries in the FDB.
Remote Age	The number of seconds used to age out FDB entries learned on an SDP. These entries correspond to MAC addresses learned on remote SAPs.
Local Age	The number of seconds used to age out FDB entries learned on local SAPs.
High WaterMark	The utilization of the FDB table of this service at which a 'table full' alarm is raised by the agent.
Low WaterMark	The utilization of the FDB table of this service at which a 'table empty' alarm is raised by the agent.
Mac Learning	Specifies whether the MAC learning process is enabled in this service.
Discard Unknown	Specifies whether frames received with an unknown destination MAC are discarded in this service.
MAC Aging	Specifies whether the MAC aging process is enabled in this service.
MAC Pinning	Specifies whether MAC Pinning is enabled in this service.
Relearn Only	When enabled, indicates that either the FDB table of this service is full or that the maximum system-wide number of MACs supported by the agent has been reached, and thus MAC learning is temporary disabled, and only MAC re-learns can take place.
Total Service FDBs	The current number of service FDBs configured on this node.
Total FDB Size	The sum of configured FDBs.
Total FDB Entries In Use	The total number of entries (both learned and static) in use.

**Sample Output**

```

A:ALA-12# show service fdb-info
=====
Forwarding Database (FDB) Information
=====
Service Id      : 700                Mac Move      : Disabled
Mac Move Rate   : 10                Mac Move Timeout : 10
Table Size      : 250               Total Count    : 0
Learned Count   : 0                Static Count    : 0
Remote Age      : 900               Local Age      : 300
High WaterMark  : 95%              Low Watermark   : 90%
Mac Aging       : Enabl             Relearn Only    : False

Service Id      : 725                Mac Move      : Disabled
Mac Move Rate   : 10                Mac Move Timeout : 10
Table Size      : 250               Total Count    : 0
Learned Count   : 0                Static Count    : 0
Remote Age      : 900               Local Age      : 300
High WaterMark  : 95%              Low Watermark   : 90%
Mac Learning    : Enabl             Discard Unknown : Dsabl
Mac Aging       : Enabl             Relearn Only    : False

Service Id      : 740                Mac Move      : Disabled
Mac Move Rate   : 10                Mac Move Timeout : 10
Table Size      : 250               Total Count    : 0
Learned Count   : 0                Static Count    : 0
Remote Age      : 900               Local Age      : 300
High WaterMark  : 95%              Low Watermark   : 90%
Mac Learning    : Enabl             Discard Unknown : Dsabl
Mac Aging       : Enabl             Relearn Only    : False
...
-----
Total Service FDBs : 7
Total FDB Configured Size : 1750
Total FDB Entries In Use : 0
-----
=====
*A:ALA-48#

```

**fdb-mac**

<b>Syntax</b>	<b>fdb-mac</b> <i>ieee-address</i>
<b>Context</b>	show>service
<b>Description</b>	Displays the FDB entry for a given MAC address.
<b>Parameters</b>	<i>ieee-address</i> — The 48-bit MAC address for which to display the FDB entry in the form <i>aa:bb:cc:dd:ee:ff</i> or <i>aa-bb-cc-dd-ee-ff</i> where <i>aa</i> , <i>bb</i> , <i>cc</i> , <i>dd</i> , <i>ee</i> and <i>ff</i> are hexadecimal numbers.
<b>Output</b>	<b>Show FDB-MAC Command Output</b> — The following table describes the show FDB MAC command output fields:

Label	Description
Service ID	The value that identifies a specific service.
MAC	The specified MAC address
Source-Identifier	The location where the MAC is defined.
Type	Static — FDB entries created by management. Learned — Dynamic entries created by the learning process. OAM — Entries created by the OAM process.

### Sample Output

```
A:ALA-12# show service fdb-mac 00:99:00:00:00:00
=====
Services Using Forwarding Database Mac 00:99:00:00:00:00
=====
ServId  MAC                               Source-Identifier      Type/Age Last Change
-----
1       00:99:00:00:00:00                  sap:1/2/7:0           Static
=====
A:ALA-12#
```

## ingress-label

<b>Syntax</b>	<b>ingress-label</b> <i>ingress-label1</i> [ <i>ingress-label2</i> ]
<b>Context</b>	show>service
<b>Description</b>	<p>Display services using the range of ingress labels.</p> <p>If only the mandatory <i>ingress-label1</i> parameter is specified, only services using the specified label are displayed.</p> <p>If both <i>ingress-label1</i> and <i>ingress-label2</i> parameters are specified, the services using the range of labels X where <i>ingress-label1</i> &lt;= X &lt;= <i>ingress-label2</i> are displayed.</p> <p>Use the <b>show router ldp bindings</b> command to display dynamic labels.</p>
<b>Parameters</b>	<p><i>ingress-label1</i> — The starting ingress label value for which to display services using the label range. If only <i>ingress-label1</i> is specified, services only using <i>ingress-label1</i> are displayed.</p> <p><b>Values</b> 0, 2048 — 131071</p> <p><i>ingress-label2</i> — The ending ingress label value for which to display services using the label range.</p> <p><b>Default</b> The <i>ingress-label1</i> value.</p> <p><b>Values</b> 2048 — 131071</p>
<b>Output</b>	<b>Show Service Ingress Label</b> — The following table describes show service ingress label output fields:

Label	Description
Svc ID	The value that identifies a specific service.
SDP Id	The SDP identifier.
Type	Indicates whether the SDP is a spoke or a mesh.
I.Lbl	The ingress label used by the far-end device to send packets to this device in this service by the SDP.
E.Lbl	The egress label used by this device to send packets to the far-end device in this service by the SDP.
Number of Bindings Found	The number of SDP bindings within the label range specified.

### Sample Output

```

A:ALA-12# show service ingress label 0
=====
Martini Service Labels
=====
Svc Id      Sdp Id      Type I.Lbl      E.Lbl
-----
1           10:1        Mesh 0         0
1           20:1        Mesh 0         0
1           30:1        Mesh 0         0
1           50:1        Mesh 0         0
1           100:1       Mesh 0         0
1           101:1       Mesh 0         0
1           102:1       Mesh 0         0
1           103:1       Mesh 0         0
1           104:1       Mesh 0         0
1           105:1       Mesh 0         0
1           106:1       Mesh 0         0
1           107:1       Mesh 0         0
1           108:1       Mesh 0         0
1           300:1       Mesh 0         0
1           301:1       Mesh 0         0
1           302:1       Mesh 0         0
1           400:1       Mesh 0         0
1           500:2       Spok 131070    2001
1           501:1       Mesh 131069    2000
100         300:100     Spok 0         0
200         301:200     Spok 0         0
300         302:300     Spok 0         0
400         400:400     Spok 0         0
-----
Number of Bindings Found : 23
-----
A:ALA-12#

```



## sap-using

<b>Syntax</b>	<b>sap-using</b> [ <b>sap</b> <i>sap-id</i> ] <b>sap-using interface</b> [ <i>ip-address</i>   <i>ip-int-name</i> ] <b>sap-using</b> [ <b>ingress</b>   <b>egress</b> ] <b>filter</b> <i>filter-id</i> <b>sap-using</b> [ <b>ingress</b>   <b>egress</b> ] <b>qos-policy</b> <i>qos-policy-id</i> <b>sap-using authentication-policy</b> <i>auth-plcy-name</i>
<b>Context</b>	show>service
<b>Description</b>	<p>Displays SAP information.</p> <p>If no optional parameters are specified, the command displays a summary of all defined SAPs.</p> <p>The optional parameters restrict output to only SAPs matching the specified properties.</p>
<b>Parameters</b>	<p><b>ingress</b> — Specifies matching an ingress policy.</p> <p><b>egress</b> — Specifies matching an egress policy.</p> <p><b>qos-policy</b> <i>qos-policy-id</i> — The ingress or egress QoS Policy ID for which to display matching SAPs.</p> <p><b>Values</b> 1 — 65535</p> <p><b>filter</b> <i>filter-id</i> — The ingress or egress filter policy ID for which to display matching SAPs.</p> <p><b>Values</b> 1 — 65535</p> <p><b>authentication</b> <i>auth-plcy-name</i> — The session authentication policy for which to display matching SAPs.</p> <p><i>sap-id</i> — Specifies the physical port identifier portion of the SAP definition. See <a href="#">Common Service Commands on page 2062</a> for <i>sap-id</i> command syntax.</p> <p><b>interface</b> — Specifies matching SAPs with the specified IP interface.</p> <p><i>ip-addr</i> — The IP address of the interface for which to display matching SAPs.</p> <p><b>Values</b> 1.0.0.0 — 223.255.255.255</p> <p><i>ip-int-name</i> — The IP interface name for which to display matching SAPs.</p>
<b>Output</b>	<b>Show Service SAP</b> — The following table describes show service SAP output fields:

Label	Description
Port ID	The ID of the access port where the SAP is defined.
Svc ID	The service identifier.
SapMTU	The SAP MTU value.
I.QoS	The SAP ingress QoS policy number specified on the ingress SAP.
I.MAC/IP	The MAC or IP filter policy ID applied to the ingress SAP.
E.QoS	The SAP egress QoS policy number specified on the egress SAP.

Label	Description (Continued)
E.Mac/IP	The MAC or IP filter policy ID applied to the egress SAP
A.Pol	The accounting policy ID assigned to the SAP.
Adm	The administrative state of the SAP.
Opr	The operational state of the SAP.

### Sample Output

```
A:ALA-12# show service sap-using sap 1/1
=====
Service Access Points
=====
PortId          SvcId      SapMTU  I.QoS  I.Mac/IP  E.QoS  E.Mac/IP  A.Pol  Adm  Opr
-----
1/1/7:0         1          1518    10     8          10     none      none   Up   Up
1/1/11:0        100         1514    1     none       1     none      none   Down Down
1/1/7:300       300         1518    10     none       10     none      1000   Up   Up
-----
Number of SAPs : 3
-----
A:ALA-12#
ALA-1# show service sap-using authentication-policy dhcp-radius-plcy
=====
Service Access Point Using Authentication Policy dhcp-radius-plcy
=====
PortId          SvcId      I.QoS  I.Fltr  E.QoS  E.Fltr  Acc.Pol  Adm  Opr
-----
1/1/1           1          1      none    1      none    none     Up   Down
-----
Saps : 1
-----
A:ALA-1#
```

## sdp

<b>Syntax</b>	<b>sdp</b> [ <i>sdp-id</i>   <b>far-end</b> <i>ip-addr</i> ] [ <b>detail</b> ]
<b>Context</b>	show>service>id
<b>Description</b>	Displays information for the SDPs associated with the service. If no optional parameters are specified, a summary of all associated SDPs is displayed.
<b>Parameters</b>	<p><i>sdp-id</i> — Displays only information for the specified SDP ID.</p> <p><b>Default</b> All SDPs</p> <p><b>Values</b> 1 — 17407</p> <p><b>far-end</b> <i>ip-addr</i> — Displays only SDPs matching with the specified far-end IP address.</p> <p><b>Default</b> SDPs with any far-end IP address.</p>

**detail** — Displays detailed SDP information.

**Output Show Service-ID SDP** — The following table describes show service-id SDP output fields:

Label	Description
Sdp Id	The SDP identifier.
Type	Indicates whether the SDP is a spoke or a mesh.
Split Horizon Group	Name of the split horizon group that the SAP belongs.
VC Type	Displays the VC type: ether, vlan, or vpls.
VC Tag	Displays the explicit dot1Q value used when encapsulating to the SDP far end.
I. Lbl	The VC label used by the far-end device to send packets to this device in this service by the SDP.
Admin Path MTU	The operating path MTU of the SDP is equal to the admin path MTU (when one is set) or the dynamically computed tunnel MTU, when no admin path MTU is set (the default case.)
Oper Path MTU	The actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Far End	Specifies the IP address of the remote end of the GRE or MPLS tunnel defined by this SDP.
Delivery	Specifies the type of delivery used by the SDP: GRE or MPLS.
Admin State	The administrative state of this SDP.
Oper State	The current status of the KeepAlive protocol.
Ingress Label	The label used by the far-end device to send packets to this device in this service by this SDP.
Egress Label	The label used by this device to send packets to the far-end device in this service by the SDP.
Last Changed	The date and time of the most recent change to the SDP.
Signaling	Specifies the signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received on this SDP.
Admin State	The administrative state of the keepalive process.
Oper State	The operational state of the keepalive process.
Hello Time	Specifies how often the SDP echo request messages are transmitted.

Label	Description (Continued)
Max Drop Count	Specifies the maximum number of consecutive SDP echo request messages that can be unacknowledged before the keepalive protocol reports a fault.
Hello Msg Len	Specifies the length of the transmitted SDP echo request messages.
Hold Down Time	Specifies the amount of time to wait before the keepalive operating status is eligible to enter the alive state.
I. Fwd. Pkts.	Specifies the number of forwarded ingress packets.
I. Dro. Pkts	Specifies the number of dropped ingress packets.
E. Fwd. Pkts.	Specifies the number of forwarded egress packets.
E. Fwd. Octets	Specifies the number of forwarded egress octets.
Associated LSP List	When the SDP type is MPLS, a list of LSPs used to reach the far-end router displays. All the LSPs in the list must terminate at the IP address specified in the far end field. If the SDP type is GRE, then the following message displays: SDP delivery mechanism is not MPLS

### Sample Output

```
*A:ALA-48# show service id 700 sdp 2
=====
Service Destination Point (Sdp Id : 2)
=====
SdpId          Type IP address    Adm    Opr      I.Lbl    E.Lbl
-----
2:222          Spok 10.10.10.104    Up     Down     0        0
2:700          Mesh 10.10.10.104    Up     Down     0        0
=====
*A:ALA-48#
```

## sdp-using

<b>Syntax</b>	<b>sdp-using</b> [ <i>sdp-id[:vc-id]</i>   <i>far-end ip-address</i> ]
<b>Context</b>	show>service
<b>Description</b>	Display services using SDP or far-end address options.
<b>Parameters</b>	<i>sdp-id</i> — Displays only services bound to the specified SDP ID. <div> <b>Values</b> 1 — 17407 </div> <i>vc-id</i> — The virtual circuit identifier. <div> <b>Values</b> 1 — 4294967295 </div>

**far-end ip-address** — Displays only services matching with the specified far-end IP address.

**Default** Services with any far-end IP address.

**Output** **Show Service SDP Using X** — The following table describes **sdp-using** output fields.

Label	Description
Svc ID	The value identifying a service.
Spd ID	The SPD identifier.
Type	Type of SDP: Spoke or Mesh.
Far End	The far-end address of the SDP.
Oper State	The operational state of the service.
Ingress Label	The label used by the far-end device to send packets to this device in this service by this SDP.
Egress Label	The label used by this device to send packets to the far-end device in this service by this SDP.

### Sample Output

```
A:ALA-1# show service sdp-using 300
=====
Service Destination Point (Sdp Id : 300)
=====
SvcId      SdpId      Type Far End      Opr State I.Label  E.Label
-----
1          300:1      Mesh 10.0.0.13    Up       131071   131071
2          300:2      Spok 10.0.0.13    Up       131070   131070
100        300:100    Mesh 10.0.0.13    Up       131069   131069
101        300:101    Mesh 10.0.0.13    Up       131068   131068
102        300:102    Mesh 10.0.0.13    Up       131067   131067
-----
Number of SDPs : 5
=====
A:ALA-1#
```

## service-using

**Syntax** **service-using [ies] [vpls] [sdp sdp-id]**

**Context** show>service

**Description** This command displays the services matching certain usage properties.  
If no optional parameters are specified, all services defined on the system are displayed.

**Parameters** **ies** — Displays matching IES instances.  
**vpls** — Displays matching VPLS instances.

**sdp** *sdp-id* — Displays only services bound to the specified SDP ID.

**Default** Services bound to any SDP ID.

**Values** 1 — 17407

**Output** **Show Service Service-Using** — The following table describes show **service-using** output fields:

Label	Description
Service Id	The value that identifies a service.
Type	Specifies the service type configured for the service ID.
Adm	The administrative state of the service.
Opr	The operational state of the service.
CustomerID	The value that identifies a specific customer.
Last Mgmt Change	The date and time of the most recent management-initiated change to this service.

### Sample Output

```
*A:ALA-48>show>service# service-using vpls
=====
Services [vpls]
=====
ServiceId    Type    Adm    Opr    CustomerId    Last Mgmt Change
-----
700          VPLS    Up     Down    7             04/11/2007 09:36:36
725          VPLS    Down   Down    7             04/11/2007 09:36:36
740          VPLS    Down   Down    1             04/11/2007 09:36:36
750          VPLS    Down   Down    7             04/11/2007 09:36:36
1730         VPLS    Down   Down    1730          04/11/2007 09:36:36
9000         VPLS    Up     Down    6             04/11/2007 09:36:36
90001        VPLS    Up     Down    6             04/11/2007 09:36:36
-----
Matching Services : 7
=====
*A:ALA-48>show>service#
```

## active-subscribers

**Syntax** **active-subscribers detail**  
**active-subscribers mirror**  
**active-subscribers summary**

**Context** show>service

**Description** This command displays active subscriber information.

**Parameters** **summary** — Displays active subscriber information in a brief format.

**detail** — Displays detailed output.

### Sample Output

```
*A:Dut-C# show service active-subscribers
=====
Active Subscribers
=====
Subscriber hpolSub1 (hpolSubProf2)
-----
(1) SLA Profile Instance sap:lag-1:2000.1 - sla:hpolSlaProf1
-----
IP Address
MAC Address      PPPoE-SID Origin
-----
200.1.4.194
00:01:00:00:03:c1 1      IPCP
-----
(2) SLA Profile Instance sap:lag-1:2000.1 - sla:hpolSlaProf2
-----
IP Address
MAC Address      PPPoE-SID Origin
-----
200.1.4.35
00:01:00:00:03:22 N/A      ARP-Host
200.1.4.195
00:01:00:00:03:c2 1      IPCP
-----
Subscriber hpolSub16 (hpolSubProf1)
-----
(1) SLA Profile Instance sap:[lag-1:2000.2] - sla:hpolSlaProf1
-----
IP Address
MAC Address      PPPoE-SID Origin
-----
200.1.4.224
00:01:00:00:03:df 1      IPCP
-----
Subscriber hpolSub2 (hpolSubProf1)
-----
(1) SLA Profile Instance sap:lag-1:2000.1 - sla:hpolSlaProf1
-----
IP Address
MAC Address      PPPoE-SID Origin
-----
200.1.4.196
00:01:00:00:03:c3 1      IPCP
-----
Number of active subscribers : 3
=====
*A:Dut-C#
```

## credit-control

<b>Syntax</b>	<b>credit-control credit-control</b> [ <b>subscriber</b> <i>sub-ident-string</i> ] <b>credit-control out-of-credit</b> [ <b>action</b> <i>action</i> ] [ <b>summary</b> ]
<b>Context</b>	show>service>active-subscribers
<b>Description</b>	This command displays active subscriber credit control information.

## filter

<b>Syntax</b>	<b>filter</b> [ <b>subscriber</b> <i>sub-ident-string</i> ] [ <b>origin</b> <i>origin</i> ]
<b>Context</b>	show>service>active-subscribers
<b>Description</b>	This command displays active subscriber filter information.

## hierarchy

<b>Syntax</b>	<b>hierarchy</b> [ <b>subscriber</b> <i>sub-ident-string</i> ]
<b>Context</b>	show>service>active-subscribers
<b>Description</b>	This command displays active subscriber hierarchy information.

## host-tracking

<b>Syntax</b>	<b>host-tracking</b> [ <b>subscriber</b> <i>sub-ident-string</i> ] <b>host-tracking</b> [ <b>subscriber</b> <i>sub-ident-string</i> ] <b>detail</b> <b>host-tracking</b> [ <b>subscriber</b> <i>sub-ident-string</i> ] <b>summary</b> <b>host-tracking</b> [ <b>subscriber</b> <i>sub-ident-string</i> ] <b>statistics</b>
<b>Context</b>	show>service>active-subscribers
<b>Description</b>	This command displays active subscriber host tracking information.

## groups

<b>Syntax</b>	<b>groups</b> [ <b>group</b> <i>group-ip-address</i> ] <b>groups group</b> <i>group-ip-address</i> ] <b>detail</b> <b>groups group</b> <i>group-ip-address</i> ] <b>summary</b>
<b>Context</b>	show>service>active-subscribers
<b>Description</b>	This command displays active subscriber host tracking groups information.



## igmp

<b>Syntax</b>	<b>igmp [subscriber <i>sub-ident-string</i>][detail]</b>
<b>Context</b>	show>service>active-subscribers
<b>Description</b>	This command displays active subscriber IGMP information.

**Sample Output**

```
*B:Dut-C# show service active-subscribers igmp
=====
Active Subscribers
=====
Subscriber                               IGMP-Policy
  HostAddr                               GrpItf                               NumGroups
-----
sub_1                                     poll
  112.112.1.1                             gi_1_1                               1
  112.112.1.2                             gi_1_1                               2
sub_2                                     poll
  112.112.1.3                             gi_1_2                               0
-----
Number of Subscribers : 2
=====
*B:Dut-C#

*B:Dut-C# show service active-subscribers igmp detail
=====
Active Subscribers Detail
=====
Subscriber                               IGMP-Policy
  HostAddr                               GrpItf                               NumGroups
  GrpAddr                               Type                               Up-Time                               Mode
  SrcAddr                               Type                               Blk/Fwd
-----
sub_1                                     poll
  112.112.1.1                             gi_1_1                               1
  225.0.0.1                             Dynamic                             0d 00:00:53                         Include
  11.11.0.1                             Dynamic                             Fwd
  11.11.0.2                             Dynamic                             Fwd
  112.112.1.2                             gi_1_1                               2
  225.0.0.1                             Dynamic                             0d 00:00:44                         Exclude
  11.11.0.1                             Dynamic                             Blk
  225.0.0.2                             Dynamic                             0d 00:00:44                         Exclude
  11.11.0.1                             Dynamic                             Blk
sub_2                                     poll
  112.112.1.3                             gi_1_2                               0
-----
Number of Subscribers : 2
=====
*B:Dut-C#

*B:Dut-C# show service active-subscribers igmp subscriber "sub_1" detail
=====
Active Subscribers Detail
=====
Subscriber                               IGMP-Policy
```

## Show Commands

```

HostAddr      GrpAddr      SrcAddr      GrpItf      Type      Up-Time      NumGroups
              Type      Mode
              Type      Blk/Fwd
-----
sub_1          poll
112.112.1.1    gi_1_1
225.0.0.1      Dynamic      0d 00:01:04    1
11.11.0.1      Dynamic      Fwd
11.11.0.2      Dynamic      Fwd
112.112.1.2    gi_1_1
225.0.0.1      Dynamic      0d 00:00:55    2
11.11.0.1      Dynamic      Exclude
225.0.0.2      Dynamic      0d 00:00:55    Blk
11.11.0.1      Dynamic      Exclude
11.11.0.1      Dynamic      Blk
-----
Number of Subscribers : 1
=====
B:Dut-C#

```

## subscriber

**Syntax**

```

subscriber sub-ident-string
subscriber sub-ident-string detail
subscriber sub-ident-string mirror
subscriber sub-ident-string sap sap-id sla-profile sla-profile-name
subscriber sub-ident-string sap sap-id sla-profile sla-profile-name detail
subscriber sub-ident-string sap sap-id sla-profile sla-profile-name mirror

```

**Context** show>service>active-subscribers

**Description** This command displays active subscriber information for a subscriber.

### Sample Output

```

*A:Dut-C# show service active-subscribers subscriber "hpolSub1"
=====
Active Subscribers
=====
Subscriber hpolSub1 (hpolSubProf2)
-----
(1) SLA Profile Instance sap:lag-1:2000.1 - sla:hpolSlaProf1
-----
IP Address      MAC Address      PPPoE-SID Origin
-----
200.1.4.194      00:01:00:00:03:c1 1      IPCP
-----
(2) SLA Profile Instance sap:lag-1:2000.1 - sla:hpolSlaProf2
-----
IP Address      MAC Address      PPPoE-SID Origin
-----
200.1.4.35

```

```

00:01:00:00:03:22 N/A      ARP-Host
200.1.4.195
00:01:00:00:03:c2 1      IPCP
=====
*A:Dut-C#

*A:Dut-C# show service active-subscribers subscriber "hpolSub1"
sap lag-1:2000.1 sla-profile "hpolSlaProf2"
=====
Active Subscribers
-----
Subscriber hpolSub1 (hpolSubProf2)
-----
(1) SLA Profile Instance sap:lag-1:2000.1 - sla:hpolSlaProf2
-----
IP Address
MAC Address      PPPoE-SID Origin
-----
200.1.4.35
00:01:00:00:03:22 N/A      ARP-Host
200.1.4.195
00:01:00:00:03:c2 1      IPCP
=====
*A:Dut-C#

*A:Dut-C# show service active-subscribers subscriber "hpolSub1"
sap lag-1:2000.1 sla-profile "hpolSlaProf1" detail
=====
Active Subscribers
-----
Subscriber hpolSub1 (hpolSubProf2)
-----
I. Sched. Policy : N/A
E. Sched. Policy : N/A
I. Policer Ctrl. : N/A
E. Policer Ctrl. : N/A
Q Frame-Based Ac*: Disabled
Acct. Policy      : N/A
Rad. Acct. Pol.   : hpolRadAcctPol
Dupl. Acct. Pol.  : N/A
ANCP Pol.         : N/A
HostTrk Pol.      : N/A
IGMP Policy       : N/A
Sub. MCAC Policy  : N/A
NAT Policy        : N/A
Def. Encap Offset : none
Avg Frame Size    : 120
Sub. ANCP-String  : "hpolSub1"
Sub. Int Dest Id  : "2000"
Host Trk Rate Adj : N/A
RADIUS Rate-Limit : 10220541
Oper-Rate-Limit   : 10220541
-----
(1) SLA Profile Instance
- sap:lag-1:2000.1 (VPRN 2000 - grp-Vprn-1)
- sla:hpolSlaProf1

```

## Show Commands

```

-----
Description          : SLA Profile Id hpolSlaProf1
Host Limit           : No Limit
Ingress Qos-Policy   : 2                      Egress Qos-Policy : 2 (vport)
Ingress Queuing Type : Service-queuing (non 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  : N/A
Egress Report-Rate   : N/A
Egress Remarking     : from Sap Qos
Credit Control Pol. : N/A
-----

IP Address          MAC Address          PPPoE-SID Origin
-----
200.1.4.194        00:01:00:00:03:c1 1          IPCP
-----

SLA Profile Instance statistics
-----
                Packets                Octets
Off. HiPrio      : 0                    0
Off. LowPrio     : 0                    0
Off. Uncolor     : 0                    0

Queueing Stats (Ingress QoS Policy 2)
Dro. HiPrio      : 0                    0
Dro. LowPrio     : 0                    0
For. InProf      : 0                    0
For. OutProf     : 0                    0

Queueing Stats (Egress QoS Policy 2)
Dro. InProf      : 0                    0
Dro. OutProf     : 0                    0
For. InProf      : 0                    0
For. OutProf     : 0                    0
-----

SLA Profile Instance per Queue statistics
-----
                Packets                Octets

Egress Queue 1
Dro. InProf      : 0                    0
Dro. OutProf     : 0                    0
For. InProf      : 0                    0
For. OutProf     : 0                    0

Egress Queue 2
Dro. InProf      : 0                    0
Dro. OutProf     : 0                    0
For. InProf      : 0                    0
For. OutProf     : 0                    0

Egress Queue 3
Dro. InProf      : 0                    0
Dro. OutProf     : 0                    0
For. InProf      : 0                    0
For. OutProf     : 0                    0

Egress Queue 4
Dro. InProf      : 0                    0

```

```

Dro. OutProf      : 0          0
For. InProf       : 0          0
For. OutProf      : 0          0

```

## Egress Queue 5

```

Dro. InProf       : 0          0
Dro. OutProf      : 0          0
For. InProf       : 0          0
For. OutProf      : 0          0

```

## Egress Queue 6

```

Dro. InProf       : 0          0
Dro. OutProf      : 0          0
For. InProf       : 0          0
For. OutProf      : 0          0

```

## Egress Queue 7

```

Dro. InProf       : 0          0
Dro. OutProf      : 0          0
For. InProf       : 0          0
For. OutProf      : 0          0

```

## Egress Queue 8

```

Dro. InProf       : 0          0
Dro. OutProf      : 0          0
For. InProf       : 0          0
For. OutProf      : 0          0

```

-----  
SLA Profile Instance per Policer statistics

```

-----
                Packets          Octets
Ingress Policer 1 (Stats mode: minimal)
Off. All        : 0              0
Dro. All        : 0              0
For. All        : 0              0

```

```

=====
*A:Dut-C#

```

id

<b>Syntax</b>	<b>id</b> <i>service-id</i>
<b>Context</b>	show>service
<b>Description</b>	Enables the context to display information for a particular service-id.
<b>Parameters</b>	<p><i>service-id</i> — The unique service identification number that identifies the service in the service domain.</p> <p><b>all</b> — Displays detailed information about the service.</p> <p><b>arp</b> — Displays ARP entries for the service.</p> <p><b>arp-host</b> — Displays ARP host related information.</p> <p><b>base</b> — Displays basic service information.</p> <p><b>fdb</b> — Displays FDB entries.</p>

- host** — Displays static hosts configured on the specified service.
- labels** — Displays labels being used by this service.
- sap** — Displays SAPs associated to the service.
- sdp** — Displays SDPs associated with the service.
- split-horizon-group** — Display split horizon group information.
- stp** — Displays STP information.

all

<b>Syntax</b>	<b>all</b>
<b>Context</b>	show>service>id
<b>Description</b>	Displays detailed information for all aspects of the service.
<b>Output</b>	<b>Show All Service-ID Output</b> — The following table describes the show all service-id command output fields:

Label	Description
Service Id	The value that identifies a service.
VPN Id	The number that identifies the VPN.
Service Type	Specifies the type of service.
SDP Id	The SDP identifier.
Description	Text string describing general information about the service.
Customer Id	The customer identifier.
Last Mgmt Change	The date and time of the most recent management-initiated change to this customer.
SAP Count	The number of SAPs specified for this service.
SDP Bind Count	The number of SDPs bound to this service.
Split Horizon Group	Name of the split horizon group for this service.
Description	Text string describing the split horizon group.
Last Changed	The date and time of the most recent management-initiated change to this split horizon group.
SDP Id	The SDP identifier.
Type	Indicates whether this Service SDP binding is a spoke or a mesh.

Label	Description (Continued)
Admin Path MTU	The desired largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Oper Path MTU	The actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Delivery	Specifies the type of delivery used by the SDP: GRE or MPLS.
Admin State	The administrative state of this SDP.
Oper State	The operational status of the KeepAlive protocol.
Ingress Label	The label used by the far-end device to send packets to this device in this service by this SDP.
Egress Label	The label used by this device to send packets to the far-end device in this service by this SDP.
Ingress Filter	The ID of the ingress filter policy.
Egress Filter	The ID of the egress filter policy.
Far End	Specifies the IP address of the remote end of the GRE or MPLS tunnel defined by this SDP.
Last Changed	The date and time of the most recent change to this customer.
Signaling	Specifies the signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received on this SDP.
Admin State	Specifies the operating status of the SDP.
Oper State	The operational state of the SDP.
Hello Time	Specifies how often the SDP echo request messages are transmitted on this SDP.
Hello Msg Len	Specifies the length of the SDP echo request messages transmitted on this SDP.
Max Drop Count	Specifies the maximum number of consecutive SDP echo request messages that can be unacknowledged before the keepalive protocol reports a fault.
Hold Down Time	Specifies the amount of time to wait before the keepalive operating status is eligible to enter the alive state.
SDP Delivery Mechanism	When the SDP type is MPLS, a list of LSPs used to reach the far-end router displays. All the LSPs in the list must terminate at the IP address specified in the Far End field. If the SDP type is GRE, then the following message displays: SDP Delivery Mechanism is not MPLS

Label	Description (Continued)
Number of SDPs	The total number SDPs applied to this service ID.
Service Access Points	
Service Id	The value that identifies a service.
Port Id	The ID of the access port where this SAP is defined.
Description	Generic information about the SAP.
Encap Value	The value of the label used to identify this SAP on the access port.
Admin State	The desired state of the SAP.
Oper State	The operating state of the SAP.
Last Changed	The date and time of the last change.
Admin MTU	The desired largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Oper MTU	The actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Ingress qos-pol- icy	The SAP ingress QoS policy ID.
Egress qos-policy	The SAP egress QoS policy ID.
Ingress Filter-Id	The SAP ingress filter policy ID.
Egress Filter-Id	The SAP egress filter policy ID.
Multi Svc Site	Indicates the multi-service site in which the SAP is a member.
Ingress sched- policy	Indicates the ingress QoS scheduler for the SAP.
Egress sched-pol- icy	Indicates the egress QoS scheduler for the SAP.
Acct. Pol	Indicates the accounting policy applied to the SAP.
Collect Stats	Specifies whether accounting statistics are collected on the SAP.
Dropped	The number of packets or octets dropped.
Offered Hi Prior- ity	The number of high priority packets, as determined by the SAP ingress QoS policy.
Offered Low Prior- ity	The number of low priority packets, as determined by the SAP ingress QoS policy.
Forwarded In Pro- file	The number of in-profile packets or octets (rate below CIR) forwarded.



Label	Description (Continued)
Forwarded Out Profile	The number of out-of-profile packets or octets (rate above CIR) forwarded.
Dropped In Profile	The number of in-profile packets or octets discarded.
Dropped Out Profile	The number of out-of-profile packets or octets discarded.
Forwarded In Profile	The number of in-profile packets or octets (rate below CIR) forwarded.
Forwarded Out Profile	The number of out-of-profile packets or octets (rate above CIR) forwarded.
Ingress Queue 1	The index of the ingress QoS queue of this SAP.
High priority offered	The packets or octets count of the high priority traffic for the SAP.
High priority dropped	The number of high priority traffic packets/octets dropped.
Low priority offered	The packets or octets count of the low priority traffic.
Low priority dropped	The number of low priority traffic packets/octets dropped.
In profile forwarded	The number of in-profile packets or octets (rate below CIR) forwarded.
Out profile forwarded	The number of out-of-profile octets (rate above CIR) forwarded.
Egress Queue 1	The index of the egress QoS queue of the SAP.
In profile forwarded	The number of in-profile packets or octets (rate below CIR) forwarded.
In profile dropped	The number of in-profile packets or octets dropped for the SAP.
Out profile forwarded	The number of out-of-profile packets or octets (rate above CIR) forwarded.
Out profile dropped	The number of out-of-profile packets or octets discarded.
State	Specifies whether DHCP relay is enabled on this SAP.
Info Option	Specifies whether Option 82 processing is enabled on this SAP.
Action	Specifies the Option 82 processing on this SAP or interface: keep, replace or drop.
Circuit ID	Specifies whether the If index is inserted in Circuit ID sub-option of Option 82.
Remote ID	Specifies whether the far-end MAC address is inserted in remote ID sub-option of Option 82

Label	Description (Continued)
Managed by Service	Specifies the service-id of the management VPLS managing this SAP.
Managed by SAP	Specifies the sap-id inside the management VPLS managing this SAP.
Prune state	Specifies the STP state inherited from the management VPLS.
Managed by Service	Specifies the service-id of the management VPLS managing this spoke SDP.
Managed by Spoke	Specifies the sap-id inside the management VPLS managing this spoke SDP.
Prune state	Specifies the STP state inherited from the management VPLS.

### Sample Output

```

A:ALA-12# show service id 9000 all
=====
Service Detailed Information
=====
Service Id       : 9000                Vpn Id           : 0
Service Type     : VPLS
Description      : This is a distributed VPLS.
Customer Id      : 6                  Last Mgmt Change  : 01/18/2007 10:31:58
Adm              : Up                  Oper              : Down
MTU              : 1514                Def. Mesh VC Id   : 750
SAP Count        : 1                  SDP Bind Count    : 3
-----
Split Horizon Group specifics
-----
Split Horizon Group : splitgroup1
-----
Description       : Split horizon group 1
Instance Id       : 1                  Last Changed      : 01/18/2007 10:31:58
-----
Service Destination Points (SDPs)
-----
Sdp Id 2:22  -(10.10.10.104)
-----
Description       : GRE-10.10.10.104
SDP Id           : 2:22                Type              : Spoke
VC Type          : Ether              VC Tag            : n/a
Admin Path MTU   : 0                  Oper Path MTU     : 0
Far End          : 10.10.10.104        Delivery          : GRE
Flags            : SdpOperDown
                  NoIngVCLabel NoEgrVCLabel
                  PathMTUTooSmall
Admin State      : Up                  Oper State         : Down
Ingress Label    : 0                  Egress Label      : 0
Ingress Filter   : n/a                Egress Filter     : n/a
Last Changed     : 01/18/2007 10:31:58 Signaling          : TLDP
-----
KeepAlive Information :
Admin State       : Disabled           Oper State         : Disabled
Hello Time        : 10                 Hello Msg Len      : 0
Max Drop Count    : 3                 Hold Down Time     : 10

```

```

Statistics          :
I. Fwd. Pkts.      : 0
E. Fwd. Pkts.      : 0
I. Dro. Pkts.      : 0
E. Fwd. Octets     : 0

```

```

Associated LSP LIST :
SDP Delivery Mechanism is not MPLS

```

```

-----
Rstp Service Destination Point specifics
-----

```

```

Mac Move           : Blockable
Rstp Admin State   : Up
Core Connectivity   : Down
Port Role          : N/A
Port Number        : 2049
Port Path Cost     : 10
Admin Edge         : Disabled
Link Type          : Pt-pt
Designated Bridge  : N/A
Active Protocol    : N/A

Rstp Oper State    : Down
Port State         : Discarding
Port Priority      : 128
Auto Edge         : Enabled
Oper Edge         : N/A
BPDU Encap        : Dot1d
Designated Port Id: 0

Fwd Transitions    : 0
Cfg BPDUs rcvd     : 0
TCN BPDUs rcvd     : 0
RST BPDUs rcvd     : 0
Bad BPDUs rcvd     : 0
Cfg BPDUs tx       : 0
TCN BPDUs tx       : 0
RST BPDUs tx       : 0

```

```

-----
Sdp Id 5:750 - (128.251.10.49)
-----

```

```

SDP Id            : 5:750
VC Type           : Ether
Admin Path MTU    : 0
Far End           : 128.251.10.49
Flags             : SdpOperDown
                  NoIngVCLabel NoEgrVCLabel
                  PathMTUTooSmall
Admin State       : Up
Ingress Label     : 0
Ingress Filter    : n/a
Last Changed      : 01/18/2007 10:31:58
Type              : Mesh
VC Tag            : n/a
Oper Path MTU     : 0
Delivery          : GRE
Oper State        : Down
Egress Label      : 0
Egress Filter     : n/a
Signaling         : TLDP

KeepAlive Information :
Admin State         : Disabled
Hello Time          : 10
Max Drop Count      : 3
Oper State          : Disabled
Hello Msg Len       : 0
Hold Down Time      : 10

```

```

Statistics          :
I. Fwd. Pkts.      : 0
E. Fwd. Pkts.      : 0
I. Dro. Pkts.      : 0
E. Fwd. Octets     : 0

```

```

Associated LSP LIST :
SDP Delivery Mechanism is not MPLS

```

```

-----
Sdp Id 7:750 - (10.10.10.106)
-----

```

```

Description       : to-MPLS-10.10.10.49
SDP Id            : 7:750
VC Type           : Ether
Admin Path MTU    : 0
Far End           : 10.10.10.106
Flags             : SdpOperDown
                  NoIngVCLabel NoEgrVCLabel
Type              : Mesh
VC Tag            : n/a
Oper Path MTU     : 0
Delivery          : MPLS

```

## Show Commands

```

PathMTUTooSmall
Admin State      : Up
Ingress Label    : 0
Ingress Filter   : n/a
Last Changed     : 01/18/2007 10:31:58
Oper State       : Down
Egress Label     : 0
Egress Filter    : n/a
Signaling        : TLDP

KeepAlive Information :
Admin State       : Disabled
Hello Time        : 10
Max Drop Count    : 3
Oper State        : Disabled
Hello Msg Len     : 0
Hold Down Time    : 10

Statistics        :
I. Fwd. Pkts.     : 0
E. Fwd. Pkts.     : 0
I. Dro. Pkts.     : 0
E. Fwd. Octets    : 0

Associated LSP LIST :
Lsp Name          : to-49
Admin State       : Down
Time Since Last Tr*: 02h01m08s
Oper State        : Down
-----
Number of SDPs : 3
-----
Service Access Points
-----
SAP 1/2/5:0
-----
Service Id        : 9000
SAP               : 1/2/5:0
Dot1Q Ethertype   : 0x8100
Encap             : q-tag
QinQ Ethertype    : 0x8100

Admin State       : Up
Flags             : PortOperDown
Last Status Change : 04/11/2007 15:56:40
Last Mgmt Change  : 04/11/2007 17:24:54
Max Nbr of MAC Addr: No Limit
Learned MAC Addr  : 0
Admin MTU         : 1518
Ingress qos-policy : 1
Ingress Filter-Id : n/a
ARP Reply Agent   : Enabled
Mac Learning      : Enabled
Mac Aging         : Enabled
Mac Pinning       : Enabled
Total MAC Addr    : 0
Static MAC Addr   : 0
Oper MTU          : 1518
Egress qos-policy : 1
Egress Filter-Id  : n/a
Discard Unkwn Srce: Disabled

Multi Svc Site    : None
Acct. Pol         : None
Collect Stats     : Disabled
-----
Rstp Service Access Point specifics
-----
Mac Move          : Blockable
Rstp Admin State  : Up
Core Connectivity : Down
Port Role         : N/A
Port Number       : 2048
Port Path Cost    : 10
Admin Edge        : Disabled
Link Type         : Pt-pt
Designated Bridge : N/A
Active Protocol   : N/A
Rstp Oper State   : Down
Port State        : Discarding
Port Priority     : 128
Auto Edge        : Enabled
Oper Edge         : N/A
BPDU Encap       : Dot1d
Designated Port Id: 0

Forward transitions: 0
Bad BPDUs rcvd    : 0

```

Cfg BPDUs rcvd	: 0	Cfg BPDUs tx	: 0
TCN BPDUs rcvd	: 0	TCN BPDUs tx	: 0
RST BPDUs rcvd	: 0	RST BPDUs tx	: 0

-----  
Sap Statistics

	Packets	Octets
Forwarding Engine Stats		
Dropped	: 0	0
Off. HiPrio	: 0	0
Off. LowPrio	: 0	0
Off. Uncolor	: 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. InProf	: 0	0
Dro. OutProf	: 0	0
For. InProf	: 0	0
For. OutProf	: 0	0

-----  
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	0
Off. LoPrio	: 0	0
Dro. HiPrio	: 0	0
Dro. LoPrio	: 0	0
For. InProf	: 0	0
For. OutProf	: 0	0

## Egress Queue 1

For. InProf	: 0	0
For. OutProf	: 0	0
Dro. InProf	: 0	0
Dro. OutProf	: 0	0

-----  
VPLS Rapid Spanning Tree Information

VPLS oper state	: Down	Core Connectivity	: Down
Rstp Admin State	: Up	Rstp Oper State	: Down
Mode	: Rstp	Vcp Active Prot.	: N/A

Bridge Id	: 80:01:14:30:ff:00:00:01	Bridge Instance Id	: 1
Bridge Priority	: 32768	Tx Hold Count	: 6
Topology Change	: Inactive	Bridge Hello Time	: 2
Last Top. Change	: 0d 00:00:00	Bridge Max Age	: 20
Top. Change Count	: 0	Bridge Fwd Delay	: 15

```

Root Bridge      : N/A
Primary Bridge   : N/A

Root Path Cost   : 0
Rcvd Hello Time  : 2
Root Priority     : 32769

Root Forward Delay: 15
Root Max Age     : 20
Root Port        : N/A
-----
Forwarding Database specifics
-----
Service Id       : 9000
Mac Move Rate    : 10
Table Size       : 250
Learned Count    : 0
Remote Age       : 900
High WaterMark   : 95%
Mac Learning     : Enabl
Mac Aging        : Enabl

Mac Move         : Disabled
Mac Move Timeout : 10
Total Count      : 0
Static Count     : 0
Local Age        : 300
Low Watermark    : 90%
Discard Unknown  : Dsabl
=====
*A:ALA-48#

```

## arp

- Syntax** `arp [ip-address] | [mac ieee-address] | [sap sap-id] | [interface ip-int-name]`
- Context** `show>service>id`
- Description** Displays the ARP table for the IES instance.
- Parameters**
- ip-address* — Displays only ARP entries in the ARP table with the specified IP address.
  - Default** All IP addresses.
  - mac ieee-address* — Displays only ARP entries in the ARP table with the specified 48-bit MAC address. The MAC address can be expressed in the form *aa:bb:cc:dd:ee:ff* or *aa-bb-cc-dd-ee-ff* where *aa*, *bb*, *cc*, *dd*, *ee* and *ff* are hexadecimal numbers.
  - Default** All MAC addresses.
  - sap sap-id* — Displays SAP information for the specified SAP ID. See [Common Service Commands on page 2062](#) for *sap-id* command syntax.
  - ip-int-name* — The IP interface name for which to display matching ARPs.
- Output** **Show Service-ID ARP** — The following table describes show service-id ARP output fields.

Label	Description
Service ID	The value identifying the service.
MAC	The specified MAC address
Source-Identifier	The location the MAC is defined.
Type	Static — FDB entries created by management.

Label	Description (Continued)
	Learned — Dynamic entries created by the learning process.
	OAM — Entries created by the OAM process.
Age	The time lapsed since the service was enabled.
Interface	The interface applied to the service.
Port	The port where the SAP is applied.

### Sample Output

```
A:ALA-12# show service id 2 arp
=====
ARP Table
=====
IP Address      MAC Address      Type   Age      Interface      Port
-----
190.11.1.1      00:03:fa:00:08:22 Other   00:00:00 ies-100-190.11.1 1/1/11:0
=====
A:ALA-12#
```

## arp-host

**Syntax** **arp-host** [wholesaler service-id] [sap sap-id | interface interface-name | ip-address ip-address[/mask] | mac ieee-address | [[port port-id] [no-inter-dest-id | inter-dest-id inter-dest-id]]] [detail]  
**arp-host statistics** [sap sap-id | interface interface-name]  
**arp-host summary** [interface interface-name]

**Context** show>service>id

**Description** This command displays ARP host related information.

### Sample Output

```
*A:Dut-C# show service id 2 arp-host
=====
ARP host table, service 2
=====
IP Address      Mac Address      Sap Id      Remaining      MC
                  Time                                     Stdbby
-----
128.128.1.2      00:80:00:00:00:01 2/1/5:2      00h04m41s
128.128.1.3      00:80:00:00:00:02 2/1/5:2      00h04m42s
128.128.1.4      00:80:00:00:00:03 2/1/5:2      00h04m43s
128.128.1.5      00:80:00:00:00:04 2/1/5:2      00h04m44s
128.128.1.6      00:80:00:00:00:05 2/1/5:2      00h04m45s
128.128.1.7      00:80:00:00:00:06 2/1/5:2      00h04m46s
128.128.1.8      00:80:00:00:00:07 2/1/5:2      00h04m47s
128.128.1.9      00:80:00:00:00:08 2/1/5:2      00h04m48s
128.128.1.10     00:80:00:00:00:09 2/1/5:2      00h04m49s
```

## Show Commands

```
128.128.1.11      00:80:00:00:00:0a 2/1/5:2      00h04m50s
-----
Number of ARP hosts : 10
=====
*A:Dut-C#

*A:Dut-C# show service id 2 arp-host ip-address 128.128.1.2 detail
=====
ARP hosts for service 2
=====
Service ID           : 2
IP Address           : 128.128.1.2
MAC Address          : 00:80:00:00:00:01
SAP                  : 2/1/5:2
Remaining Time       : 00h04m58s

Sub-Ident            : "alu_1_2"
Sub-Profile-String   : ""
SLA-Profile-String   : ""
App-Profile-String   : ""
ARP host ANCP-String : ""
ARP host Int Dest Id : ""
RADIUS-User-Name     : "128.128.1.2"

Session Timeout (s)  : 301
Start Time           : 02/09/2009 16:35:07
Last Auth            : 02/09/2009 16:36:34
Last Refresh         : 02/09/2009 16:36:38
Persistence Key      : N/A
-----
Number of ARP hosts : 1
=====
*A:Dut-C#

*A:Dut-C# show service id 2 arp-host statistics
=====
ARP host statistics
=====
Num Active Hosts      : 20
Received Triggers     : 70
Ignored Triggers      : 10
Ignored Triggers (overload) : 0
SHCV Checks Forced    : 0
Hosts Created         : 20
Hosts Updated         : 40
Hosts Deleted         : 0
Authentication Requests Sent : 40
=====
*A:Dut-C#

*A:Dut-C# show service id 2 arp-host summary
=====
ARP host Summary, service 2
=====
Sap                  Used      Provided   Admin State
-----
sap:2/1/5:2         20       8000      inService
-----
```



```
Number of SAPs : 1
```

```
-----
=====
*A:Dut-C#
```

## base

**Context** show>service>id

This command displays basic information about the service ID including service type, description, SAPs and SDPs.

**Output** **Show Service-ID Base** — The following table describes show service-id base output fields:

Label	Description
Service Id	The service identifier.
Vpn Id	Specifies the VPN ID assigned to the service.
Service Type	The type of service.
Description	Generic information about the service.
Customer Id	The customer identifier.
Last Mgmt Change	The date and time of the most recent management-initiated change to this customer.
Adm	The administrative state of the service.
Oper	The operating state of the service.
Mtu	The largest frame size (in octets) that the service can handle.
Def. Mesh VC Id	This object is only valid in services that accept mesh SDP bindings. It is used to validate the VC ID portion of each mesh SDP binding defined in the service.
SAP Count	The number of SAPs defined on the service.
SDP Bind Count	The number of SDPs bound to the service.
Identifier	Specifies the service access (SAP) and destination (SDP) points.
Type	Specifies the signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received on the SDP.
AdmMTU	Specifies the desired largest service frame size (in octets) that can be transmitted through this SDP to the far-end ESR, without requiring the packet to be fragmented.

Label	Description (Continued)
OprMTU	Specifies the actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end ESR, without requiring the packet to be fragmented.
Opr	The operational state of the SDP.

Sample Output

```
*A:ALA-48>show>service>id# base
=====
Service Basic Information
=====
Service Id       : 750                Vpn Id           : 750
Service Type     : VPLS
Description      : Distributed VPLS services.
Customer Id      : 7
Last Status Change: 04/11/2007 09:36:33
Last Mgmt Change  : 04/11/2007 09:36:36
Admin State      : Down               Oper State       : Down
MTU              : 1514               Def. Mesh VC Id  : 750
SAP Count        : 1                 SDP Bind Count   : 2
-----
Service Access & Destination Points
-----
Identifier                Type      AdmMTU  OprMTU  Adm    Opr
-----
sap:1/1/7:0               q-tag    1518    1518    Up     Down
sdp:1:22 S(10.10.10.49)    TLDP     0       0       Up     Down
sdp:8:750 M(10.10.10.104)  TLDP     0       0       Up     Down
=====
*A:ALA-48>show>service>id#
```

fdb

- Syntax** fdb [sap sap-id] | [sdp sdp-id] | [mac ieee-address] | [detail]
- Context** show>service>id  
show>service>fdb-mac
- Description** This command displays FDB entry for a given MAC address.
- Parameters** sap sap-id — Specifies the physical port identifier portion of the SAP definition. See [Common Service Commands on page 2062](#) for sap-id command syntax.  
sdp sdp-id — The SDP identifier.
- Values** 1 — 17407
- mac ieee-address — Specifies to display information pertaining to the MAC address.
- detail — Displays detailed information.

Sample Output

```
*A:ALA-48>show>service>id# fdb mac
=====
Service Forwarding Database
=====
ServId      MAC                Source-Identifier    Type/Age  Last Change
-----
6           00:aa:00:00:00:00  sap:lag-2           L/0       04/11/2007
15:04:31
6           00:aa:00:00:00:01  sap:lag-2           L/0       04/11/2007
15:04:31
6           00:aa:00:00:00:02  sap:lag-2           L/0       04/11/2007
15:04:31
6           00:aa:00:00:00:03  sap:lag-2           L/0       04/11/2007
15:04:31
6           00:aa:00:00:00:04  sap:lag-2           L/0       04/11/2007
15:04:31
10          12:12:12:12:12:12  sap:1/1/1:100       S         04/11/2007
10:03:29
=====
*A:ALA-48>show>service>id#
```

## host

<b>Syntax</b>	<b>host [sap sap-id]</b>
<b>Context</b>	show>service>id
<b>Description</b>	This command displays static hosts configured on this service.
<b>Parameters</b>	<i>sap-id</i> — Specifies the physical port identifier portion of the SAP definition. See <a href="#">Common Service Commands on page 2062</a> for <i>sap-id</i> command syntax.

### Sample Output

```
*A:ALA-48>config>service>vpls>sap# show service id 700 host sap 1/1/9:0
=====
Static Hosts for service 700
=====
Sap          IP Address      Configured MAC      Dynamic MAC
Subscriber
-----
1/1/9:0      10.10.10.104    N/A                 N/A
N/A
-----
Number of static hosts : 1
=====
*A:ALA-48>config>service>vpls>sap#
```

## labels

<b>Syntax</b>	<b>labels</b>
<b>Context</b>	show>service>id
<b>Description</b>	Displays the labels being used by the service.
<b>Output</b>	<b>Show Service-ID Labels</b> — The following table describes show service-id labels output fields:

Label	Description
Svc Id	The service identifier.
Sdp Id	The SDP identifier.
Type	Indicates whether the SDP is a spoke or a mesh.
I. Lbl	The VC label used by the far-end device to send packets to this device in this service by the SDP.
E. Lbl	The VC label used by this device to send packets to the far-end device in this service by the SDP.

## Sample Output

```

A:ALA-12# show service id 1 labels
=====
Martini Service Labels
=====
Svc Id      Sdp Id      Type I.Lbl      E.Lbl
-----
1           10:1        Mesh 0        0
1           20:1        Mesh 0        0
1           30:1        Mesh 0        0
1           40:1        Mesh 130081    131061
1           60:1        Mesh 131019    131016
1           100:1       Mesh 0        0
-----
Number of Bound SDPs : 6
-----
A:ALA-12#

```

## mfib

<b>Syntax</b>	<b>mfib brief</b> <b>mfib [group <i>grp-address</i>]</b> <b>mfib statistics [group <i>grp-address</i>]</b>
<b>Context</b>	show>service>id
<b>Description</b>	This command displays the multicast FIB on the VPLS service.
<b>Parameters</b>	<b>group <i>grp-ip-address</i></b> — Displays the multicast FIB for a specific multicast group address. <b>statistics</b> — Displays statistics on the multicast FIB.
<b>Output</b>	<b>Show Output</b> — The following table describes the command output fields:

Label	Description
Source Address	IPv4 unicast source address.
Group Address	IPv4 multicast group address.
SAP/SDP ID	Indicates the SAP/SDP to which the corresponding multicast stream will be forwarded/blocked.
Forwarding/Blocking	Indicates whether the corresponding multicast stream will be blocked/forwarded.
Number of Entries	Number of entries in the MFIB.
Forwarded Packets	Indicates the number of multicast packets forwarded for the corresponding source/group.
Svc ID	Indicates the service to which the corresponding multicast stream will be forwarded/blocked. Local means that the multicast stream will be forwarded/blocked to a SAP or SDP local to the service.

### Sample Output

```
A:ALA-1>show>service>id # mfib
=====
Multicast FIB, Service 10
=====
Source Address  Group Address  Sap/Sdp Id          Svc Id  Fwd/Blk
-----
*               226.1.1.1      sap:1/1/1:10        Local   Fwd
                  sap:1/1/1:2         20       Fwd
                  sap:1/1/4:100        20       Fwd
                  sap:1/1/4:200        20       Fwd
*               226.1.1.2      sap:1/1/4:200        20       Fwd
-----
Number of entries: 2
=====
A:ALA-1>

*A:rbac_C# show service id 1 mfib statistics
=====
```

## Show Commands

```
Multicast FIB Statistics, Service 1
=====
Source Address  Group Address      Matched Pkts      Matched Octets
-----
*               *               838               50280
11.11.0.0       225.0.0.0           8                 480
11.11.0.0       225.0.0.1           0                 0
*               * (MLD)           0                 0
*               33:33:00:00:00:01 2650              159000
*               33:33:00:00:00:02 0                  0
-----
Number of entries: 6
=====
*A:rbae_C#
```

## mld-snooping

<b>Syntax</b>	<b>mld-snooping</b>
<b>Context</b>	show>service>id
<b>Description</b>	This command displays MLD snooping information.

## all

<b>Syntax</b>	<b>all</b>
<b>Context</b>	show>service>id>mld-snooping
<b>Description</b>	This command displays detailed information about MLD snooping.

**Sample Output**

```
*A:rbae_C# show service id 1 mld-snooping all
=====
MLD Snooping info for service 1
-----
MLD Snooping Base info
-----
Admin State : Up
Querier      : FE80::12 on SAP 2/1/5
-----
Sap/Sdp      Oper   MRtr  Send   Max Num  MVR      Num
Id           State  Port  Queries Groups  From-VPLS Groups
-----
sap:1/1/4    Up     No    Disabled No Limit Local     0
sap:2/1/5    Up     Yes   Disabled No Limit Local     0
sdp:31:1     Up     No    Disabled No Limit N/A      0
sdp:36:1     Up     No    Disabled No Limit N/A      0
-----
MLD Snooping Querier info
-----
Sap Id       : 2/1/5
IP Address   : FE80::12
Expires      : 11s
Up Time      : 0d 00:05:05
Version      : 2

General Query Interval : 10s
Query Response Interval : 1.0s
Robust Count           : 2
-----
MLD Snooping Multicast Routers
-----
MRouter
Sap/Sdp Id           Up Time      Expires      Version
-----
FE80::12
2/1/5                0d 00:05:05  11s          2
-----
Number of mrouter: 1
```

## Show Commands

```

-----
MLD Snooping Proxy-reporting DB
-----
Group Address          Mode      Up Time      Num Sources
-----
Number of groups: 0
-----
MLD Snooping SAP 1/1/4 Port-DB
-----
Group Address          Mode      Type      From-VPLS  Up Time      Expires  Num  MC
                               Src      Stdby
-----
Number of groups: 0
-----
MLD Snooping SAP 2/1/5 Port-DB
-----
Group Address          Mode      Type      From-VPLS  Up Time      Expires  Num  MC
                               Src      Stdby
-----
Number of groups: 0
-----
MLD Snooping SDP 31:1 Port-DB
-----
Group Address          Mode      Type      From-VPLS  Up Time      Expires  Num Src
-----
Number of groups: 0
-----
MLD Snooping SDP 36:1 Port-DB
-----
Group Address          Mode      Type      From-VPLS  Up Time      Expires  Num Src
-----
Number of groups: 0
-----
MLD Snooping Static Source Groups
-----
MLD Snooping Statistics
-----
Message Type          Received      Transmitted  Forwarded
-----
General Queries       43           0           129
Group Queries         0           0           0
Group-Source Queries  0           0           0
V1 Reports            0           0           0
V2 Reports            4           35          0
V1 Done               0           0           0
Unknown Type          0           N/A         0
-----
Drop Statistics
-----
Bad Length           : 0
Bad MLD Checksum     : 0
Bad Encoding          : 0
No Router Alert      : 0
Zero Source IP       : 0
Wrong Version         : 0
Lcl-Scope Packets    : 0

```



```

Rsvd-Scope Packets      : 0

Send Query Cfg Drops    : 0
Import Policy Drops     : 0
Exceeded Max Num Groups : 0
MCAC Policy Drops       : 0
MCS Failures            : 0

MVR From VPLS Cfg Drops : 0
MVR To SAP Cfg Drops    : 0
-----
MLD Snooping Multicast VPLS Registration info
-----
MLD Snooping Admin State : Up

MVR Admin State          : Down
MVR Policy               : None
-----
Local SAPs/SDPs
-----
Svc Id      Sap/Sdp      Oper      From      Num Local
            Id           State     VPLS      Groups
-----
1           sap:1/1/4     Up        Local      0
1           sap:2/1/5     Up        Local      0
1           sdp:31:1      Up        N/A        0
1           sdp:36:1      Up        N/A        0
-----
MVR SAPs (from-vpls=1)
-----
Svc Id      Sap/Sdp      Oper      From      Num MVR
            Id           State     VPLS      Groups
-----
No MVR SAPs found.
=====
*A:rbac_C#

```

## base

<b>Syntax</b>	<b>base</b>
<b>Context</b>	show>service>id>mld-snooping
<b>Description</b>	This command displays basic MLD snooping information.

## mrouters

<b>Syntax</b>	<b>mrouters [detail]</b>
<b>Context</b>	show>service>id>mld-snooping
<b>Description</b>	This command displays all multicast routers.

**Sample Output**

```

*A:rbae_C# show service id 1 mld-snooping mrouter
=====
MLD Snooping Multicast Routers for service 1
=====
MRouter
      Sap/Sdp Id          Up Time      Expires    Version
-----
FE80::12
      2/1/5              0d 00:05:55    11s        2
-----
Number of mrouter: 1
=====
*A:rbae_C#

*A:rbae_C# show service id 1 mld-snooping mrouter detail
=====
MLD Snooping Multicast Routers for service 1
-----
MRouter FE80::12
-----
Sap Id      : 2/1/5
Expires     : 18s
Up Time     : 0d 00:06:28
Version     : 2
-----
Number of mrouter: 1
=====

```

**mvr**

<b>Syntax</b>	<b>mvr</b>
<b>Context</b>	show>service>id>mld-snooping
<b>Description</b>	This command displays multicast VPLS registration information.

**port-db**

<b>Syntax</b>	<b>port-db sap <i>sap-id</i></b> <b>port-db sap <i>sap-id</i> detail</b> <b>port-db sap <i>sap-id</i> group <i>grp-ipv6-address</i></b> <b>port-db sdp <i>sdp-id:vc-id</i> [detail]</b> <b>port-db sdp <i>sdp-id:vc-id</i> group <i>grp-ipv6-address</i></b>
<b>Context</b>	show>service>id>mld-snooping
<b>Description</b>	This command displays MLD snooping information related to a specific SAP.

**Sample Output**

```

*A:rbae_C# show service id 1 mld-snooping port-db sap 1/1/4
=====
MLD Snooping SAP 1/1/4 Port-DB for service 1
=====
Group Address
      Mode      Type      From-VPLS  Up Time      Expires  Num  MC
                               Src      Stdby
-----
FF04::1
      include dynamic local      0d 00:00:19  0s      1
FF04::2
      include dynamic local      0d 00:00:18  0s      1
-----
Number of groups: 2
=====
*A:rbae_C#

*A:rbae_C# show service id 1 mld-snooping port-db sap 1/1/4 detail
=====
MLD Snooping SAP 1/1/4 Port-DB for service 1
-----
MLD Group FF04::1
-----
Mode           : include           Type           : dynamic
Up Time        : 0d 00:00:33       Expires        : 0s
Compat Mode    : MLD Version 2
Vl Host Expires : 0s
MVR From-VPLS  : local            MVR To-SAP     : local
MC Standby     : no
-----
Source Address
      Up Time      Expires  Type      Fwd/Blk
-----
2011::1
      0d 00:00:33  20s      dynamic  Fwd
-----
MLD Group FF04::2
-----
Mode           : include           Type           : dynamic
Up Time        : 0d 00:00:32       Expires        : 0s
Compat Mode    : MLD Version 2
Vl Host Expires : 0s
MVR From-VPLS  : local            MVR To-SAP     : local
MC Standby     : no
-----
Source Address
      Up Time      Expires  Type      Fwd/Blk
-----
2011::1
      0d 00:00:33  20s      dynamic  Fwd
-----
Number of groups: 2
=====
*A:rbae_C#

```

proxy-db

<b>Syntax</b>	<b>proxy-db [detail]</b> <b>proxy-db group</b> <i>grp-ipv6-address</i>
<b>Context</b>	show>service>id>mld-snooping
<b>Description</b>	This command displays proxy-reporting database entries.

Sample Output

```
*A:rbae_C# show service id 1 mld-snooping proxy-db
=====
MLD Snooping Proxy-reporting DB for service 1
=====
Group Address
Mode          Up Time      Num Sources
-----
FF04::1
include       0d 00:01:01    1
FF04::2
include       0d 00:01:00    1
-----
Number of groups: 2
=====
*A:rbae_C#

*A:rbae_C# show service id 1 mld-snooping proxy-db detail
=====
MLD Snooping Proxy-reporting DB for service 1
-----
MLD Group FF04::1
-----
Up Time : 0d 00:01:03          Mode : include
-----
Source Address                Up Time
-----
2011::1                      0d 00:01:03
-----
MLD Group FF04::2
-----
Up Time : 0d 00:01:02          Mode : include
-----
Source Address                Up Time
-----
2011::1                      0d 00:01:02
-----
Number of groups: 2
=====
*A:rbae_C#
```

## querier

<b>Syntax</b>	<b>querier</b>
<b>Context</b>	show>service>id>mld-snooping
<b>Description</b>	This command displays information about the current querier.

**Sample Output**

```
*A:rbac_C# show service id 1 mld-snooping querier
=====
MLD Snooping Querier info for service 1
=====
Sap Id           : 2/1/5
IP Address       : FE80::12
Expires          : 11s
Up Time          : 0d 00:13:35
Version          : 2

General Query Interval : 10s
Query Response Interval : 1.0s
Robust Count          : 2
=====
*A:rbac_C#
```

## static

<b>Syntax</b>	<b>static [sap sap-id   sdp sdp-id:vc-id]</b>
<b>Context</b>	show>service>id>mld-snooping
<b>Description</b>	This command displays MLD snooping static group membership data.

**Sample Output**

```
*A:rbac_C# show service id 1 mld-snooping static
=====
MLD Snooping Static Source Groups for service 1
-----
MLD Snooping Static Source Groups for SDP 36:1
-----
Source           Group
-----
2011::1          FF04::2
*
                  FF04::3
-----
Static (*,G)/(S,G) entries: 2
=====
*A:rbac_C#
```

## statistics

<b>Syntax</b>	<b>statistics[sap sap-id   sdp sdp-id:vc-id]</b>
<b>Context</b>	show>service>id>mld-snooping
<b>Description</b>	This command displays MLD snooping statistics.

## Sample Output

```
*A:rbae_C# show service id 1 mld-snooping statistics
=====
MLD Snooping Statistics for service 1
=====
Message Type           Received      Transmitted   Forwarded
-----
General Queries        109           0             327
Group Queries           0             0             0
Group-Source Queries    0             0             0
V1 Reports              0             0             0
V2 Reports              438           87            0
V1 Done                 0             0             0
Unknown Type            0             N/A           0
-----
Drop Statistics
-----
Bad Length              : 0
Bad MLD Checksum        : 0
Bad Encoding             : 0
No Router Alert         : 0
Zero Source IP          : 0
Wrong Version           : 0
Lcl-Scope Packets       : 0
Rsvd-Scope Packets      : 0

Send Query Cfg Drops    : 0
Import Policy Drops     : 0
Exceeded Max Num Groups : 0
MCAC Policy Drops       : 0
MCS Failures            : 0

MVR From VPLS Cfg Drops : 0
MVR To SAP Cfg Drops    : 0
=====
*A:rbae_C#
```

## mstp-configuration

<b>Syntax</b>	<b>mstp-configuration</b>
<b>Context</b>	show>service>id
<b>Description</b>	This command displays the MSTP specific configuration data. This command is only valid on a managment VPLS.

## sap

<b>Syntax</b>	<b>sap sap-id [detail]</b>
<b>Context</b>	show>service>id
<b>Description</b>	This command displays information for the SAPs associated with the service. If no optional parameters are specified, a summary of all associated SAPs is displayed.
<b>Parameters</b>	<i>sap-id</i> — The ID that displays SAP information. See <a href="#">Common Service Commands on page 2062</a> for <i>sap-id</i> command syntax. <b>detail</b> — Displays detailed information for the SAP.
<b>Output</b>	<b>Show Service-ID SAP</b> — The following table describes show service SAP fields:

Label	Description
Service Id	The service identifier.
SAP	The SAP and qtag.
Encap	The encapsulation type of the SAP.
Ethertype	Specifies an Ethernet type II Ethertype value.
Admin State	The administrative state of the SAP.
Oper State	The operational state of the SAP.
Flags	Specifies the conditions that affect the operating status of this SAP.
Last Status Change	Specifies the time of the most recent operating status change to this SAP
Last Mgmt Change	Specifies the time of the most recent management-initiated change to this SAP.
Admin MTU	The desired largest service frame size (in octets) that can be transmitted through the SAP to the far-end router, without requiring the packet to be fragmented.

Label	Description (Continued)
Oper MTU	The actual largest service frame size (in octets) that can be transmitted through the SAP to the far-end router, without requiring the packet to be fragmented.
Ingress qos-policy	The ingress QoS policy ID assigned to the SAP.
Egress qos-policy	The egress QoS policy ID assigned to the SAP.
Ingress Filter-Id	The ingress filter policy ID assigned to the SAP.
Egress Filter-Id	The egress filter policy ID assigned to the SAP.
Acct. Pol	The accounting policy ID assigned to the SAP.
Collect Stats	Specifies whether collect stats is enabled.
Dropped	The number of packets and octets dropped due to SAP state, ingress MAC or IP filter, same segment discard, bad checksum, etc.
Off. HiPrio	The number of high priority packets and octets, as determined by the SAP ingress QoS policy, offered by the Pchip to the Qchip.
Off. LowPrio	The number of low priority packets and octets, as determined by the SAP ingress QoS policy, offered by the Pchip to the Qchip.
Off. Uncolor	The number of uncolored packets and octets, as determined by the SAP ingress QoS policy, offered by the Pchip to the Qchip.
Dro. HiPrio	The number of high priority packets and octets, as determined by the SAP ingress QoS policy, dropped by the Qchip due to: MBS exceeded, buffer pool limit exceeded, etc.
Dro. LowPrio	The number of low priority packets and octets, as determined by the SAP ingress QoS policy, dropped by the Qchip due to: MBS exceeded, buffer pool limit exceeded, etc.
For. InProf	The number of in-profile packets and octets (rate below CIR) forwarded by the ingress Qchip.
For. OutProf	The number of out-of-profile packets and octets discarded by the egress Qchip due to MBS exceeded, buffer pool limit exceeded, etc.
Dro. InProf	The number of in-profile packets and octets discarded by the egress Qchip due to MBS exceeded, buffer pool limit exceeded, etc.
Dro. OutProf	The number of out-of-profile packets and octets discarded by the egress Qchip due to MBS exceeded, buffer pool limit exceeded, etc.
For. InProf	The number of in-profile packets and octets (rate below CIR) forwarded by the egress Qchip.
For. OutProf	The number of out-of-profile packets and octets (rate above CIR) forwarded by the egress Qchip.



```

Max Nbr of MAC Addr: No Limit
Learned MAC Addr : 0
Admin MTU : 1518
Ingress qos-policy : 1
Ingress Filter-Id : n/a
Mac Learning : Enabled
Mac Aging : Enabled

Total MAC Addr : 0
Static MAC Addr : 0
Oper MTU : 1518
Egress qos-policy : 1
Egress Filter-Id : n/a
Discard Unkwn Srce: Disabled

Multi Svc Site : West
E. Sched Pol : SLA1
Acct. Pol : None
Collect Stats : Disabled
=====
*A:ALA-48>show>service>id#

*A:ALA-48>show>service>id# sap 1/1/7:0 detail
=====
Service Access Points(SAP)
=====
Service Id : 750
SAP : 1/1/7:0
Dot1Q Ethertype : 0x8100
Encap : q-tag
QinQ Ethertype : 0x8100

Admin State : Up
Flags : PortOperDown
Last Status Change : 04/09/2007 09:23:26
Last Mgmt Change : 04/09/2007 09:23:28
Max Nbr of MAC Addr: No Limit
Learned MAC Addr : 0
Admin MTU : 1518
Ingress qos-policy : 100
Shared Q plcy : default
Ingr IP Fltr-Id : n/a
Ingr Mac Fltr-Id : n/a
Ingr IPv6 Fltr-Id : n/a
tod-suite : None
ARP Reply Agent : Enabled
Mac Learning : Enabled
Mac Aging : Enabled
L2PT Termination : Disabled

Total MAC Addr : 0
Static MAC Addr : 0
Oper MTU : 1518
Egress qos-policy : 1
Multipoint shared : Enabled
Egr IP Fltr-Id : 10
Egr Mac Fltr-Id : n/a
Egr IPv6 Fltr-Id : n/a
qinq-pbit-marking : both
Host Conn Verify : Enabled
Discard Unkwn Srce: Disabled
Mac Pinning : Disabled
BPDU Translation : Disabled

Multi Svc Site : None
I. Sched Pol : SLA1
E. Sched Pol : SLA1
Acct. Pol : None
Collect Stats : Disabled

Anti Spoofing : None
MCAC Policy Name :
MCAC Max Unconst BW: no limit
MCAC In use Mand BW: 0
MCAC In use Opnl BW: 0
Auth Policy : none
Egr MCast Grp :

Nbr Static Hosts : 1
MCAC Const Adm St : Enable
MCAC Max Mand BW : no limit
MCAC Avail Mand BW: unlimited
MCAC Avail Opnl BW: unlimited

-----
Rstp Service Access Point specifics
-----
Mac Move : Blockable
Rstp Admin State : Up
Core Connectivity : Down
Port Role : N/A
Port Number : 2048

Rstp Oper State : Down
Port State : Discarding
Port Priority : 128

```

## Show Commands

```

Port Path Cost      : 10
Admin Edge          : Disabled
Link Type           : Pt-pt
Designated Bridge   : N/A
Active Protocol     : N/A
Auto Edge           : Enabled
Oper Edge           : N/A
BPDU Encap          : Dot1d
Designated Port Id: 0

```

```

Forward transitions: 0
Cfg BPDUs rcvd     : 0
TCN BPDUs rcvd     : 0
RST BPDUs rcvd     : 0
Bad BPDUs rcvd     : 0
Cfg BPDUs tx       : 0
TCN BPDUs tx       : 0
RST BPDUs tx       : 0

```

### Sap Statistics

```

-----
Packets      Octets
Forwarding Engine Stats
Dropped      : 0      0
Off. HiPrio  : 0      0
Off. LowPrio : 0      0
Off. Uncolor : 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. InProf  : 0      0
Dro. OutProf : 0      0
For. InProf  : 0      0
For. OutProf : 0      0

```

### 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      0
Off. LoPrio  : 0      0
Dro. HiPrio  : 0      0
Dro. LoPrio  : 0      0
For. InProf  : 0      0
For. OutProf : 0      0

```

### Egress Queue 1

```

For. InProf  : 0      0
For. OutProf : 0      0
Dro. InProf  : 0      0
Dro. OutProf : 0      0

```

```

=====
*A:ALA-48>show>service>id#

```

## sdp

- Syntax** **sdp** [*sdp-id* | **far-end** *ip-addr*] [**detail**]
- Context** show>service>id
- Description** Displays information for the SDPs associated with the service.  
If no optional parameters are specified, a summary of all associated SDPs is displayed.
- Parameters** *sdp-id* — Displays only information for the specified SDP ID.  
     **Default** All SDPs.  
     **Values** 1 — 17407  
*far-end ip-addr* — Displays only SDPs matching with the specified far-end IP address.  
     **Default** SDPs with any far-end IP address.  
**detail** — Displays detailed SDP information.
- Output** **Show Service-ID SDP** — The following table describes show service-id SDP output fields:

Label	Description
Sdp Id	The SDP identifier.
Type	Indicates whether the SDP is a spoke or a mesh.
Split Horizon Group	Name of the split horizon group that the SAP belongs to.
VC Type	Displays the VC type: ether, vlan, or vpls.
VC Tag	Displays the explicit dot1Q value used when encapsulating to the SDP far end.
I. Lbl	The VC label used by the far-end device to send packets to this device in this service by the SDP.
Admin Path MTU	The operating path MTU of the SDP is equal to the admin path MTU (when one is set) or the dynamically computed tunnel MTU, when no admin path MTU is set (the default case.)
Oper Path MTU	The actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Far End	Specifies the IP address of the remote end of the GRE or MPLS tunnel defined by this SDP.
Delivery	Specifies the type of delivery used by the SDP: GRE or MPLS.
Admin State	The administrative state of this SDP.
Oper State	The current status of the SDP.

Label	Description (Continued)
Ingress Label	The label used by the far-end device to send packets to this device in this service by this SDP.

### Sample Output

```

A:ALA-12# show service id 9000 sdp 2:22 detail
=====
Service Destination Point (Sdp Id : 2:22) Details
-----
  Sdp Id 2:22  -(10.10.10.103)
-----
Description      : GRE-10.10.10.103
SDP Id           : 2:22                               Type           : Spoke
Split Horiz Grp  : (DSL-group1)
VC Type          : Ether                               VC Tag          : n/a
Admin Path MTU   : 4462                               Oper Path MTU   : 4462
Far End          : 10.10.10.103                       Delivery        : GRE
Admin State      : Up                                 Oper State      : TLDP Down
Ingress Label    : 0                                  Egress Label    : 0
Ingress Filter   : n/a                               Egress Filter   : n/a
Last Changed     : 04/11/2007 11:48:20                Signaling       : TLDP

KeepAlive Information :
Admin State          : Disabled                       Oper State       : Disabled
Hello Time           : 10                             Hello Msg Len    : 0
Max Drop Count       : 3                             Hold Down Time   : 10

Statistics           :
I. Fwd. Pkts.        : 0                             I. Dro. Pkts.    : 0
E. Fwd. Pkts.        : 0                             E. Fwd. Octets   : 0

Associated LSP LIST :
SDP Delivery Mechanism is not MPLS
-----
Rstp Service Destination Point specifics
-----
Mac Move            : Disabled
Rstp Admin State    : Up                               Rstp Oper State   : Down
Core Connectivity   : Down
Port Role           : N/A                             Port State        : Discarding
Port Number         : 2049                             Port Priority     : 128
Port Path Cost      : 10                               Auto Edge        : Enabled
Admin Edge          : Disabled                         Oper Edge         : N/A
Link Type           : Pt-pt                            BPDU Encap       : Dot1d
Designated Bridge   : N/A                             Designated Port Id: 0
Active Protocol     : N/A

Fwd Transitions     : 0                               Bad BPDUs rcvd   : 0
Cfg BPDUs rcvd      : 0                               Cfg BPDUs tx     : 0
TCN BPDUs rcvd      : 0                               TCN BPDUs tx     : 0
RST BPDUs rcvd      : 0                               RST BPDUs tx     : 0
-----
Number of SDPs : 1
=====
A:ALA-12#

```

## gsmp

<b>Syntax</b>	<b>gsmp</b>
<b>Context</b>	show>service>id
<b>Description</b>	This enables the command to display GSMP information.

## neighbors

<b>Syntax</b>	<b>neighbors group</b> [ <i>name</i> ] [ <i>ip-address</i> ]
<b>Context</b>	show>service>id>gsmp
<b>Description</b>	This command display GSMP neighbor information.
<b>Parameters</b>	<p><b>group</b> — A GSMP group defines a set of GSMP neighbors which have the same properties.</p> <p><i>name</i> — Specifies a GSMP group name is unique only within the scope of the service in which it is defined.</p> <p><i>ip-address</i> — Specifies the ip-address of the neighbor.</p>
<b>Output</b>	<b>Show Service-ID GSMP Neighbors Group</b> — The following table describes show service-id gsmp neighbors group output fields:

Label	Description
Group	Displays the group name.
Neighbor	Displays the neighbor IP address.
AdminState	Displays the administrative state of the neighbor.
Sessions	Displays the number of sessions (open TCP connections) for each configured neighbor.

## Sample Output

These commands show the configured neighbors per service, regardless that there exists an open TCP connection with this neighbor. The admin state is shown because for a neighbor to be admin enabled, the service, gsmp node, group node and the neighbor node in this service must all be in 'no shutdown' state. Session gives the number of session (open TCP connections) for each configured neighbor.

```
A:active>show>service>id>gsmp# neighbors
=====
GSMP neighbors
=====
Group                Neighbor                AdminState  Sessions
-----
dslam1                192.168.1.2            Enabled     0
dslam1                192.168.1.3            Enabled     0
```

```
-----
Number of neighbors shown: 2
=====
A:active>show>service>id>gsmp#

A:active>show>service>id>gsmp# neighbors group dslam1
=====
GSMP neighbors
=====
Group                               Neighbor                AdminState  Sessions
-----
dslam1                             192.168.1.2            Enabled     0
dslam1                             192.168.1.3            Enabled     0
-----
Number of neighbors shown: 2
=====
A:active>show>service>id>gsmp#

A:active>show>service>id>gsmp# neighbors group dslam1 192.168.1.2
=====
GSMP neighbors
=====
Group                               Neighbor                AdminState  Sessions
-----
dslam1                             192.168.1.2            Enabled     0
=====
A:active>show>service>id>gsmp#
```

sessions

Syntax	<b>sessions</b> [ <i>group name</i> ] <b>neighbor</b> <i>ip-address</i> ] [ <i>port port-number</i> ] [ <i>association</i> ] [ <i>statistics</i> ]
Context	show>service>id>gsmp
Description	This command displays GSMP sessions information.
Parameters	<b>group</b> — A GSMP group defines a set of GSMP neighbors which have the same properties. <i>name</i> — Specifies a GSMP group name is unique only within the scope of the service in which it is defined. <i>ip-address</i> — Specifies the ip-address of the neighbor. <i>port</i> — Specifies the neighbor TCP port number use for this ANCP session. <b>Values</b> 0 — 65535 <b>association</b> — Displays to what object the ANCP-string is associated. <b>statistics</b> — Displays statistics information about an ANCP session known to the system.
Output	<b>Show Service-ID GSMP sessions</b> — The following table describes service ID GSMP sessions output fields:

Label	Description
Port	Displays the port ID number.
Ngbr-IPAddr	Displays the neighbor IP address.
GsmP-Group	Displays the GSMP group ID.
State	The GSMP state of this TCP connection.
Peer Instance	Together with the peer port and peer name output, displays a unique GSMP ID for each end of the GSMP connection.
Peer Port	Together with the peer instance and peer name output, displays a unique GSMP ID for each end of the GSMP connection.
Peer Name	Together with the peer port and peer instance output, displays a unique GSMP ID for each end of the GSMP connection.
timeouts	Displays the number of successive timeouts for this session.
Peer Timer	Displays the GSMP keepalive timer.
Capabilities	Displays the ANCP capabilities negotiated for this session.
Conf Capabilities	Displays the ANCP capabilities configured for this session.
Priority Marking	Displays the priority marking configured for this session.
Local Addr	Displays the IP address used by the box's side of the TCP connection.
Conf. Local Addr.	Displays the configured IP address used by the box's side of the TCP connection.
Sender Instance	The instance sent to the neighbor in this session.
Sender Port	The port sent to the neighbor in this session.
Sender Name	The name sent to the neighbor in this session.
Max. Timeouts	The maximum number of successive timeouts configured for this session.
Sender Timer	Indicates the timeout value that will be announced towards the neighbor. The neighbor uses this timeout value while waiting for an acknowledgment from this system.

### Sample Output

This show command gives information about the open TCP connections with DSLAMs.

```
A:active>show>service>id>gsmP# sessions
=====
GSMP sessions for service 999
=====
Port    Ngbr-IPAddr    GsmP-Group
```

## Show Commands

```
-----
40590 192.168.1.2 dslam1
-----
Number of GSMP sessions : 1
=====
A:active>show>service>id>gsmp#

A:active>show>service>id>gsmp# sessions neighbor 192.168.1.2 port 40590
=====
GSMP sessions for service 999 , neighbor 192.168.1.2, Port 40590
=====
State           : Established
Peer Instance   : 1                      Sender Instance : a3cf58
Peer Port       : 0                      Sender Port      : 0
Peer Name       : 12:12:12:12:12:12     Sender Name      : 00:00:00:00:00:00
Timeouts        : 0                      Max. Timeouts    : 3
Peer Timer      : 100                    Sender Timer     : 100
Capabilities     : DTD OAM
Conf Capabilities : DTD OAM
Priority Marking  : dscp nc2
Local Addr.     : 192.168.1.4
Conf Local Addr. : N/A
=====
A:active>show>service>id>gsmp#
A:active>show>service>id>gsmp# sessions neighbor 192.168.1.2 port 40590 association
=====
ANCP-Strings
=====
ANCP-String                                     Assoc. State
-----
No ANCP-Strings found
=====
A:active>show>service>id>gsmp#
A:active>show>service>id>gsmp# sessions neighbor 192.168.1.2 port 40590 statistics
=====
GSMP session stats, service 999, neighbor 192.168.1.2, Port 40590
=====
Event                                     Received Transmitted
-----
Dropped                                0         0
Syn                                    1         1
Syn Ack                                1         1
Ack                                    14        14
Rst Ack                                0         0
Port Up                                0         0
Port Down                              0         0
OAM Loopback                           0         0
=====
A:active>show>service>id>gsmp#
```

Note: The association command gives an overview of each ANCP string received from this session.

```
A:active>show>service>id>gsmp# sessions neighbor 192.168.1.2 port 40590 association
=====
ANCP-Strings
=====
ANCP-String                                     Assoc.
State
-----
7330-ISAM-E47 atm 1/1/01/01:19425.64048          ANCP    Up
```



```
-----
Number of ANCP-Strings : 1
=====
A:active>show>service>id>gsm
```

## host

- Syntax** **host** [**sap** *sap-id*] [**wholesaler** *service-id*] [**port** *port-id*] [**inter-dest-id** *intermediate-destination-id*] [**detail**]  
**host** [**sap** *sap-id*] [**wholesaler** *service-id*] [**port** *port-id*] **no-inter-dest-id** [**detail**]  
**host summary**
- Context** show>service>id
- Description** This command displays information about static host configured on this service.
- Parameters** **sap** *sap-id* — Specifies the physical port identifier portion of the SAP definition. See [Common Service Commands on page 2062](#) for *sap-id* command syntax.
- intermediate-destination-id* — Specifies the intermediate destination identifier which is encoded in the identification strings.
- Values** Up to 32 characters maximum
- summary** — Displays summary static host information.
- detail** — Displays detailed static host information.

## interface

- Syntax** **interface** [{*ip-address*|*ip-int-name*] [*interface-type*] [**detail**] [**family** *family*]|**summary**]
- Context** show>service>id
- Description** This command displays service interface information.
- Parameters** *ip-address* — Displays information for the specified IP address.
- ip-int-name* — Displays information for the specified interface name.
- detail** — Displays detailed interface information.
- family** *family* — Specifies to display only information belonging to the address family IPv4 or IPv6. Only these two values will be accepted.
- summary** — Displays detailed interface information.

### Sample Output

```
A:cses-E11>config>service>vprn# show service id 10 interface "gi-2-01" detail
=====
Interface Table
=====
Interface
```

## Show Commands

```
-----
If Name       : gi-2-01
Sub If Name   : si-2
Red If Name   :
Admin State   : Up                               Oper (v4/v6) : Up/Down
Protocols     : None
-----

Details
-----
Description   : (Not Specified)
If Index      : 6                               Virt. If Index : 6
Last Oper Chg: 11/27/2012 13:19:28             Global If Index : 380
Mon Oper Grp  : None
Srrp En RtnG : Disabled                         Hold time      : N/A
Group Port    : 1/1/2
TOS Marking    : Trusted                       If Type        : VPRN Grp
SNTP B.Cast    : False
MAC Address    : d2:30:01:01:00:02             Mac Accounting  : Disabled
Ingress stats : Disabled
Arp Timeout    : 14400                         IPv6 Nbr ReachTi*: 30
IP Oper MTU     : 1500                         ICMP Mask Reply : True
Arp Populate    : Disabled                     Host Conn Verify : Disabled
Cflowd         : None
LdpSyncTimer   : None
LSR Load Bal*  : system
uRPF Chk       : disabled
uRPF Ipv6 Chk  : disabled
Rx Pkts        : 0                             Rx Bytes       : 0
Rx V4 Pkts     : 0                             Rx V4 Bytes    : 0
Rx V6 Pkts     : 0                             Rx V6 Bytes    : 0
Tx Pkts        : 32                             Tx Bytes       : 3392
Tx V4 Pkts     : 32                             Tx V4 Bytes    : 3392
Tx V4 Discar*  : 0                             Tx V4 Discard By*: 0
Tx V6 Pkts     : 0                             Tx V6 Bytes    : 0
Tx V6 Discar*  : 0                             Tx V6 Discard By*: 0

Proxy ARP Details
Rem Proxy ARP  : Disabled                       Local Proxy ARP : Disabled
Policies       : none

Proxy Neighbor Discovery Details
Local Pxy ND   : Disabled
Policies       : none

DHCP no local server

DHCP Details
Description    : (Not Specified)
Admin State    : Down                           Lease Populate  : 0
Gi-Addr       : Not configured                  Gi-Addr as Src Ip: Disabled
Action        : Keep                           Trusted        : Disabled

DHCP Proxy Details
Admin State    : Down
Lease Time     : N/A
Emul. Server   : Not configured

Subscriber Authentication Details
Auth Policy    : None

DHCP6 Relay Details
```

```

Description   : (Not Specified)
Admin State   : Down
Oper State    : Down
If-Id Option  : None
Src Addr      : Not configured
Lease Populate : 0
Nbr Resolution : Disabled
Remote Id     : Disabled

DHCP6 Server Details
Admin State   : Down
Max. Lease States: 8000

ISA Tunnel redundant next-hop information
Static Next-*:
Dynamic Next*:

ICMP Details
Redirects      : Number - 100
Unreachables   : Number - 100
TTL Expired    : Number - 100
Time (seconds) : 10
Time (seconds) : 10
Time (seconds) : 10

IPCP Address Extension Details
Peer IP Addr*  : Not configured
Peer Pri DNS*  : Not configured
Peer Sec DNS*  : Not configured
-----
Qos Details
-----
Ing Qos Poli*  : (none)
Ingress FP Q*  : (none)
Ing FP QGrp *  : (none)
Egr Qos Policy : (none)
Egress Port QGrp : (none)
Egr Port QGrp In* : (none)
-----
Interfaces : 1
=====
* indicates that the corresponding row element may have been truncated.
A:cses-E11>config>service>vprn#

```

## retailers

**Syntax**     **retailers**

**Context**    show>service>id

**Description** This command displays service retailer information.

### Sample Output

```

*A:ALA-48>config# show service id 101 retailers
=====
Retailers for service 101
=====
Retailer Svc ID          Num Static Hosts      Num Dynamic Hosts
-----
102                      3                      1
105                      0                      1
-----
Number of retailers : 2
=====

```

```
*A:ALA-48>config#
```

## wholesalers

<b>Syntax</b>	<b>wholesalers</b>
<b>Context</b>	show>service>id
<b>Description</b>	This command displays service wholesaler information.

### Sample Output

```
*A:ALA-48>config# show service id 102 wholesalers
=====
Wholesalers for service 102
=====
Wholesaler Svc ID          Num Static Hosts      Num Dynamic Hosts
-----
101                        3                      1
-----
Number of wholesalers : 1
=====
*A:ALA-48>config#
```

Wholesaler information can also be displayed in the lease-state context.

```
*A:ALA-48>config# show service id 105 dhcp lease-state wholesaler 101
=====
DHCP lease state table, service 105
=====
IP Address      Mac Address      Sap/Sdp Id      Remaining Lease MC
                  LifeTime      Origin      Stdbby
-----
Wholesaler 101 Leases
-----
103.3.2.62      00:00:1f:bd:00:c6 lag-1:105      00h00m39s  RADIUS
-----
Number of lease states : 1
=====
*A:ALA-48>config#
```

## split-horizon-group

<b>Syntax</b>	<b>split-horizon-group</b> [ <i>group-name</i> ]
<b>Context</b>	show>service>id
<b>Description</b>	This command displays service split horizon groups.
<b>Parameters</b>	<i>group-name</i> — Specifies a group name up to 32 characters in length.

**Sample Output**

```

A:ALA-1# show service id 700 split-horizon-group
=====
Service: Split Horizon Group
=====
Name                               Description
-----
R DSL-group1                       Split horizon group for DSL
-----
R = Residential Split Horizon Group
No. of Split Horizon Groups: 1
=====
A:ALA-1#

A:ALA-1# show service id 700 split-horizon-group DSL-group1
=====
Service: Split Horizon Group
=====
Name                               Description
-----
R DSL-group1                       Split horizon group for DSL
-----
Associations
-----
SAP                                1/1/3:1

SDP                                108:1
SDP                                109:1
-----
R = Residential Split Horizon Group
SAPs Associated : 1                SDPs Associated : 2
=====
A:ALA-1#

```

static-host

<b>Syntax</b>	<b>static-host</b> [ <b>sap</b> <i>sap-id</i> ] [ <b>wholesaler</b> <i>service-id</i> ] [ <b>port</b> <i>port-id</i> ][ <b>inter-dest-id</b> <i>intermediate-destination-id</i> ] [ <b>detail</b> ] <b>static-host</b> [ <b>sap</b> <i>sap-id</i> ] [ <b>wholesaler</b> <i>service-id</i> ] [ <b>port</b> <i>port-id</i> ] <b>no-inter-dest-id</b> [ <b>detail</b> ] <b>static-host summary</b>
<b>Context</b>	show>service>id
<b>Description</b>	This command displays Display static hosts configured on this service.
<b>Parameters</b>	<b>sap</b> <i>sap-id</i> — Specifies the physical port identifier portion of the SAP definition. See <a href="#">Common Service Commands on page 2062</a> for <i>sap-id</i> command syntax.  <i>intermediate-destination-id</i> — Specifies the intermediate destination identifier which is encoded in the identification strings.  <b>Values</b> Up to 32 characters maximum  <b>summary</b> — Displays summary static host information.  <b>detail</b> — Displays detailed static host information.

Sample Output

```
*A:ALA-48# show service id 88 static-host
=====
Static Hosts for service 88
=====
Sap              IP Address      Configured MAC   Dynamic MAC
Subscriber                               Admin State      Fwding State
-----
1/2/20:0         10.10.10.104    N/A              N/A
N/A              Down            Not Fwding
3/2/4:50/5       143.144.145.1  N/A              N/A
N/A              Up              Fwding
-----
Number of static hosts : 2
=====
*A:ALA-48#
```

## stp

<b>Syntax</b>	<b>stp [detail]</b>
<b>Context</b>	show>service>id
<b>Description</b>	Displays information for the spanning tree protocol instance for the service.
<b>Parameters</b>	<b>detail</b> — Displays detailed information.
<b>Output</b>	<b>Show Service-ID STP Output</b> — The following table describes show service-id STP output fields:

Label	Description
RSTP Admin State	Indicates the administrative state of the Rapid Spanning Tree Protocol instance associated with this service.
Core Connectivity	Indicates the connectivity status to the core.
RSTP Oper State	Indicates the operational state of the Rapid Spanning Tree Protocol instance associated with this service. This field is applicable only when STP is enabled on the router.
Bridge-id	Specifies the MAC address used to identify this bridge in the network.
Hold Time	Specifies the interval length during which no more than two Configuration BPDUs shall be transmitted by this bridge.
Bridge fwd delay	Specifies how fast a bridge changes its state when moving toward the forwarding state.
Bridge Hello time	Specifies the amount of time between the transmission of Configuration BPDUs.
Bridge max age	Specifies the maximum age of spanning tree protocol information learned from the network on any port before it is discarded. This is the actual value that this bridge is currently using.
Bridge priority	Defines the priority of the spanning tree protocol instance associated with this service.
Topology change	Specifies whether a topology change is currently in progress.
Last Top. change	Specifies the time (in hundredths of a second) since the last time a topology change was detected by the Spanning Tree Protocol instance associated with this service.
Top. change count	Specifies the total number of topology changes detected by the Spanning Tree Protocol instance associated with this service since the management entity was last reset or initialized.

**Sample Output**

```

A:ALA-12# show service id 1 stp
=====
Spanning Tree Information
=====
VPLS Spanning Tree Information
-----
RSTP Admin State   : Up                RSTP Oper State   : Down
Core Connectivity   : Down
Bridge-id           : 04:67:ff:00:00:01

Hold Timer          : 1                  Bridge fwd delay   : 15
Bridge Hello time    : 1                  Bridge max age      : 20
Bridge priority      : 1                  Topology change     : Inactive
Last Top. change     : 0d 00:00:00        Top. change count   : 0

Root bridge-id      : 00:03:fa:00:00:00

Root path cost       : 1                  Root forward delay  : 15
Root hello time      : 1                  Root max age        : 20
Root priority        : 0                  Root port           : vcp

-----
Spanning Tree Specifics
-----
SAP Identifier       : 1/1/7:0           RSTP State         : Down
STP Port State       : Forwarding         BPDU encap         : dot1d
Port Number          : 2048               Priority            : 128
Cost                  : 10                 Fast Start          : Disabled
Designated Port      : 34816              Designated Bridge   : 02:fa:00:04:54:01
=====
A:ALA-12#
ALA-1# show service id 1 stp
=====
Inherited Rstp State (from mVPLS), Service 1
=====
Sap/Spoke Id        Oper-   Prune-   Port-   Managed by   Managed by
                    State    State    State    Service      Sap/spoke
-----
1/1/1:2             Down    Pruned   Discarding 101          1/1/1:1
100:1               Up      Pruned   Discarding 100          100:100
=====

```

**authentication**

<b>Syntax</b>	<b>authentication</b>
<b>Context</b>	show>service>id
<b>Description</b>	This command enables the context to show session authentication information.



## statistics

<b>Syntax</b>	<b>statistics</b> [ <i>policy name</i> ] [ <b>sap</b> <i>sap-id</i> ]
<b>Context</b>	show>service>id>authentication
<b>Description</b>	This command displays session authentication statistics for this service.
<b>Parameters</b>	<p><b>policy name</b> — Specifies an existing authentication policy name.</p> <p><b>sap-id</b> — Specifies the physical port identifier portion of the SAP definition. See <a href="#">Common Service Commands on page 2062</a> for <i>sap-id</i> command syntax.</p>

**Sample Output**

```
*A:ALA-48# show service id 700 authentication statistics
=====
Authentication Statistics for service 700
=====
Client Packets Authenticate Fail   : 0
Client Packets Authenticate Ok    : 0
=====
*A:ALA-48#
```

## subscriber-hosts

<b>Syntax</b>	<b>subscriber-hosts</b> [ <b>sap</b> <i>sap-id</i> ] [ <b>ip</b> <i>ip-address</i> [/ <i>mask</i> ]] [ <b>mac</b> <i>ieee-address</i> ] [ <b>sub-profile</b> <i>sub-profile-name</i> ] [ <b>sla-profile</b> <i>sla-profile-name</i> ] [ <b>detail</b> ]
<b>Context</b>	show>service>id
<b>Description</b>	This command displays subscriber host information.
<b>Parameters</b>	<p><b>sap</b> <i>sap-id</i> — Specifies the physical port identifier portion of the SAP definition. See <a href="#">Common Service Commands on page 2062</a> for <i>sap-id</i> command syntax.</p> <p><b>ip</b> <i>ip-address</i>[/<i>mask</i>] — Shows information for the specified IP address and mask.</p> <p><b>mac</b> <i>ieee-address</i> — Displays information only for the specified 48-bit MAC address. The MAC address can be expressed in the form <i>aa:bb:cc:dd:ee:ff</i> or <i>aa-bb-cc-dd-ee-ff</i> where <i>aa</i>, <i>bb</i>, <i>cc</i>, <i>dd</i>, <i>ee</i> and <i>ff</i> are hexadecimal numbers.</p> <p><b>profile</b> <i>sub-profile-name</i> — Displays an existing subscriber profile name to be associated with the static subscriber host. The subscriber profile is configured in the <b>config&gt;subscr-mgmt&gt;sub-profile</b> context.</p> <p><b>sla-profile</b> <i>sla-profile-name</i> — Displays this optional parameter to specify an existing SLA profile name.</p> <p><b>detail</b> — Displays detailed information.</p>
<b>Output</b>	<b>Show Service-ID subscriber hosts</b> — The following table describes show service-id subscriber hosts output fields:

Label	Description
Sap	Displays the SAP ID number.
IP Address	Displays the IP address.
MAC Address	Displays the MAC address
Origin Subscriber	The ID of the originating subscriber.
Redirection filter id	Displays the Redirection Filter ID number.
Status: active/inactive	Displays the status of one-time HTTP redirection.
Filter-id-source	Displays source of the HTTP filter.

### Sample Output

```
*A:Dut-C># show service id 3 subscriber-hosts
=====
Subscriber Host table
=====
Sap              Subscriber
  IP Address
  MAC Address    PPPoE-SID   Origin      Fwding State
-----
2/1/5:2          TEACAHEH74
  11.11.1.61
    00:80:00:00:00:0a      N/A      ARP-Host    Fwding
[pw-11:11]
  11.11.1.2
    00:00:11:11:01:02      N/A      ARP-Host    Fwding
[pw-11:12]
  11.11.1.3
    00:00:11:11:01:03      N/A      ARP-Host    Fwding
[pw-11:13]
  11.11.1.4
    00:00:11:11:01:04      N/A      ARP-Host    Fwding
[pw-22:22]
  22.22.1.2
    00:00:22:22:01:02      N/A      ARP-Host    Fwding
[pw-33:33]
  33.33.1.2
    00:00:33:33:01:02      N/A      ARP-Host    Fwding
-----
Number of subscriber hosts: 6
=====
*A:Dut-C>#
```

```
A:Dut-A# show service id 100 subscriber-hosts ip 10.100.1.5
=====
Subscriber Host table
=====
Sap              IP Address    MAC Address    Origin(*) Subscriber
-----
1/2/1:102        10.100.1.5    00:10:00:00:00:03 -/D/-  alcatel_100
```

```

-----
Number of subscriber hosts : 1
=====
(*) S=Static Host, D=DHCP Lease, N=Non-Sub-Traffic
A:Dut-A#

```

## sdp

**Syntax**

```

sdp sdp-id pw-port [pw-port-id]
sdp sdp-id pw-port
sdp sdp-id pw-port [pw-port-id] [statistics]
sdp [consistent | inconsistent | na] egressifs
sdp sdp-id keep-alive-history
sdp far-end ip-address | ipv6-address keep-alive-history
sdp [sdp-id] detail
sdp far-end ip-address | ipv6-address detail

```

**Context** show>service>sdp

**Description** This command displays SDP information.

If no optional parameters are specified, a summary SDP output for all SDPs is displayed.

**Parameters** *sdp-id* — The SDP ID for which to display information.

**Default** All SDPs.

**Values** 1 — 17407

*pw-port pw-port-id* — Displays the SAP identifier for PW-SAPs.

**Values** 1 — 10239

*far-end ip-address* — Displays only SDPs matching with the specified far-end IP address.

**Default** SDPs with any far-end IP address.

*detail* — Displays detailed SDP information.

**Default** SDP summary output.

*keep-alive-history* — Displays the last fifty SDP keepalive events for the SDP.

**Default** SDP summary output.

### Sample Output

```

=====
*A:ALA-12>config>service# show service sdp 1 pw-port
=====
Service Destination Point (sdp Id 1 Pw-Port)
=====
Pw-port   VC-Id   Adm     Encap    Opr     VC Type   Egr      Monitor
          Shaper  Oper
          VPort  Group
-----
1          1       up      dot1q    up      ether

```

## Show Commands

```
2          2          up      qinq      up      ether
3          3          up      dot1q     up      ether
4          4          up      qinq      up      ether
-----
Entries found : 4
=====

*A:ALA-12>config>service# show service sdp 1 pw-port 3
=====
Service Destination Point (Sdp Id 1 Pw-Port 3)
=====
SDP Binding port      : lag-1
VC-Id                 : 3                      Admin Status       : up
Encap                 : dot1q                  Oper Status        : up
VC Type               : ether
Oper Flags            : (Not Specified)
Monitor Oper-Group    : (Not Specified)
=====

*A:ALA-12>config>service# show service sdp 1 pw-port 3 statistics
=====
Service Destination Point (Sdp Id 1 Pw-Port 3)
=====
SDP Binding port      : lag-1
VC-Id                 : 3                      Admin Status       : up
Encap                 : dot1q                  Oper Status        : up
VC Type               : ether
Oper Flags            : (Not Specified)
Monitor Oper-Group    : (Not Specified)

Statistics            :
I. Fwd. Pkts.         : 0                      I. Dro. Pkts.      : 0
I. Fwd. Octs.         : 0                      I. Dro. Octs.      : 0
E. Fwd. Pkts.         : 0                      E. Fwd. Octets     : 0
=====
```

## subscriber-using

<b>Syntax</b>	<b>subscriber-using</b> [ <b>service-id</b> <i>service-id</i> ] [ <b>sap-id</b> <i>sap-id</i> ] [ <b>interface</b> <i>ip-int-name</i> ] [ <b>ip</b> <i>ip-address</i> [/ <i>mask</i> ]] [ <b>mac</b> <i>ieee-address</i> ] [ <b>sub-profile</b> <i>sub-profile-name</i> ] [ <b>sla-profile</b> <i>sla-profile-name</i> ] [ <b>app-profile</b> <i>app-profile-name</i> ] [ <b>port</b> <i>port-id</i> ] [ <b>no-inter-dest-id</b>   <b>inter-dest-id</b> <i>intermediate-destination-id</i> ]
<b>Context</b>	show>service
<b>Description</b>	This command displays selective subscriber information using specific options.
<b>Parameters</b>	<p><b>service-id</b> <i>service-id</i> — Displays information for the specifies ID that uniquely identifies a service.</p> <p><b>sap-id</b> <i>sap-id</i> — Displays the physical port identifier portion of the SAP definition. See <a href="#">Common Service Commands on page 2062</a> for <i>sap-id</i> command syntax.</p> <p><b>interface</b> <i>ip-int-name</i> — Shows DHCP statistics on the specified interface.</p> <p><b>port</b> <i>port-id</i> — Indicates the SAP or SDP for which this entry contains information.</p> <p><b>ip</b> <i>ip-address</i>[/<i>mask</i>] — Shows information for the specified IP address and mask.</p> <p><b>mac</b> <i>ieee-address</i> — Displays information only for the specified 48-bit MAC address. The MAC address can be expressed in the form <i>aa:bb:cc:dd:ee:ff</i> or <i>aa-bb-cc-dd-ee-ff</i> where <i>aa</i>, <i>bb</i>, <i>cc</i>, <i>dd</i>, <i>ee</i> and <i>ff</i> are hexadecimal numbers.</p> <p><b>sub-profile</b> <i>sub-profile-name</i> — Displays an existing subscriber profile name to be associated with the static subscriber host. The subscriber profile is configured in the <b>config&gt;subscr-mgmt&gt;sub-profile</b> context.<i>sla-profile-name</i></p> <p><b>sla-profile</b> — Displays this optional parameter to specify an existing SLA profile name.</p> <p><b>app-profile</b> — Displays the application specified profile.</p> <p><b>inter-dest-id</b> <i>intermediate-destination-id</i> — Indicates the intermediate destination identifier received from either the DHCP or the RADIUS server or the local user database.</p>

## Sample Output

```
A:Dut-A# show service subscriber-using service-id 100
=====
Subscribers
=====
Subscriber                               Sub Profile
-----
alcatel_100                             sub_prof100
-----
Matching Subscribers : 1
=====
A:Dut-A#

A:Dut-A# show service subscriber-using
=====
Subscribers
=====
Subscriber                               Sub Profile
-----
alcatel_100                             sub_prof100
alcatel_110                             sub_prof110
-----
```

Show Commands

```
alcatel_120                sub_prof120
alcatel_130                sub_prof130
alcatel_80                 sub_prof80
alcatel_90                 sub_prof90
client_PC1                 sub_profPC1
static                     sub_default
-----
Matching Subscribers : 8
=====
A:Dut-A#
```

redundancy

<b>Syntax</b>	<b>redundancy</b>
<b>Context</b>	show
<b>Description</b>	This command enables the context to show multi-chassis redundancy information.

multi-chassis

<b>Syntax</b>	<b>multi-chassis all</b> <b>multi-chassis mc-lag peer</b> <i>ip-address</i> [ <b>lag</b> <i>lag-id</i> ] <b>multi-chassis mc-lag</b> [ <b>peer</b> <i>ip-address</i> [ <b>lag</b> <i>lag-id</i> ]] <b>statistics</b> <b>multi-chassis sync</b> [ <b>peer</b> <i>ip-address</i> ] [ <b>detail</b> ] <b>multi-chassis sync</b> [ <b>peer</b> <i>ip-address</i> ] <b>statistics</b>
<b>Context</b>	show>redundancy
<b>Description</b>	This command displays multi-chassis redundancy information.
<b>Parameters</b>	<b>all</b> — Displays all multi-chassis information. <b>mc-lag</b> — Displays multi-chassis LAG information. <b>peer</b> <i>ip-address</i> — Displays the address of the multi-chassis peer. <b>lag</b> <i>lag-id</i> — Displays the specified LAG ID on this system that forms an multi-chassis LAG configuration with the indicated peer. <b>statistics</b> — Displays statistics for the multi-chassis peer. <b>sync</b> — Displays synchronization information. <b>detail</b> — Displays detailed information.

Sample Output

```
A:pc1# show redundancy multi-chassis all
=====
Multi-Chassis Peers
=====
Peer IP          Src IP          Auth          Peer Admin
MCS Admin       MCS Oper       MCS State     MC-LAG Admin   MC-LAG Oper
```

```

-----
10.10.10.102      10.10.10.101      hash      Enabled
      Enabled      Enabled      inSync      Enabled      Enabled
10.10.20.1        0.0.0.0        None      Disabled
      --          --          --          Disabled      Disabled
=====
A:pc1#

*A:Dut-C# show redundancy multi-chassis mc-lag peer 10.10.10.1
=====
Multi-Chassis MC-Lag Peer 10.10.10.1
=====
Last State chg: 09/24/2007 07:58:03
Admin State: Up      Oper State : Up
KeepAlive: 10 deci-seconds      Hold On Ngbr Failure : 3
-----
Lag Id LACP Key Remote Lag Id System Id Sys Prio Last State Changed
-----
1      326661      00:00:00:33:33:33  32888  09/24/2007 07:56:35
-----
Number of LAGs : 1
=====
*A:Dut-C#

A:pc1# show redundancy multi-chassis mc-lag statistics
=====
Multi-Chassis Statistics
=====
Packets Rx      : 129816
Packets Rx Keepalive      : 129798
Packets Rx Config      : 3
Packets Rx Peer Config      : 5
Packets Rx State      : 10
Packets Dropped KeepaliveTask      : 0
Packets Dropped Packet Too Short      : 0
Packets Dropped Verify Failed      : 0
Packets Dropped Tlv Invalid Size      : 0
Packets Dropped Out of Seq      : 0
Packets Dropped Unknown Tlv      : 0
Packets Dropped Tlv Invalid LagId      : 0
Packets Dropped MD5      : 0
Packets Dropped Unknown Peer      : 0
Packets Tx      : 77918
Packets Tx Keepalive      : 77879
Packets Tx Config      : 6
Packets Tx Peer Config      : 26
Packets Tx State      : 7
Packets Tx Failed      : 0
=====
A:pc1#

A:pc1# show redundancy multi-chassis mc-lag peer 10.10.10.102 lag 2 statistics
=====
Multi-Chassis Statistics, Peer 10.10.10.102 Lag 2
=====
Packets Rx Config      : 1
Packets Rx State      : 4
Packets Tx Config      : 2
Packets Tx State      : 3

```

## Show Commands

```
Packets Tx Failed          : 0
=====
A:pc1#

A:pc1#show redundancy multi-chassis mc-lag peer 10.10.10.102 statistics
=====
Multi-Chassis Statistics, Peer 10.10.10.102
=====
Packets Rx                  : 129918
Packets Rx Keepalive        : 129900
Packets Rx Config           : 3
Packets Rx Peer Config      : 5
Packets Rx State            : 10
Packets Dropped State Disabled : 0
Packets Dropped Packets Too Short : 0
Packets Dropped Tlv Invalid Size : 0
Packets Dropped Tlv Invalid LagId : 0
Packets Dropped Out of Seq   : 0
Packets Dropped Unknown Tlv  : 0
Packets Dropped MD5          : 0
Packets Tx                  : 77979
Packets Tx Keepalive         : 77940
Packets Tx Peer Config       : 26
Packets Tx Failed           : 0
=====
A:pc1#

A:pc1# show redundancy multi-chassis sync
=====
Multi-chassis Peer Table
=====
Peer
-----
Peer IP Address      : 10.10.10.102
Description          : CO1
Authentication       : Enabled
Source IP Address    : 10.10.10.101
Admin State          : Enabled
-----
Sync-status
-----
Client Applications   :
Sync Admin State     : Up
Sync Oper State      : Up
DB Sync State        : inSync
Num Entries          : 0
Lcl Deleted Entries  : 0
Alarm Entries        : 0
Rem Num Entries       : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries     : 0
=====
Peer
-----
Peer IP Address      : 10.10.20.1
Authentication       : Disabled
Source IP Address    : 0.0.0.0
Admin State          : Disabled
=====
```



A:pc1#

pc1# show redundancy multi-chassis sync peer 10.10.10.102

=====

Multi-chassis Peer Table

=====

Peer

-----

Peer IP Address : 10.10.10.102  
 Description : CO1  
 Authentication : Enabled  
 Source IP Address : 10.10.10.101  
 Admin State : Enabled

-----

Sync-status

-----

Client Applications :  
 Sync Admin State : Up  
 Sync Oper State : Up  
 DB Sync State : inSync  
 Num Entries : 0  
 Lcl Deleted Entries : 0  
 Alarm Entries : 0  
 Rem Num Entries : 0  
 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 : 0  
 Lcl Deleted Entries : 0  
 Alarm Entries : 0

-----

Rem Num Entries : 0  
 Rem Lcl Deleted Entries : 0  
 Rem Alarm Entries : 0

-----

Application : srrp  
 Num Entries : 0  
 Lcl Deleted Entries : 0  
 Alarm Entries : 0

-----

## Show Commands

```
Rem Num Entries      : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries    : 0
=====
A:pc1#

A:pc1# show redundancy multi-chassis sync peer 10.10.10.102 detail
=====
Multi-chassis Peer Table
=====
Peer
-----
Peer IP Address      : 10.10.10.102
Description          : CO1
Authentication       : Enabled
Source IP Address    : 10.10.10.101
Admin State          : Enabled
-----
Sync-status
-----
Client Applications  :
Sync Admin State     : Up
Sync Oper State      : Up
DB Sync State        : inSync
Num Entries          : 0
Lcl Deleted Entries  : 0
Alarm Entries        : 0
Rem Num Entries      : 0
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          : 0
Lcl Deleted Entries  : 0
Alarm Entries        : 0
-----
Rem Num Entries      : 0
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
=====
Ports synced on peer 10.10.10.102
=====
Port/Encap           Tag
-----
1/1/1
  1-2                r1
=====
A:pc1#

A:pc1# show redundancy multi-chassis sync statistics
=====
Multi-chassis Peer Sync Stats
=====
Peer IP Address      : 10.10.10.102
Packets Tx Total     : 511
Packets Tx Hello     : 510
Packets Tx Data      : 0
Packets Tx Other     : 1
Packets Tx Error     : 0
Packets Rx Total     : 511
Packets Rx Hello     : 510
Packets Rx Data      : 0
Packets Rx Other     : 1
Packets Rx Error     : 0
Packets Rx Header Err : 0
Packets Rx Body Err  : 0
Packets Rx Seq Num Err : 0
=====
Peer IP Address      : 10.10.20.1
Packets Tx Total     : 0
Packets Tx Hello     : 0
Packets Tx Data      : 0
Packets Tx Other     : 0
Packets Tx Error     : 0
Packets Rx Total     : 0
Packets Rx Hello     : 0
Packets Rx Data      : 0
Packets Rx Other     : 0
Packets Rx Error     : 0
Packets Rx Header Err : 0
Packets Rx Body Err  : 0
Packets Rx Seq Num Err : 0
=====
A:pc1#

A:pc1# show redundancy multi-chassis sync peer 10.10.10.102 statistics
=====
Multi-chassis Peer Sync Stats
=====
Peer IP Address      : 10.10.10.102
Packets Tx Total     : 554
Packets Tx Hello     : 553
Packets Tx Data      : 0

```

## Show Commands

```
Packets Tx Other      : 1
Packets Tx Error      : 0
Packets Rx Total      : 554
Packets Rx Hello      : 553
Packets Rx Data       : 0
Packets Rx Other      : 1
Packets Rx Error      : 0
Packets Rx Header Err : 0
Packets Rx Body Err   : 0
Packets Rx Seq Num Err : 0
=====
A:pc1#
```

## mc-ipsec

- Syntax** **mc-ipsec peer *addr* [tunnel-group *group-id*]**
- Context** show>redundancy>multi-chassis
- Description** This command displays the IPsec multi-chassis states. Optionally, only the states of the specified tunnel-groups will be displayed.
- Parameters** *addr* — Specifies the address of the peer.  
*group-id* — Specifies the tunnel-group ID.

### Sample Output

```
=====
Multi-Chassis MC-IPSec
=====
Peer Name      : (Not Specified)
Peer Addr      : 2.2.2.2
Keep Alive Intvl: 10                Hold on Nbr Fail      : 3
BFD Intf Name  : None
BFD Dest Addr  :
Last update    : 03/20/2012 22:48:55

=====
Multi-Chassis IPsec Multi Active Tunnel-Group Table
=====
ID              Peer Group      Priority  Preempt   Admin State  Mastership
-----
1               1                100      Disabled  Up           eligible
-----
Multi Active Tunnel Group Entries found: 1
=====
```

## mc-ring

<b>Syntax</b>	<b>mc-ring peer <i>ip-address</i> statistics</b> <b>mc-ring peer <i>ip-address</i> [ring <i>sync-tag</i> [detail statistics]]</b> <b>mc-ring peer <i>ip-address</i> ring <i>sync-tag</i> ring-node [<i>ring-node-name</i> [detail statistics]]</b> <b>mc-ring global-statistics</b>
<b>Context</b>	show>redundancy>multi-chassis
<b>Description</b>	This command displays multi-chassis ring information.
<b>Parameters</b>	<p><i>ip-address</i> — Specifies the address of the multi-chassis peer to display.</p> <p><b>ring <i>sync-tag</i></b> — Specifies a synchronization tag to be displayed that was used while synchronizing this port with the multi-chassis peer.</p> <p><b>node <i>ring-node-name</i></b> — Specifies a ring-node name.</p> <p><b>global-statistics</b> — Displays global statistics for the multi-chassis ring.</p> <p><b>detail</b> — Displays detailed peer information for the multi-chassis ring.</p>
<b>Output</b>	<b>Show mc-ring peer ip-address ring Output</b> — The following table describes mc-ring peer ip-address ring output fields.

Label	Description
Sync Tag	Displays the synchronization tag that was used while synchronizing this port with the multi-chassis peer.
Oper State	<p><b>noPeer</b> — The peer has no corresponding ring configured.</p> <p><b>connected</b> — The inband control connection with the peer is operational.</p> <p><b>broken</b> — The inband control connection with the peer has timed out.</p> <p><b>conflict</b> — The inband control connection with the peer has timed out but the physical connection is still OK; the failure of the inband signaling connection is caused by a misconfiguration. For example, a conflict between the configuration of this system and its peer, or a misconfiguration on one of the ring access node systems.</p> <p><b>testingRing</b> — The inband control connection with the peer is being set up. Waiting for result.</p> <p><b>waitingForPeer</b> — Verifying if this ring is configured on the peer.</p> <p><b>configErr</b> — The ring is administratively up, but a configuration error prevents it from operating properly.</p> <p><b>halfBroken</b> — The inband control connection indicates that the ring is broken in one direction (towards the peer).</p>

Label	Description (Continued)
	localBroken — The inband control connection with the peer is known to be broken due to local failure or local administrative action.
	shutdown — The ring is shutdown.
Failure Reason	
No. of MC Ring entries	

### Sample Output

```
show redundancy multi-chassis mc-ring peer 10.0.0.2 ring ring11 detail
=====
Multi-Chassis MC-Ring Detailed Information
=====
Peer           : 10.0.0.2
Sync Tag       : ring11
Port ID        : 1/1/3
Admin State    : inService
Oper State     : connected
Admin Change   : 01/07/2008 21:40:07
Oper Change    : 01/07/2008 21:40:24
Failure Reason : None
-----
In Band Control Path
-----
Service ID     : 10
Interface Name : to_an1
Oper State     : connected
Dest IP        : 10.10.0.2
Src IP         : 10.10.0.1
-----
VLAN Map B Path Provisioned
-----
range 13-13
range 17-17
-----
VLAN Map Excluded Path Provisioned
-----
range 18-18
-----
VLAN Map B Path Operational
-----
range 13-13
range 17-17
-----
VLAN Map Excluded Path Operational
-----
range 18-18
=====
```

```
*A:ALA-48>show>redundancy>multi-chassis# mc-ring peer 192.251.10.104
```

```
MC Ring entries
```

```
=====
```

Sync Tag	Oper State	Failure Reason
-----		

```
No. of MC Ring entries: 0
```

```
show redundancy multi-chassis mc-ring peer 10.0.0.2
```

```
MC Ring entries
```

```
=====
```

Sync Tag	Oper State	Failure Reason
-----		

ring11	connected	None
ring12	shutdown	None

```
No. of MC Ring entries: 4
```

```
show redundancy multi-chassis mc-ring peer 10.0.0.2 ring ring11 ring-node an1 detail
```

```
Multi-Chassis MC-Ring Node Detailed Information
```

```
=====
```

Peer	: 10.0.0.2
Sync Tag	: ring11
Node Name	: an1
Oper State Loc	: connected
Oper State Rem	: notTested
In Use	: True
Admin Change	: 01/07/2008 21:40:07
Oper Change	: 01/07/2008 21:40:25
Failure Reason	: None

```
Ring Node Connectivity Verification
```

```
=====
```

Admin State	: inService
Service ID	: 11
VLAN Tag	: 11
Dest IP	: 10.11.3.1
Src IP	: None
Interval	: 1 minutes
Src MAC	: None

```
show redundancy multi-chassis mc-ring peer 10.0.0.2 ring ring11 ring-node
```

```
MC Ring Node entries
```

```
=====
```

Name	Loc Oper St.	Failure Reason
In Use	Rem Oper St.	
-----		
an1	connected	None
Yes	notTested	
an2	connected	None
Yes	notTested	

No. of MC Ring Node entries: 2

=====

**show redundancy multi-chassis ring peer statistics Output** — The following table describes multi-chassis ring peer output fields

Label	Description
Message	Displays the message type.
Received	Indicates the number of valid MC-Ring signaling messages received from the peer.
Transmitted	Indicates the number of valid MC-Ring signaling messages transmitted from the peer.
MCS ID Request	Displays the number of valid MCS ID requests were received from the peer.
MCS ID Response	Displays the number of valid MCS ID responses were received from the peer.
Ring Exists Request	Displays the number of valid 'ring exists' requests were received from the peer.
Ring Exists Response	Displays the number of valid ring exists' responses were received from the peer.
Keepalive	Displays the number of valid MC-Ring control packets of type 'keep-alive' were received from the peer.

```
*A:ALA-48>show>redundancy>multi-chassis# mc-ring peer 192.251.10.104 statistics
=====
MC Ring statistics for peer 192.251.10.104
=====
Message                                Received      Transmitted
-----
MCS ID Request                         0             0
MCS ID Response                        0             0
Ring Exists Request                    0             0
Ring Exists Response                   0             0
Keepalive                             0             0
-----
Total                                  0             0
=====
*A:ALA-48>show>redundancy>multi-chassis#
```

#### show mc-ring ring-node Output

Label	Description
Oper State	Displays the state of the connection verification (both local and remote).



Label	Description (Continued)
	notProvisioned — Connection verification is not provisioned.
	configErr — Connection verification is provisioned but a configuration error prevents it from operating properly.
	notTested — Connection verification is administratively disabled or is not possible in the current situation.
	testing — Connection Verification is active, but no results are yet available.
	connected — The ring node is reachable.
	disconnected — Connection verification has timed out.
In Use	Displays “True” if the ring node is referenced on an e-pipe or as an inter-dest-id on a static host or dynamic lease.

**show mc-ring global-statistics Output**

Label	Description
Rx	Displays the number of MC-ring signaling packets were received by this system.
Rx Too Short	Displays the number of MC-ring signaling packets were received by this system that were too short.
Rx Wrong Authentication	Displays the number of MC-ring signaling packets were received by this system with invalid authentication.
Rx Invalid TLV	Displays the number of MC-ring signaling packets were received by this system with invalid TLV.
Rx Incomplete	Displays the number of MC-ring signaling packets were received by this system that were incomplete.
Rx Unknown Type	Displays the number of MC-ring signaling packets were received by this system that were of unknown type.
Rx Unknown Peer	Displays the number of MC-ring signaling packets were received by this system that were related to an unknown peer.
Rx Unknown Ring	Displays the number of MC-ring signaling packets were received by this system that were related to an unknown ring.
Rx Unknown Ring Node	Displays the number of MC-ring signaling packets were received by this system that were related to an unknown ring node.

Label	Description (Continued)
Tx	Displays the number of MC-ring signaling packets were transmitted by this system.
Tx No Buffer	Displays the number of MC-ring signaling packets could not be transmitted by this system due to a lack of packet buffers.
Tx Transmission Failed	Displays the number of MC-ring signaling packets could not be transmitted by this system due to a transmission failure.
Tx Unknown Destination	Displays the number of MC-ring <b>unknown destination</b> signaling packets were transmitted by this system.
Missed Configuration Events	Displays the number of missed configuration events on this system.
Missed BFD Events	Displays the number of missed BFD events on this system.

### Sample Output

```
*A:ALA-48>show>redundancy>multi-chassis# mc-ring global-statistics
=====
Global MC Ring statistics
=====
Rx                               : 0
Rx Too Short                     : 0
Rx Wrong Authentication          : 0
Rx Invalid TLV                   : 0
Rx Incomplete                    : 0
Rx Unknown Type                  : 0
Rx Unknown Peer                  : 0
Rx Unknown Ring                  : 0
Rx Unknown Ring Node             : 0
Tx                               : 36763
Tx No Buffer                      : 0
Tx Transmission Failed           : 0
Tx Unknown Destination           : 0
Missed Configuration Events      : 0
Missed BFD Events                : 0
=====
*A:ALA-48>show>redundancy>multi-chassis#
```

## lease-state

<b>Syntax</b>	<b>lease-state</b> [ <b>sap</b> <i>sap-id</i>   <b>sdp</b> <i>sdp-id:vc-id</i>   <b>interface</b> <i>interface-name</i>   <b>ip-address</b> <i>ip-address[/mask]</i>   <b>chaddr</b> <i>ieee-address</i>   <b>mac</b> <i>ieee-address</i>   {[ <b>port</b> <i>port-id</i> ] [ <b>no-inter-dest-id</b>   <b>inter-dest-id</b> <i>inter-dest-id</i> ]}] [ <b>detail</b> ]		
<b>Context</b>	show>service>id>dhcp		
<b>Description</b>	This command displays DHCP lease state information.		
<b>Parameters</b>	<b>Values</b>	service-id: 1 — 214748364	

**sap** *sap-id* — Specifies the physical port identifier portion of the SAP definition. See [Common Service Commands on page 2062](#) for *sap-id* command syntax.

**sdp** *sdp-id* — The SDP identifier.

**Values** 1 — 17407

**vc-id** — The virtual circuit identifier. This value is used to validate the VC ID portion of each mesh SDP binding defined in the service. The default value of this object is equal to the service ID.

**Values** 1 — 4294967295

**ip** *ip-address[/mask]* — Shows information for the specified IP address and mask.

**port** *port-id* — The DHCP lease state local specifies that the DHCP lease state is learned by either a SAP or SDP. When the value is SAP, the value indicates the SAP for which this entry contains information.

**chaddr** — Specifies the MA address of the DHCP lease state.

**interface** *interface-name* — Shows information for the specified IP interface.

**detail** — Displays detailed lease state information.

**inter-dest-id** — Indicates the intermediate destination identifier received from either the DHCP or the RADIUS server or the local user database.

### Sample Output

```
*A:ALA-48>config# show service id 101 dhcp lease-state
=====
DHCP lease state table, service 101
=====
IP Address      Mac Address      Sap/Sdp Id      Remaining   Lease   MC
                  LifeTime      Origin      Stdby
-----
102.1.1.52      00:00:1f:bd:00:bb lag-1:101      00h02m56s   DHCP-R
103.3.2.62      00:00:1f:bd:00:c6 lag-1:105      00h02m59s   RADIUS
-----
Number of lease states : 2
=====
*A:ALA-48>config#

*A:ALA-48>config# show service id 105 dhcp lease-state wholesaler 101
=====
DHCP lease state table, service 105
=====
IP Address      Mac Address      Sap/Sdp Id      Remaining   Lease   MC
                  LifeTime      Origin      Stdby
-----
Wholesaler 101 Leases
-----
103.3.2.62      00:00:1f:bd:00:c6 lag-1:105      00h00m39s   RADIUS
-----
Number of lease states : 1
=====
*A:ALA-48>config#
```

## statistics

<b>Syntax</b>	<b>statistics</b> <b>[[sap sap-id]][sdp sdp-id:vc-id]][interface interface-name]]</b>
<b>Context</b>	show>service>id>dhcp
<b>Description</b>	This command displays DHCP relay statistics.
<b>Parameters</b>	<b>interface ip-int-name</b> — Display DHCP statistics on the specified interface. <b>interface interface-name</b> — Displays DHCP statistics for the specified interface name. <b>sap sap-id</b> — Displays DHCP statistics for the specified SAP. See <a href="#">Common Service Commands on page 2062</a> for <i>sap-id</i> command syntax.

## Sample Output

```
*A:ALA-48# show service id 88 dhcp statistics interface SpokeTerm
=====
DHCP Statistics for interface SpokeTerm
=====
Rx Packets                               : 0
Tx Packets                               : 0
Rx Malformed Packets                     : 0
Rx Untrusted Packets                     : 0
Client Packets Discarded                  : 0
Client Packets Relayed                    : 0
Client Packets Snooped                    : 0
Client Packets Proxied (RADIUS)           : 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
=====
*A:ALA-48#
```

## summary

<b>Syntax</b>	<b>summary</b> <b>[interface interface-name   saps]</b>
<b>Context</b>	show>service>id>dhcp
<b>Description</b>	This command displays DHCP configuration summary information.
<b>Parameters</b>	<b>interface interface-name</b> — Displays summary information for the specified existing interface. <b>sap</b> — Displays summary information for SAPs per interface.

## Sample Output

```
IES:
*A:ALA-48>show>service>id>dhcp# summary
=====
```

DHCP Summary, service 700

```
=====
Sap/Sdp                Snoop  Used/  Info  Admin
                       Provided Option State
-----
sap:1/1/9:0            No     0/0    Keep  Down
sap:1/1/25:0           No     0/0    Keep  Down
sdp:8:700              No     N/A    N/A   N/A
-----
Number of Entries : 3
-----
*A:ALA-48>show>service>id>dhcp#
```

VPLS:

\*A:ALA-49>show>service# id 700 dhcp summary

```
=====
DHCP Summary, service 700
=====
Sap/Sdp                Snoop  Used/  Arp Reply  Info  Admin
                       Provided Agent   Option State
-----
sap:1/1/9:0            No     0/0    No          Keep  Down
sdp:2:222              No     N/A    N/A         N/A   N/A
sdp:2:700              No     N/A    N/A         N/A   N/A
-----
Number of Entries : 3
-----
*A:ALA-49>show>service#
```

\*A:ALA-48# show service id 88 dhcp summary saps

```
=====
DHCP Summary, service 88
=====
Interface Name          Arp    Used/  Info  Admin
SapId/Sdp              Populate Provided Option State
-----
SpokeTerm               No     0/0    Keep  Up
  sdp:spoke-3:3         0/0
new-if                  No     0/1    Keep  Up
  sap:1/2/19:0          0/1
test123                 No     0/0    Keep  Up
  sap:3/2/4:50/5        0/0
testabc                 No     0/0    Keep  Up
  sap:1/2/20:0          0/0
-----
Interfaces: 4
=====
*A:ALA-48#
```

\*A:ALA-48# show service id 88 dhcp summary interface SpokeTerm

```
=====
DHCP Summary, service 88
=====
Interface Name          Arp    Used/  Info  Admin
SapId/Sdp              Populate Provided Option State
-----
SpokeTerm               No     0/0    Keep  Up
  sdp:spoke-3:3         0/0
```

```
-----  
Interfaces: 1  
-----  
*A:ALA-48#
```

statistics

- Syntax** **statistics** [**interface** *ip-int-name* | *ip-address*]
- Context** show>router>dhcp
- Description** This command displays statistics for DHCP relay and DHCP snooping. If no IP address or interface name is specified, then all configured interfaces are displayed. If an IP address or interface name is specified, then only data regarding the specified interface is displayed.
- Parameters** **interface** *ip-int-name* | *ip-address* — Displays statistics for the specified IP interface or IP address.
- Output** **Show DHCP Statistics Output** — The following table describes the output fields for DHCP statistics.

Label	Description
Received Packets	The number of packets received from the DHCP clients.
Transmitted Packets	The number of packets transmitted to the DHCP clients.
Received Malformed Packets	The number of malformed packets received from the DHCP clients.
Received Untrusted Packets	The number of untrusted packets received from the DHCP clients.
Client Packets Discarded	The number of packets received from the DHCP clients that were discarded.
Client Packets Relayed	The number of packets received from the DHCP clients that were forwarded.
Client Packets Snooped	The number of packets received from the DHCP clients that were snooped.
Server Packets Discarded	The number of packets received from the DHCP server that were discarded.
Server Packets Relayed	The number of packets received from the DHCP server that were forwarded.
Server Packets Snooped	The number of packets received from the DHCP server that were snooped.
Client packets proxied (RADIUS)	The number of packets that were generated from RADIUS data and not relayed from a server

Label	Description (Continued)
Client packets proxied (Lease-Split)	Indicates the total number of client packets proxied by the DHCP relay agent based on data received from a RADIUS server
DHCP RELEASES spoofed	Indicates the total number of DHCP release messages spoofed by the DHCP relay agent to the DHCP server.
DHCP FORCE-RENEWs spoofed	The number of DHCP force-renew packets sent to DHCP clients.

### Sample Output

```
A:SUB-Dut-A# show router 1000 dhcp statistics
=====
DHCP Global Statistics (Service: 1000)
=====
Rx Packets                : 16000
Tx Packets                : 15041
Rx Malformed Packets     : 0
Rx Untrusted Packets     : 0
Client Packets Discarded  : 423
Client Packets Relayed    : 0
Client Packets Snooped    : 0
*Client Packets Proxied (RADIUS) : 0
Client Packets Proxied (Lease-Split) : 0 *
Server Packets Discarded  : 0
Server Packets Relayed    : 0
Server Packets Snooped    : 0
*DHCP RELEASES Spoofed   : 0
DHCP FORCE-RENEWs Spoofed : 0 *
=====
A:SUB-Dut-A#
```

## summary

<b>Syntax</b>	<b>summary</b>
<b>Context</b>	show>router>dhcp
<b>Description</b>	Display the status of the DHCP Relay and DHCP Snooping functions on each interface.
<b>Output</b>	<b>Show DHCP Summary Output</b> — The following table describes the output fields for DHCP summary.

Label	Description
Interface Name	Name of the router interface.
ARP Populate	Indicates whether ARP populate is enabled.
Used/Provided	Indicates the number of used and provided DHCP leases.

Label	Description (Continued)
Info Option	Indicates whether Option 82 processing is enabled on the interface.
Admin State	Indicates the administrative state.

**Sample Output**

```
*A:ALA-48>show>router>dhcp# summary
=====
Interface Name           Arp      Used/    Info   Admin
                        Populate Provided Option State
-----
ccaiesif                 No        0/0     Keep   Down
ccanet6                  No        0/0     Keep   Down
iesBundle                 No        0/0     Keep   Up
spokeSDP-test            No        0/0     Keep   Down
test                     No        0/0     Keep   Up
test1                    No        0/0     Keep   Up
test2                    No        0/0     Keep   Up
testA                     No        0/0     Keep   Up
testB                     No        0/0     Keep   Up
testIES                   No        0/0     Keep   Up
to-web                   No        0/0     Keep   Up
-----
Interfaces: 11
=====
*A:ALA-48>show>router>dhcp#
```



## IGMP Snooping Show Commands

### igmp-snooping

<b>Syntax</b>	<b>igmp-snooping</b>
<b>Context</b>	show>service>id>
<b>Description</b>	This command enables the context to display IGMP snooping information.

### all

<b>Syntax</b>	<b>all</b>
<b>Context</b>	show>service>id>igmp-snooping
<b>Description</b>	Displays detailed information for all aspects of IGMP snooping on the VPLS service.
<b>Output</b>	<b>Show All Service-ID</b> — The following table describes the show all service-id command output fields:

Label	Description
Admin State	The administrative state of the IGMP instance.
Querier	Displays the address of the IGMP querier on the IP subnet to which the interface is attached.
Sap/Sdp Id	Displays the SAP and SDP IDs of the service ID.
Oper State	Displays the operational state of the SAP and SDP IDs of the service ID.
Mrtr Port	Specifies there the port is a multicast router port.
Send Queries	Specifies whether the send-queries command is enabled or disabled.
Max Num Groups	Specifies the maximum number of multicast groups that can be joined on this SAP or SDP.
MVR From VPLS	Specifies MVR from VPLS.
Num Groups	Specifies the actual number of multicast groups that can be joined on this SAP or SDP.

## Sample Output

```

*A:ALA-42>show>service>id>igmp-snooping>snooping# all
=====
IGMP Snooping info for service 100
=====
IGMP Snooping Base info
-----
Admin State : Up
Querier      : 10.20.1.6 on SAP 2/2/1:100
-----
Sap/Sdp      Oper    MRtr  Send    Max Num  MVR      Num
Id           State   Port  Queries Groups   From-VPLS Groups
-----
sap:1/1/4    Up      No    Disabled No Limit Local     1
sap:2/2/1:100 Up      No    Disabled No Limit Local     0
-----
IGMP Snooping Querier info
-----
Sap Id       : 2/2/1:100
IP Address   : 10.20.1.6
Expires      : 3s
Up Time      : 0d 00:15:23
Version      : 3

General Query Interval : 2s
Query Response Interval : 1.0s
Robust Count   : 2
-----
IGMP Snooping Multicast Routers
-----
MRouter      Sap/Sdp Id           Up Time           Expires           Version
-----
10.20.1.6    2/2/1:100           0d 00:15:24      2s               3
-----
Number of mrouters: 1
-----
IGMP Snooping Proxy-reporting DB
-----
Group Address Mode    Up Time           Num Sources
-----
225.0.0.0    exclude 0d 00:00:04      0
-----
Number of groups: 1
-----
IGMP Snooping SAP 1/1/4 Port-DB
-----
Group Address Mode    Type    From-VPLS Up Time           Expires           Num Src
-----
225.0.0.0    exclude dynamic local          0d 00:00:05      4s               0
-----
Number of groups: 1
-----
IGMP Snooping SAP 2/2/1:100 Port-DB
-----
Group Address Mode    Type    From-VPLS Up Time           Expires           Num Src
-----
Number of groups: 0
-----

```

## IGMP Snooping Static Source Groups

## IGMP Snooping Statistics

Message Type	Received	Transmitted	Forwarded
General Queries	463	0	463
Group Queries	0	0	0
Group-Source Queries	0	0	0
V1 Reports	0	0	0
V2 Reports	0	0	0
V3 Reports	4	4	0
V2 Leaves	0	0	0
Unknown Type	0	N/A	0

## Drop Statistics

Bad Length : 0  
 Bad IP Checksum : 0  
 Bad IGMP Checksum : 0  
 Bad Encoding : 0  
 No Router Alert : 0  
 Zero Source IP : 0

Send Query Cfg Drops : 0  
 Import Policy Drops : 0  
 Exceeded Max Num Groups : 0

MVR From VPLS Cfg Drops : 0  
 MVR To SAP Cfg Drops : 0

## IGMP Snooping Multicast VPLS Registration info

IGMP Snooping Admin State : Up

MVR Admin State : Down  
 MVR Policy : None

## Local SAPs/SDPs

Svc Id	Sap/Sdp Id	Oper State	From VPLS	Num Local Groups
100	sap:1/1/4	Up	Local	1
100	sap:2/2/1:100	Up	Local	0

## MVR SAPs (from-vpls=100)

Svc Id	Sap/Sdp Id	Oper State	From VPLS	Num MVR Groups
--------	------------	------------	-----------	----------------

No MVR SAPs found.

\*A:ALA-42>show>service>id>snooping#

mrouters

<b>Syntax</b>	<b>mrouters [detail]</b>
<b>Context</b>	show>service>id>igmp-snooping
<b>Description</b>	Displays all multicast routers.
<b>Parameters</b>	<b>detail</b> — Displays detailed information.
<b>Output</b>	<b>Show igmp-snooping mrouters</b> — The following table describes the show igmp-snooping mrouters output fields:

Label	Description
MRouter	Specifies the multicast router port.
Sap/Sdp Id	Specifies the SAP and SDP ID multicast router ports.
Up Time	Displays the length of time the mrouter has been up.
Expires	Displays the amount of time left before the query interval expires.
Version	Displays the configured version of IGMP running on this interface.

Sample Output

```
*A:ALA-48# show service id 700 igmp-snooping mrouters
=====
IGMP Snooping Multicast Routers for service 700
=====
MRouter      Sap/Sdp Id      Up Time      Expires      Version
-----
Number of mrouters: 0
=====
*A:ALA-48#
```

mvr

<b>Syntax</b>	<b>mvr</b>
<b>Context</b>	show>service>id>igmp-snooping
<b>Description</b>	Displays Multicast VPLS Registration (MVR) information.
<b>Output</b>	<b>Show igmp-snooping mvr</b> — The following table describes the show igmp-snooping mvr output fields:

Label	Description
IGMP Snooping Admin State	Displays the IGMP snooping administrative state.
MVR Admin State	Displays the MVR administrative state.
MVR Policy	Displays the MVR policy name.
Svc ID	Displays the service ID.
Sap/SDP	Displays the SAP/SDP ID.
Oper State	Displays the operational state.
From VPLS	Displays the originating VPLS name.
Num Local Groups	Displays the number of local groups.

### Sample Output

```

A:ALA-1>show>service>id>snooping# mvr
=====
IGMP Snooping Multicast VPLS Registration info for service 10
=====
IGMP Snooping Admin State : Up

MVR Admin State           : Up
MVR Policy                 : mvr-policy
-----
Local SAPs/SDPs
-----
Svc Id      Sap/Sdp      Oper      From      Num Local
            Id           State     VPLS      Groups
-----
100         sap:1/1/10:10        Up        Local     100
100         sap:1/1/10:20        Up        Local     100
-----
MVR SAPs (from-vpls=10)
-----
Svc Id      Sap/Sdp      Oper      From      Num MVR
            Id           State     VPLS      Groups
-----
20          sap:1/1/4:100        Up        10        100
30          sap:1/1/31:10.10    Up        10        100
=====
A:ALA-1>show>service>id>snooping#

```

## port-db

<b>Syntax</b>	<b>port-db sap <i>sap-id</i> [detail]</b> <b>port-db sap <i>sap-id</i> group <i>grp-address</i></b> <b>port-db sdp-<i>id</i>:<i>vc-id</i> [detail]</b> <b>port-db sdp <i>sdp-id</i>:<i>vc-id</i> group <i>grp-address</i></b>
<b>Context</b>	show>service>id>igmp-snooping
<b>Description</b>	This command displays information on the IGMP snooping port database for the VPLS service.
<b>Parameters</b>	<p><b>group <i>grp-ip-address</i></b> — Displays the IGMP snooping port database for a specific multicast group address.</p> <p><b>sap <i>sap-id</i></b> — Displays the IGMP snooping port database for a specific SAP. See <a href="#">Common Service Commands on page 2062</a> for <i>sap-id</i> command syntax.</p> <p><b>sdp <i>sdp-id</i></b> — Displays only IGMP snooping entries associated with the specified mesh SDP or spoke SDP. For a spoke SDP, the VC ID must be specified; for a mesh SDP, the VC ID is optional.</p> <p><b>Values</b> 1 — 17407</p> <p><b><i>vc-id</i></b> — The virtual circuit ID on the SDP ID for which to display information.</p> <p><b>Default</b> For mesh SDPs only, all VC IDs</p> <p><b>Values</b> 1 — 4294967295</p> <p><b>group <i>grp-address</i></b> — Displays IGMP snooping statistics matching the specified group address.</p> <p><b>source <i>ip-address</i></b> — Displays IGMP snooping statistics matching one particular source within the multicast group.</p>

## Sample Output

```
A:ALA-1>show>service>id>snooping# port-db sap 1/1/2
=====
IGMP Snooping SAP 1/1/2 Port-DB for service 10
=====
Group Address      Mode      Type      From-VPLS  Up Time      Expires    Num Src
-----
225.0.0.1          include   dynamic   Local      0d 00:04:44  0s         2
-----
Number of groups: 1
=====

A:ALA-1>show>service>id>snooping# port-db sap 1/1/2 detail
=====
IGMP Snooping SAP 1/1/2 Port-DB for service 10
=====
IGMP Group 225.0.0.1
-----
Mode                : include                Type                : dynamic
Up Time             : 0d 00:04:57             Expires              : 0s
Compat Mode         : IGMP Version 3
V1 Host Expires     : 0s                    V2 Host Expires     : 0s
-----
```

```

Source Address  Up Time      Expires  Type      Fwd/Blk
-----
1.1.1.1.1      0d 00:04:57  20s      dynamic   Fwd
1.1.1.1.2      0d 00:04:57  20s      dynamic   Fwd
-----
Number of groups: 1
=====
A:ALA-1>show>service>id>snooping#

```

## proxy-db

<b>Syntax</b>	<b>proxy-db [detail]</b> <b>proxy-db group <i>grp-address</i></b>
<b>Context</b>	show>service>id>igmp-snooping
<b>Description</b>	Displays information on the IGMP snooping proxy reporting database for the VPLS service.
<b>Parameters</b>	<b>group <i>grp-ip-address</i></b> — Displays the IGMP snooping proxy reporting database for a specific multicast group address.

### Sample Output

```

A:ALA-1>show>service>id>snooping# proxy-db
=====
IGMP Snooping Proxy-reporting DB for service 10
=====
Group Address      Mode      Up Time      Num Sources
-----
225.0.0.1          include   0d 00:05:40    2
-----
Number of groups: 1
=====

A:ALA-1>show>service>id>snooping# proxy-db detail
=====
IGMP Snooping Proxy-reporting DB for service 10
=====
IGMP Group 225.0.0.1
-----
Up Time : 0d 00:05:54          Mode : include
-----
Source Address  Up Time
-----
1.1.1.1.1      0d 00:05:54
1.1.1.1.2      0d 00:05:54
-----
Number of groups: 1
=====
A:ALA-1>show>service>id>snooping#

```

## querier

<b>Syntax</b>	<b>querier</b>
<b>Context</b>	show>service>id>igmp-snooping
<b>Description</b>	Displays information on the IGMP snooping queriers for the VPLS service.

**Sample Output**

```
A:ALA-1>show>service>id>snooping# querier
=====
IGMP Snooping Querier info for service 10
=====
Sap Id           : 1/1/1
IP Address       : 10.10.10.1
Expires          : 6s
Up Time          : 0d 00:56:50
Version          : 3

General Query Interval : 5s
Query Response Interval : 2.0s
Robust Count           : 2
=====
A:ALA-1>show>service>id>snooping#
```

## static

<b>Syntax</b>	<b>static [sap <i>sap-id</i>   sdp <i>sdp-id:vc-id</i>]</b>
<b>Context</b>	show>service>id>igmp-snooping
<b>Description</b>	Displays information on static IGMP snooping source groups for the VPLS service.
<b>Parameters</b>	<p><b>sap <i>sap-id</i></b> — Displays static IGMP snooping source groups for a specific SAP. See <a href="#">Common Service Commands on page 2062</a> for <i>sap-id</i> command syntax.</p> <p><b>sdp <i>sdp-id</i></b> — Displays the IGMP snooping source groups for a specific spoke or mesh SDP.</p> <p><b>Values</b>      1 — 17407</p> <p><b>vc-id</b> — The virtual circuit ID on the SDP ID for which to display information.</p> <p><b>Default</b>      For mesh SDPs only, all VC IDs</p> <p><b>Values</b>      1 — 4294967295</p>

**Sample Output**

```
A:ALA-1>show>service>id>snooping# static
=====
IGMP Snooping Static Source Groups for SAP 1/1/2
-----
Source           Group
-----
*                225.0.0.2
```



```

*                               225.0.0.3
-----
Static (*,G)/(S,G) entries: 2
-----
IGMP Snooping Static Source Groups for SDP 10:10
-----
Source          Group
-----
1.1.1.1         225.0.0.10
-----
Static (*,G)/(S,G) entries: 1
=====
A:ALA-1>show>service>id>snooping#

```

## statistics

- Syntax** `statistics [sap sap-id | sdp sdp-id:vc-id]`
- Context** `show>service>id>igmp-snooping`
- Description** Displays IGMP snooping statistics for the VPLS service.
- Parameters**
- sap *sap-id*** — Displays IGMP snooping statistics for a specific SAP. See [Common Service Commands on page 2062](#) for *sap-id* command syntax.
  - sdp *sdp-id*** — Displays the IGMP snooping statistics for a specific spoke or mesh SDP.
- Values** 1 — 17407
- vc-id*** — The virtual circuit ID on the SDP ID for which to display information.
- Default** For mesh SDPs only, all VC IDs
- Values** 1 — 4294967295

## Sample Output

```

A:ALA-1>show>service>id>snooping# statistics
=====
IGMP Snooping Statistics for service 1
=====
Message Type          Received      Transmitted    Forwarded
-----
General Queries       4             0              4
Group Queries         0             0              0
Group-Source Queries  0             0              0
V1 Reports            0             0              0
V2 Reports            0             0              0
V3 Reports            0             0              0
V2 Leaves             0             0              0
Unknown Type          0             N/A            0
-----
Drop Statistics
-----
Bad Length            : 0
Bad IP Checksum       : 0
Bad IGMP Checksum     : 0

```

## Show Commands

```
Bad Encoding          : 0
No Router Alert       : 0
Zero Source IP        : 0

Send Query Cfg Drops  : 0
Import Policy Drops   : 0
Exceeded Max Num Groups : 0

MVR From VPLS Cfg Drops : 0
MVR To SAP Cfg Drops   : 0
=====
A:ALA-1>show>service>id>snooping#
```

## mc-ecmp-balance

<b>Syntax</b>	<b>mc-ecmp-balance [detail]</b>
<b>Context</b>	show>router>pim
<b>Description</b>	This command displays multicast balance information.
<b>Parameters</b>	<b>detail</b> — Displays detailed information.

### Sample Output

```
A:ALA-48>config>router>pim# show router pim mc-ecmp-balance
=====
PIM ECMP Balance
=====
MC-ECMP-Balance          : Disabled
Rebalance in progress    : No
Last Rebalance Time      : 11/13/2007 09:03:10
Rebalance Type           : Unknown
Optional Threshold Used   : 0
Mc Ecmp Balance Hold Time : None
=====
A:ALA-48>config>router>pim#
```

## mcast-management

<b>Syntax</b>	<b>mcast-management</b>
<b>Context</b>	show
<b>Description</b>	This command shows multicast path management related information.

## bandwidth-policy

<b>Syntax</b>	<b>bandwidth-policy <i>policy-name</i> [detail]</b>
<b>Context</b>	show>mcast-management
<b>Description</b>	This command displays multicast path management bandwidth policy information.
<b>Parameters</b>	<i>policy-name</i> — 32 char max

**Output**

```

Bandwidth Policies : 2
=====
*A:Dut-C# *show mcast-management bandwidth-policy detail*
=====
Bandwidth Policy Details
=====
Policy                : gie
-----
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
-----
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
=====
Bandwidth Policies : 2
=====
*A:Dut-C#

```

# channel

<b>Syntax</b>	<b>channel</b> [ <b>router</b> <i>router-instance</i>   <b>vpls</b> <i>service-id</i> ] [ <b>mda</b> <i>slot[/mda]</i> ] [ <b>group</b> <i>ip-address</i> [ <b>source</b> <i>ip-address</i> ]] [ <b>path</b> <i>path-type</i> ] [ <b>detail</b> ]
<b>Context</b>	show>mcast-management
<b>Description</b>	This command displays multicast path management channel related information.
<b>Parameters</b>	<b>vpls</b> <i>service-id</i> — Specifies an existing VPLS service ID. <b>Values</b> service-id: 1 — 214748364 <i>ip-address</i> — ipv4-address a.b.c.d <i>path-type</i> — Specifies the path type.

<b>Values</b>	primary, secondary, ancillary
<b>Output</b>	<pre>*A:Dut-C# *show mcast-management channel* ===== Multicast Channels ===== Legend : D - Dynamic E - Explicit ===== Source Address                               Slot/Mda   Current Bw Path      D/E Group Address                               Highest Bw ----- 10.10.4.10                                10/2      134646 Ancillary D 225.0.0.0                                134646 ===== Multicast Channels : 1 ===== *A:Dut-C#  *A:Dut-C# *show mcast-management channel detail* ===== Multicast Channels ===== Source Address      : 10.10.4.10 Group Address       : 225.0.0.0 ----- Slot/Mda            : 10/2          Current Bw          : 134646 kbps Dynamic/Explicit    : Dynamic      Current Path         : Ancillary Oper Admin Bw       : 0 kbps       Preference          : 0 Ing last highest    : 134646       Ing sec highest     : 109532 Black-hole rate     : None         Blackhole            : No Time remaining      : 30 seconds ===== Multicast Channels : 1 ===== *A:Dut-C#</pre>

## mda

- Syntax** **mda** [*slot[/mda]*] [*path path-type*]
- Context** show>mcast-management
- Description** This command displays multicast path management MDA related information.
- Parameters** *path-type* — Specifies the path type.

**Values** primary, secondary, ancillary

**Output**

```
*A:Dut-C# *show mcast-management mda 10/2*
=====
MDA 10/2
=====
S/M  Bw-policy                                Type          Limit      In-use-Bw
Admin
-----
10/2  gie                                primary        0 Ms       0 Ms       up
      gie                                secondary      0 Ms       0 Ms       up
      gie                                ancillary      0 Ms       219. 64 Ms up
=====
*A:Dut-C#
```

## group

- Syntax** **group** [*grp-ip-address*]  
**group summary**
- Context** show>router>igmp
- Description** This command displays IGMP group information.

## group-interface

- Syntax** **group-interface** [*fwd-service service-id*] [*ip-int-name*] [*detail*]
- Context** show>router>igmp
- Description** This command displays IGMP group-interface information.

## hosts

- Syntax** **hosts** [*group grp-address*] [*detail*] [*fwd-service service-id*] [*grp-interface ip-int-name*]  
**hosts** [*host ip-address*] [*group grp-address*] [*detail*]  
**hosts summary**
- Context** show>router>igmp

## Show Commands

**Description** This command displays IGMP hosts information.

## interface

**Syntax** **interface** [*ip-int-name*|*ip-address*] [**group**] [*grp-ip-address*] [**detail**]

**Context** show>router>igmp

**Description** This command displays IGMP interface information.

## ssm-translate

**Syntax** **ssm-translate** [*interface-name*]

**Context** show>router>igmp

**Description** This command displays SSM translate configuration information.

## static

**Syntax** **static** [*ip-int-name*|*ip-addr*]

**Context** show>router>igmp

**Description** This command displays IGMP static group/source configuration information.

## statistics

<b>Syntax</b>	<b>statistics</b> [ <i>ip-int-name</i>   <i>ip-address</i> ] <b>statistics host</b> [ <i>ip-address</i> ]
<b>Context</b>	show>router>igmp
<b>Description</b>	This command displaysIGMP statistics information.

## status

<b>Syntax</b>	<b>status</b>
<b>Context</b>	show>router>igmp
<b>Description</b>	This command displays IGMP status information.

## tunnel-interface

<b>Syntax</b>	<b>tunnel-interface</b>
<b>Context</b>	show>router>igmp
<b>Description</b>	This command displaysIGMP tunnel-interface information.

---

## Clear Commands

### id

<b>Syntax</b>	<b>id</b> <i>service-id</i>
<b>Context</b>	clear>service clear>service>statistics
<b>Description</b>	This command clears the identification for a specific service.
<b>Parameters</b>	<i>service-id</i> — The ID that uniquely identifies a service.  <b>Values</b> service-id: 1 — 214748364 svc-name: A string up to 64 characters in length.

### arp-host

<b>Syntax</b>	<b>arp-host</b> <b>arp-host</b> { <b>mac</b> <i>ieee-address</i>   <b>sap</b> <i>sap-id</i>   <b>ip-address</b> <i>ip-address[/mask]</i> } <b>arp-host</b> [ <b>port</b> <i>port-id</i> ] [ <b>inter-dest-id</b> <i>intermediate-destination-id</i>   <b>no-inter-dest-id</b> ] <b>arp-host statistics</b> [ <b>sap</b> <i>sap-id</i>   <b>interface</b> <i>interface-name</i> ]
<b>Context</b>	clear>service>id
<b>Description</b>	This command clears ARP host data.

### authentication

<b>Syntax</b>	<b>authentication</b>
<b>Context</b>	clear>service>id
<b>Description</b>	This command enters the context to clear session authentication information.

### msap

<b>Syntax</b>	<b>msap</b> <i>msap-id</i>
<b>Context</b>	clear>service>id
<b>Description</b>	This command clears Managed SAP information.



<b>Parameters</b>	<i>msap-id</i> — Specifies a Managed SAP ID.		
	<b>Values</b>	dot1q	[port-id   lag-id]:qtag1
		qinq	[port-id   lag-id]:qtag1.qtag2
		qtag1	0 — 4094
		qtag2	0 — 4094

## msap-policy

<b>Syntax</b>	<b>msap-policy</b> <i>msap-policy-name</i>
<b>Context</b>	clear>service>id
<b>Description</b>	This command clears Managed SAPs created by the Managed SAP policy.
<b>Parameters</b>	<i>msap-policy-name</i> — Specifies an existing MSAP policy.

## statistics

<b>Syntax</b>	<b>statistics</b>
<b>Context</b>	clear>service>id>authentication
<b>Description</b>	This command clears session authentication statistics for this service.

## statistics

<b>Syntax</b>	<b>statistics</b>
<b>Context</b>	clear>service
<b>Description</b>	This command clears the statistics for a service.

## subscriber

<b>Syntax</b>	<b>subscriber</b> <i>sub-ident-string</i>
<b>Context</b>	clear>service>statistics
<b>Description</b>	This command clears the statistics for a particular subscriber.
<b>Parameters</b>	<i>sub-ident-string</i> — Clears statistics for the specified subscriber identification string.

fdb

<b>Syntax</b>	<b>fdb</b> { <b>all</b>   <b>mac</b> <i>ieee-address</i>   <b>sap</b> <i>sap-id</i> ]   <b>mesh-sdp</b> <i>sdp-id[:vc-id]</i>   <b>spoke-sdp</b> <i>sdp-id:vc-id</i> }																		
<b>Context</b>	clear>service>id																		
<b>Description</b>	This command clears FDB entries for the service.																		
<b>Parameters</b>	<b>all</b> — Clears all FDB entries.  <b>mac</b> <i>ieee-address</i> — Clears only FDB entries in the FDB table with the specified 48-bit MAC address. The MAC address can be expressed in the form <i>aa:bb:cc:dd:ee:ff</i> or <i>aa-bb-cc-dd-ee-ff</i> where <i>aa</i> , <i>bb</i> , <i>cc</i> , <i>dd</i> , <i>ee</i> and <i>ff</i> are hexadecimal numbers.  <i>sap-id</i> — Specifies the physical port identifier portion of the SAP definition. See <a href="#">Common Service Commands on page 2062</a> for <i>sap-id</i> command syntax.  <b>mesh-sdp</b> — Clears only service FDB entries associated with the specified mesh SDP ID. For a mesh SDP, the VC ID is optional.  <b>spoke-sdp</b> — Clears only service FDB entries associated with the specified spoke SDP ID. For a spoke SDP, the VC ID must be specified.  <i>sdp-id</i> — The SDP ID for which to clear associated FDB entries.  <i>vc-id</i> — The virtual circuit ID on the SDP ID for which to clear associated FDB entries.  <table><tr><td><b>Values</b></td><td><i>sdp-id[:vc-id]</i></td><td><i>sdp-id</i></td><td>1 — 17407</td></tr><tr><td></td><td></td><td><i>vc-id</i></td><td>1 — 4294967295</td></tr><tr><td></td><td><i>sdp-id:vc-id</i></td><td><i>sdp-id</i></td><td>1 — 17407</td></tr><tr><td></td><td></td><td><i>vc-id</i></td><td>1 — 4294967295</td></tr></table>			<b>Values</b>	<i>sdp-id[:vc-id]</i>	<i>sdp-id</i>	1 — 17407			<i>vc-id</i>	1 — 4294967295		<i>sdp-id:vc-id</i>	<i>sdp-id</i>	1 — 17407			<i>vc-id</i>	1 — 4294967295
<b>Values</b>	<i>sdp-id[:vc-id]</i>	<i>sdp-id</i>	1 — 17407																
		<i>vc-id</i>	1 — 4294967295																
	<i>sdp-id:vc-id</i>	<i>sdp-id</i>	1 — 17407																
		<i>vc-id</i>	1 — 4294967295																

mesh-sdp

<b>Syntax</b>	<b>mesh-sdp</b> <i>sdp-id[:vc-id]</i> <b>ingress-vc-label</b>						
<b>Context</b>	clear>service>id						
<b>Description</b>	Clears and resets the mesh SDP bindings for the service.						
<b>Parameters</b>	<i>sdp-id</i> — The mesh SDP ID to be reset.  <table><tr><td><b>Values</b></td><td>1 — 17407</td></tr></table> <i>vc-id</i> — The virtual circuit ID on the SDP ID to be reset.  <table><tr><td><b>Default</b></td><td>All VC IDs on the SDP ID.</td></tr><tr><td><b>Values</b></td><td>1 — 4294967295</td></tr></table>	<b>Values</b>	1 — 17407	<b>Default</b>	All VC IDs on the SDP ID.	<b>Values</b>	1 — 4294967295
<b>Values</b>	1 — 17407						
<b>Default</b>	All VC IDs on the SDP ID.						
<b>Values</b>	1 — 4294967295						

## spoke-sdp

<b>Syntax</b>	<b>spoke-sdp</b> <i>sdp-id:vc-id ingress-vc-label</i>
<b>Context</b>	clear>service>id
<b>Description</b>	Clears and resets the spoke SDP bindings for the service.
<b>Parameters</b>	<i>sdp-id</i> — The spoke SDP ID to be reset. <div><b>Values</b> 1 — 17407</div> <i>vc-id</i> — The virtual circuit ID on the SDP ID to be reset. <div><b>Values</b> 1 — 4294967295</div>

## sap

<b>Syntax</b>	<b>sap</b> <i>sap-id</i> { <b>all</b>   <b>counters</b>   <b>stp</b> }
<b>Context</b>	clear>service>statistics
<b>Description</b>	Clears SAP statistics for a SAP.
<b>Parameters</b>	<i>sap-id</i> — Specifies the physical port identifier portion of the SAP definition. See <a href="#">Common Service Commands on page 2062</a> for <i>sap-id</i> command syntax. <b>all</b> — Clears all SAP queue statistics and STP statistics. <b>counters</b> — Clears all queue statistics associated with the SAP. <b>stp</b> — Clears all STP statistics associated with the SAP.

## sdp

<b>Syntax</b>	<b>sdp</b> <i>sdp-id</i> [ <b>keep-alive</b> ]
<b>Context</b>	clear>service>statistics
<b>Description</b>	Clears keepalive statistics associated with the SDP ID.
<b>Parameters</b>	<i>sdp-id</i> — The SDP ID for which to clear statistics. <div><b>Values</b> 1 — 17407</div> <b>keep-alive</b> — Clears the keepalive history.

## counters

<b>Syntax</b>	<b>counters</b>
<b>Context</b>	clear>service>statistics>id
<b>Description</b>	Clears all traffic queue counters associated with the service ID.

## sap

<b>Syntax</b>	<b>sap sap-id {all   counters   stp}</b>
<b>Context</b>	clear>service>statistics>id
<b>Description</b>	Clears statistics for the SAP bound to the service.
<b>Parameters</b>	<p><i>sap-id</i> — Specifies the SAP ID for which to clear statistics. See <a href="#">Common Service Commands on page 2062</a> for <i>sap-id</i> command syntax.</p> <p><b>all</b> — Clears all queue statistics and STP statistics associated with the SAP.</p> <p><b>counters</b> — Clears all queue statistics associated with the SAP.</p> <p><b>stp</b> — Clears all STP statistics associated with the SAP.</p>

## spoke-sdp

<b>Syntax</b>	<b>spoke-sdp sdp-id[:vc-id] {all   counters   stp}</b>
<b>Context</b>	clear>service>statistics>id
<b>Description</b>	Clears statistics for the spoke SDP bound to the service.
<b>Parameters</b>	<p><i>sdp-id</i> — The spoke SDP ID for which to clear statistics.</p> <p><b>Values</b> 1 — 17407</p> <p><i>vc-id</i> — The virtual circuit ID on the SDP ID to be reset.</p> <p><b>Values</b> 1 — 4294967295</p> <p><b>all</b> — Clears all queue statistics and STP statistics associated with the SDP.</p> <p><b>counters</b> — Clears all queue statistics associated with the SDP.</p> <p><b>stp</b> — Clears all STP statistics associated with the SDP.</p>

## stp

<b>Syntax</b>	<b>stp</b>
<b>Context</b>	clear>service>statistics>id
<b>Description</b>	Clears all spanning tree statistics for the service ID.

## detected-protocols

<b>Syntax</b>	<b>detected-protocols {all   sap <i>sap-id</i>   spoke-sdp <i>sdp-id</i>[:<i>vc-id</i>]}</b>
<b>Context</b>	clear>service>id>stp
<b>Description</b>	RSTP automatically falls back to STP mode when it receives an STP BPDU. The <b>clear detected-protocols</b> command forces the system to revert to the default RSTP mode on the SAP or spoke SDP.
<b>Parameters</b>	<p><b>all</b> — Clears all detected protocol information.</p> <p><i>sap-id</i> — Clears the specified lease state SAP information. See <a href="#">Common Service Commands on page 2062</a> for <i>sap-id</i> command syntax.</p> <p><i>sdp-id</i> — The SDP ID to be cleared.</p> <p><b>Values</b> 1 — 17407</p> <p><i>vc-id</i> — The virtual circuit ID on the SDP ID to be cleared.</p> <p><b>Values</b> 1 — 4294967295</p>

## lease-state

<b>Syntax</b>	<b>lease-state[no-dhcp-release]</b> <b>lease-state [port <i>port-id</i>] [inter-dest-id <i>intermediate-destination-id</i>] [no-dhcp-release]</b> <b>lease-state [port <i>port-id</i>] no-inter-dest-id [no-dhcp-release]</b> <b>lease-state ip-address <i>ip-address</i>[/<i>mask</i>] [no-dhcp-release]</b> <b>lease-state mac <i>ieee-address</i> [no-dhcp-release]</b> <b>lease-state sap <i>sap-id</i> [no-dhcp-release]</b> <b>lease-state sdp <i>sdp-id</i>[:<i>vc-id</i>] [no-dhcp-release]</b>
<b>Context</b>	clear>service>id>dhcp
<b>Description</b>	Clears DHCP lease state information for this service.
<b>Parameters</b>	<p><b>no-dhcp-release</b> — Specifies that the node will clear the state without sending the DHCP release message.</p> <p><b>port <i>port-id</i></b> — The DHCP lease state local specifies that the DHCP lease state is learned by either a SAP or SDP. When the value is SAP, the value indicates the SAP for which this entry contains information.</p>

*ip-address* — The IP address of the IP interface. The *ip-address* portion of the **address** command specifies the IP host address that will be used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).

*ieee-address* — Specifies the 48-bit MAC address for the static ARP in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee, and ff are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

*intermediate-destination-id* — Specifies the intermediate destination identifier which is encoded in the identification strings.

**Values** Up to 32 characters maximum

*sap-id* — Clears the specified lease state SAP information. See [Common Service Commands on page 2062](#) for *sap-id* command syntax.

*sdp-id* — The SDP ID to be cleared.

**Values** 1 — 17407

*vc-id* — The virtual circuit ID on the SDP ID to be cleared.

**Values** 1 — 4294967295

## statistics

<b>Syntax</b>	<b>statistics</b> [ <b>sap</b> <i>sap-id</i>   <b>sdp</b> [ <i>sdp-id</i> [: <i>vc-id</i> ]] ]
<b>Context</b>	clear>service>id>dhcp
<b>Description</b>	Clears DHCP statistics for this service.

## port-db

<b>Syntax</b>	<b>port-db</b> { <b>sap</b> <i>sap-id</i>   <b>sdp</b> <i>sdp-id</i> [: <i>vc-id</i> ]} [ <b>group</b> <i>grp-address</i> [ <b>source</b> <i>ip-address</i> ]]
<b>Context</b>	clear>service>id>igmp-snooping
<b>Description</b>	Clears the information on the IGMP snooping port database for the VPLS service.
<b>Parameters</b>	<p><b>sap</b> <i>sap-id</i> — Clears IGMP snooping statistics matching the specified SAP ID and optional encapsulation value. See <a href="#">Common Service Commands on page 2062</a> for <i>sap-id</i> command syntax.</p> <p><b>sdp</b> <i>sdp-id</i> — Clears only IGMP snooping entries associated with the specified mesh SDP or spoke SDP. For a spoke SDP, the VC ID must be specified; for a mesh SDP, the VC ID is optional.</p> <p><b>Values</b> 1 — 17407</p> <p><i>vc-id</i> — The virtual circuit ID on the SDP ID for which to clear information.</p> <p><b>Default</b> For mesh SDPs only, all VC IDs.</p> <p><b>Values</b> 1 — 4294967295</p> <p><b>group</b> <i>grp-address</i> — Clears IGMP snooping statistics matching the specified group address.</p>

**source** *ip-address* — Clears IGMP snooping statistics matching one particular source within the multicast group.

## querier

<b>Syntax</b>	<b>querier</b>
<b>Context</b>	clear>service>id>igmp-snooping
<b>Description</b>	Clears the information on the IGMP snooping queriers for the VPLS service.

## statistics

<b>Syntax</b>	<b>statistics {all   sap <i>sap-id</i>   sdp <i>sdp-id:vc-id</i>}</b>
<b>Context</b>	clear>service>id>igmp-snooping
<b>Description</b>	Clears IGMP snooping statistics for the VPLS service.
<b>Parameters</b>	<p><b>sap</b> <i>sap-id</i> — Clears IGMP snooping statistics matching the specified SAP ID and optional encapsulation value. See <a href="#">Common Service Commands on page 2062</a> for <i>sap-id</i> command syntax.</p> <p><b>sdp</b> <i>sdp-id</i> — Clears only IGMP snooping entries associated with the specified mesh SDP or spoke SDP. For a spoke SDP, the VC ID must be specified, for a mesh SDP, the VC ID is optional.</p> <p><b>Values</b> 1 — 17407</p> <p><i>vc-id</i> — The virtual circuit ID on the SDP ID for which to clear statistics.</p> <p><b>Default</b> For mesh SDPs only, all VC IDs</p> <p><b>Values</b> 1 — 4294967295</p>

## mfib

<b>Syntax</b>	<b>mfib</b>
<b>Context</b>	clear>service>id
<b>Description</b>	Enter the context to clear multicast FIB info for the VPLS service.

## statistics

<b>Syntax</b>	<b>statistics {all   group <i>grp-address</i>}</b>
<b>Context</b>	clear>service>id>mfib
<b>Description</b>	Clears multicast FIB statistics for the VPLS service.
<b>Parameters</b>	<i>grp-address</i> — Specifies an IGMP multicast group address that receives data on an interface.

## mld-snooping

<b>Syntax</b>	<b>mld-snooping</b>
<b>Context</b>	clear>service>id
<b>Description</b>	This command enables the context to clear MLD snooping-related data.

## port-db

<b>Syntax</b>	<b>port-db sap</b> <i>sap-id</i> [ <b>group</b> <i>grp-ipv6-address</i> ] <b>port-db sap</b> <i>sap-id</i> <b>group</b> <i>grp-ipv6-address</i> <b>source</b> <i>src-ipv6-address</i> <b>port-db sdp</b> <i>sdp-id:vc-id</i> [ <b>group</b> <i>grp-ipv6-address</i> ] <b>port-db sdp</b> <i>sdp-id:vc-id</i> <b>group</b> <i>grp-ipv6-address</i> <b>source</b> <i>src-ipv6-address</i>
<b>Context</b>	clear>service>id>mld-snooping
<b>Description</b>	This command clears MLD snooping port-db group data.

## querier

<b>Syntax</b>	<b>querier</b>
<b>Context</b>	clear>service>id>mld-snooping
<b>Description</b>	This command clears MLD snooping querier information.

## statistics

<b>Syntax</b>	<b>statistics all</b> <b>statistics sap</b> <i>sap-id</i> <b>statistics sdp</b> <i>sdp-id:vc-id</i>
<b>Context</b>	clear>service>id>mld-snooping
<b>Description</b>	This command clears MLD snooping statistics.
<b>Parameters</b>	<b>all</b> — Clears all MLD snooping statistics <b>sap</b> <i>sap-id</i> — Clears all MLD snooping SAP statistics. See <a href="#">Common Service Commands on page 2062</a> for <i>sap-id</i> command syntax. <b>sdp</b> <i>sdp-id:vc-id</i> — Clears all MLD snooping SDP statistics. See <a href="#">Common Service Commands on page 2062</a> for <i>sap-id</i> command syntax.



## arp

<b>Syntax</b>	<b>arp</b> { <b>all</b>   <i>ip-address</i> } <b>arp interface</b> [ <i>ip-int-name</i>   <i>ip-address</i> ]
<b>Context</b>	clear>router
<b>Description</b>	This command clears all or specific ARP entries. The scope of ARP cache entries cleared depends on the command line option(s) specified.
<b>Parameters</b>	<b>all</b> — Clears all ARP cache entries. <i>ip-address</i> — Clears the ARP cache entry for the specified IP address. <b>interface</b> <i>ip-int-name</i> — Clears all ARP cache entries for the IP interface with the specified name. <b>interface</b> <i>ip-addr</i> — Clears all ARP cache entries for the specified IP interface with the specified IP address.

## dhcp

<b>Syntax</b>	<b>dhcp</b>
<b>Context</b>	clear>router
<b>Description</b>	This command enables the context to clear and reset DHCP entities.

## statistics

<b>Syntax</b>	<b>statistics</b> [ <b>interface</b> <i>ip-int-name</i>   <i>ip-address</i> ]
<b>Context</b>	clear>router>dhcp
<b>Description</b>	Clears DHCP statistics.

---

## Debug Commands

### arp-host

<b>Syntax</b>	<b>[no] arp-host</b>
<b>Context</b>	debug>service>id
<b>Description</b>	This command enables and ARP host debugging. The <b>no</b> form of the command disables ARP host debugging

### mld-snooping

<b>Syntax</b>	<b>[no] mld-snooping</b>
<b>Context</b>	debug>service>id
<b>Description</b>	This command enables and configures MLD-snooping debugging. The <b>no</b> form of the command disables MLD-snooping debugging

### detail-level

<b>Syntax</b>	<b>detail-level {low medium high}</b> <b>no detail-level</b>
<b>Context</b>	debug>service>id>mld
<b>Description</b>	This command enables and configures the MLD tracing detail level. The <b>no</b> form of the command disables the MLD tracing detail level.

### mac

<b>Syntax</b>	<b>[no] mac <i>ieee-address</i></b>
<b>Context</b>	debug>service>id>mld
<b>Description</b>	This command shows MLD packets for the specified MAC address. The <b>no</b> form of the command disables the MAC debugging.

## mode

<b>Syntax</b>	<b>mode</b> {dropped-only ingr-and-dropped egr-ingr-and-dropped} <b>no mode</b>
<b>Context</b>	debug>service>id>mld
<b>Description</b>	This command enables and configures the MLD tracing mode. The <b>no</b> form of the command disables the configures the MLD tracing mode.

## sap

<b>Syntax</b>	[no] <b>sap</b> <i>sap-id</i>
<b>Context</b>	debug>service>id>mld
<b>Description</b>	This command shows MLD packets for a specific SAP. The <b>no</b> form of the command disables the debugging for the SAP.

## sdp

<b>Syntax</b>	[no] <b>sdp</b> <i>sdp-id:vc-id</i>
<b>Context</b>	debug>service>id>mld
<b>Description</b>	This command shows MLD packets for a specific SDP. The <b>no</b> form of the command disables the debugging for the SDP.



# DHCP Management

---

## In This Chapter

This chapter provides information about using DHCP, including theory, supported features and configuration process overview.

Topics in this chapter include:

- [DHCP Principles on page 314](#)
- [DHCP Features on page 316](#)
  - [DHCP Relay on page 316](#)
  - [Subscriber Identification Using Option 82 Field on page 319](#)
  - [DHCP Snooping on page 322](#)
  - [DHCP Lease State Table on page 324](#)
  - [DHCPv6 on page 331](#)
  - [DHCP Relay Enhancements on page 335](#)
  - [Virtual Subnet for DHCPv4 Hosts on page 343](#)
- [Proxy DHCP Server on page 344](#)
- [Local address assignment on page 366](#)
- [Configuring DHCP with CLI on page 367](#)

## DHCP Principles

In a Triple Play network, client devices (such as a routed home gateway, a session initiation protocol (SIP) phone or a set-top box) use Dynamic Host Configuration Protocol (DHCP) to dynamically obtain their IP address and other network configuration information.

DHCP is defined and shaped by several RFCs and drafts in the IETF DHC working group including the following:

- RFC 1534, *Interoperation Between DHCP and BOOTP*
- RFC 2131, *Dynamic Host Configuration Protocol*
- RFC 2132, *DHCP Options and BOOTP Vendor Extensions*
- RFC 3046, *DHCP Relay Agent Information Option*

The DHCP operation is illustrated in [Figure 14](#).

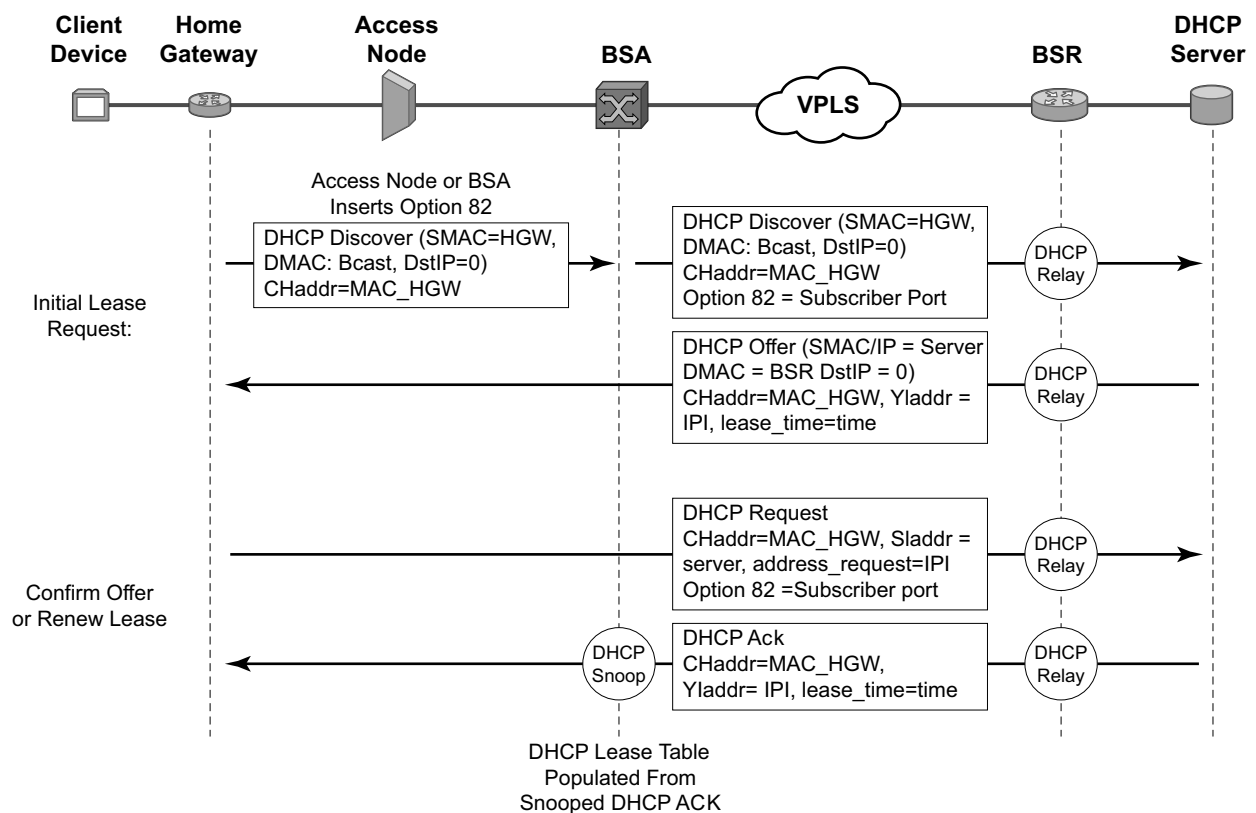


Figure 14: IP Address Assignment with DHCP

05532068

1. During boot-up, the client device sends a DHCP discover message to get an IP address from the DHCP Server. The message contains:
  - Destination MAC address — broadcast
  - Source MAC address — MAC of client device
  - Client hardware address — MAC of client device

If this message passes through a DSLAM or other access node, typically the Relay information option (Option 82) field is added, indicating shelf, slot, port, VPI, VCI etc. to identify the subscriber.

DHCP relay is enabled on the first IP interface in the upstream direction. Depending on the scenario, the DSLAM, BSA or the BSR will relay the discover message as a unicast packet towards the configured DHCP server. DHCP relay is configured to insert the giaddr in order to indicate to the DHCP server in which subnet an address should be allocated.

2. The DHCP server will lookup the client MAC address and Option 82 information in its database. If the client is recognized and authorized to access the network, an IP address will be assigned and a DHCP offer message returned. The BSA or BSR will relay this back to the client device.
3. It is possible that the discover reached more than one DHCP server, and thus that more than one offer was returned. The client selects one of the offered IP addresses and confirms it wants to use this in a DHCP request message, sent as unicast to the DHCP server that offered it.
4. The DHCP server confirms that the IP address is still available, updates its database to indicate it is now in use, and replies with a DHCP ACK message back to the client. The ACK will also contain the Lease Time of the IP address.

## DHCP Features

- [DHCP Relay on page 316](#)
  - [Subscriber Identification Using Option 82 Field on page 319](#)
  - [DHCP Snooping on page 322](#)
  - [DHCP Lease State Table on page 324](#)
- 

## DHCP Relay

Since DHCP requests are broadcast packets that normally will not propagate outside of their IP subnet, a DHCP relay agent intercepts such requests and forwards them as unicast messages to a configured DHCP server.

When forwarding a DHCP message, the relay agent sets the giaddr (gateway IP address) in the packet to the IP address of its ingress interface. This allows DHCP clients to use a DHCP server on a remote network. From both a scalability and a security point of view, it is recommended that the DHCP relay agent is positioned as close as possible to the client terminals.

DHCP relay is practical only in a Layer 3 environment, and thus is only supported in IESservices.

While DHCP relay is not implemented in a VPLS, it is still possible to insert or modify Option 82 information.


---

## DHCP Relay Enhancements

GRT-leaking can be used to relay DHCPv4 and DHCPv6 messages between a VRPN and the Global Routing Table (GRT) for network deployments where DHCP clients and server are in separate routing instances of which one is the Base routing instance.

In network deployments where DHCPv4 client subnets cannot be leaked in the DHCPv4 server routing instance, unicast renewals messages cannot be routed in the DHCPv4 server routing instance as illustrated in [Figure 15](#):





### Figure 15: DHCPv4 Server Routing Instance

With the **relay-unicast-msg** command in the DHCPv4 relay on a regular interface or group-interface, it is possible to configure the gi-address of a DHCPv4 relayed message to any local address that is configured in the same routing instance. Unicast renewals are, in this case, relayed to the DHCPv4 server. In the upstream direction, update the source IP address and add the gateway IP address (gi-address) field before sending the message to the intended DHCP server (the message is not broadcasted to all configured DHCP servers). In the downstream direction, remove the gi-address and update the destination IP address to the value of the yiaddr (your IP address) field. By default, unicast DHCPv4 release messages are forwarded transparently. The optional **release-update-src-ip** flag, updates the source IP address with the value used for relayed DHCPv4 messages.

For retail subscriber interfaces, the **relay-unicast-msg** must be configured at the subscriber-interface dhcp CLI context as shown in the Example 1.

The **relay-unicast-msg** function is not supported in combination with a double DHCPv4 relay (L3 DHCPv4 relay in front of a 7750 DHCPv4 relay with “relay-unicast-msg” enabled).

### Example 1

```
config>service>vpn>

interface "lo0" create
    address 192.168.0.1/32
    loopback
exit
subscriber-interface "sub-int-1" create
    address 10.1.0.254/24
    group-interface "group-int-1-1" create
        dhcp
        server 172.16.1.1
```

## DHCP Relay Enhancements

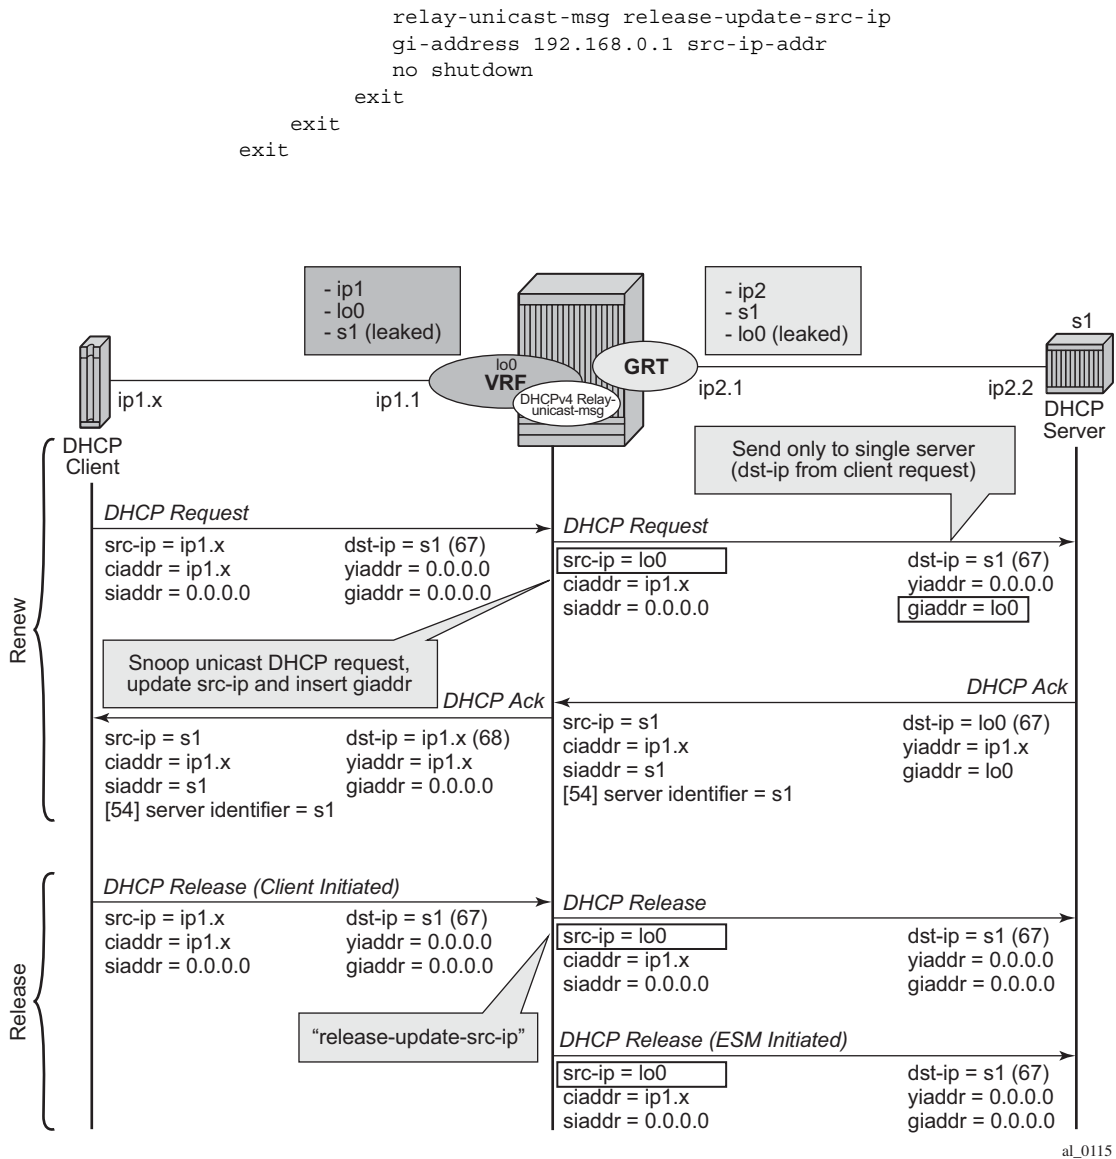


Figure 16: relay-unicast-msg Command in the DHCPv4 Relay

## Subscriber Identification Using Option 82 Field

Option 82, or the relay information option is specified in RFC 3046, *DHCP Relay Agent Information Option*, and allows the router to append some information to the DHCP request that identifies where the original DHCP request came from.

There are two sub-options under Option 82:

- *Agent Circuit ID Sub-option* (RFC 3046, section 3.1): This sub-option specifies data which must be unique to the box that is relaying the circuit.
- *Remote ID Sub-option* (RFC 3046 section 3.2): This sub-option identifies the host at the other end of the circuit. This value must be globally unique.

Both sub-options are supported by the Alcatel-Lucent 7450 ESS and can be used separately or together.

Inserting Option 82 information is supported independently of DHCP relay. However in a VPLS (when the 7450 ESS is not configured for DHCP Relay), DHCP snooping must be enabled on the SAP to be able to insert Option 82 information.

When the circuit id sub-option field is inserted by the 7450 ESS, it can take following values:

- *sap-id*: The SAP index (only under a IES service)
- *ifindex*: The index of the IP interface (only under a IES service)
- *ascii-tuple*: An ASCII-encoded concatenated tuple, consisting of [*system-name*|*service-id*|*interface-name*] (for IES) or [*system-name*|*service-id*|*sap-id*] (for VPLS).
- *vlan-ascii-tuple*: An ASCII-encoded concatenated tuple, consisting of the *ascii-tuple* followed by Dot1p bits and Dot1q tags.

When a DHCP packet is received with Option 82 information already present, the system can do one of three things. The available actions are:

- **Replace** — On ingress the existing information-option is replaced with the information-option parameter configured on the 7450 ESS. On egress (towards the customer) the information option is stripped (per the RFC).
- **Drop** — The DHCP packet is dropped and a counter is incremented.
- **Keep** — The existing information is kept on the packet and the router does not add any additional information. On egress the information option is not stripped and is sent on to the downstream node.

In accordance with the RFC, the default behavior is to keep the existing information; except if the giaddr of the packet received is identical to a local IP address on the router, then the packet is dropped and an error incremented regardless of the configured action.

The maximum packet size for a DHCP relay packet is 1500 bytes. If adding the Option 82 information would cause the packet to exceed this size, the DHCP relay request will be forwarded without the Option 82 information. This packet size limitation exists to ensure that there will be no fragmentation on the end Ethernet segment where the DHCP server attaches.

In the downstream direction, the inserted Option 82 information should not be passed back towards the client (as per RFC 3046, *DHCP Relay Agent Information Option*). To enable downstream stripping of the option 82 field, DHCP snooping should be enabled on the SDP or SAP connected to the DHCP server.

## Trusted and Untrusted

There is a case where the relay agent could receive a request where the downstream node added Option 82 information without also adding a giaddr (giaddr of 0). In this case the default behavior is for the router to drop the DHCP request. This behavior is in line with the RFC.

The 7450 ESS supports a command `trusted`, which allows the router to forward the DHCP request even if it receives one with a giaddr of 0 and Option 82 information attached. This could occur with older access equipment. In this case the relay agent would modify the request's giaddr to be equal to the ingress interface. This only makes sense when the action in the information option is `keep`, and the service is IES. In the case where the Option 82 information gets replaced by the relay agent, either through explicit configuration or the VPLS DHCP Relay case, the original Option 82 information is lost, and the reason for enabling the trusted option is lost.

## DHCP Snooping

This section discusses the Alcatel-Lucent 7450 ESS acting as a Broadband Subscriber Aggregator (BSA) with Layer 2 aggregation towards a Broadband Subscriber Router (BSR).

A typical initial DHCP scenario is:

```
client          server
    ---discover---->
<-----offer-----
    -----request---->
<-----ack-----
```

But, when the client already knows its IP address, it can skip the discover:

```
client          server
    -----request---->
<-----ack-----
```

The BSA can copy packets designated to the standard UDP port for DHCP (port 67) to its control plane for inspection, this process is called DHCP snooping.

DHCP snooping can be performed in two directions:

1. From the client to the DHCP server (Discover or Request messages):
  - to insert Option 82 information (when the system is not configured to do DHCP Relay), see [Subscriber Identification Using Option 82 Field on page 319](#));
  - to forward DHCP requests to a RADIUS server first, and not send them to the DHCP server unless the RADIUS server confirms positive identification. See section [Anti-Spoofing Filters on page 546](#).

For these applications, DHCP snooping must be enabled on the SAP towards the subscriber;

2. From the DHCP server (ACK messages):

- to remove the Option 82 field toward the client
- to build a dynamic DHCP lease state table for security purposes, see section [DHCP Lease State Table on page 324](#)
- to perform Enhanced Subscriber Management, see section [Triple Play Enhanced Subscriber Management on page 729](#)

For these applications, DHCP snooping must be enabled on both the SAP or SDP towards the network and the SAP towards the subscriber.

A major application for DHCP response snooping in the context of Triple Play is security: A malicious user A could send an IP packet (for example, requesting a big video stream) with as source the IP address of user B. Any return packets would be sent to B, and thus potentially jam the connection to B.

As the snooped information is coming straight from the operator's DHCP server, it is considered reliable. The BSA and BSR can use the snooped information to build anti-spoofing filters, populate the ARP table, send ARP replies, etc.

## DHCP Lease State Table

The DHCP lease state table has a central role in the BSA operation (Figure 17). For each SAP on each service it maintains the identities of the hosts that are allowed network access.

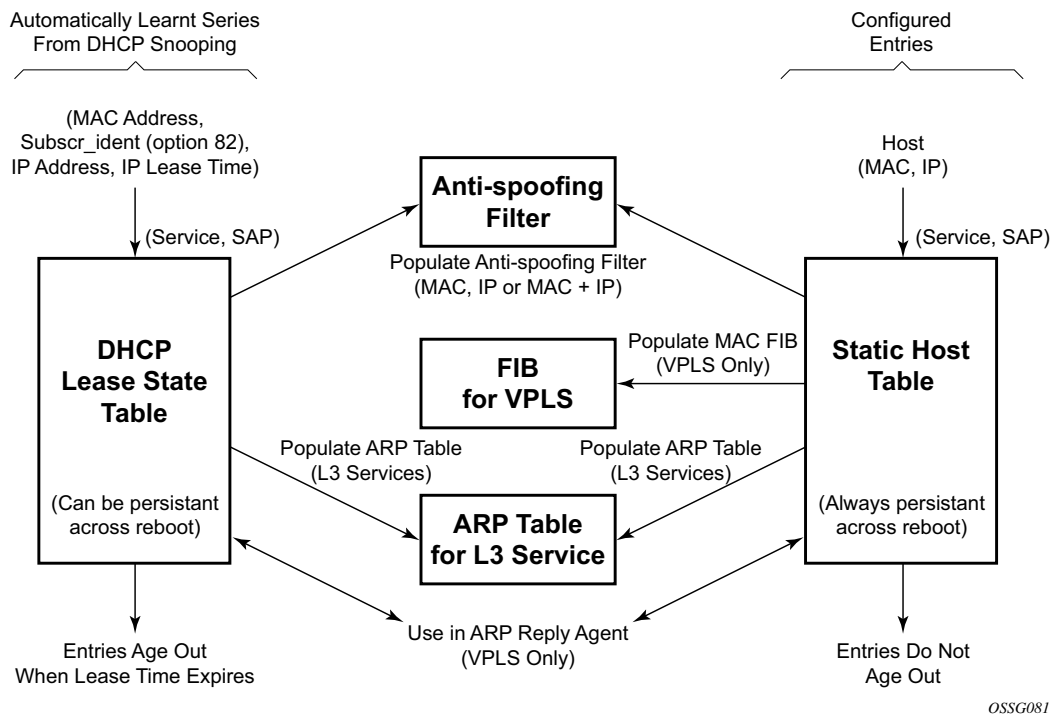


Figure 17: DHCP Lease State Table

When the command **lease-populate** is enabled on a SAP, the DHCP lease state table is populated by snooping DHCP ACK messages on that SAP, as described in the [DHCP Snooping](#) section above.

Entries in the DHCP lease state table remain valid for the duration of the IP address lease. When a lease is renewed, the expiry time is updated. If the lease expires and is not renewed, the entry is removed from the DHCP lease state table.

For VPLS, DHCP snooping must be explicitly enabled (using the **snoop** command) on the SAP and/or SDP where DHCP messages requiring snooping ingress the VPLS instance. For IES interfaces, using the **lease-populate** command also enables DHCP snooping for the subnets defined under the IP interface. Lease state information is extracted from snooped or relayed DHCP ACK messages to populate DHCP lease state table entries for the SAP or IP interface.



For IES services, if ARP populate is configured, no static ARPs are allowed. For IES services, if ARP populate is not configured, then static ARPs are allowed.

The retained DHCP lease state information representing dynamic hosts can be used in a variety of ways:

- To populate a SAP based anti-spoof filter table to provide dynamic anti-spoof filtering — Anti-spoof filtering is only available on VPLS SAPs, or IES IP interfaces terminated on a SAP.
- To populate a VPLS SAP-based arp-reply-agent table to provide dynamic ARP replies using the dynamic hosts IP assigned IP address and learned MAC address — ARP reply agent functionality is only available for static and dynamic hosts associated with a VPLS SAP.
- To populate the system's ARP cache using the **arp-populate** feature — **arp-populate** functionality is only available for static and dynamic hosts associated with IES interfaces.
- To populate managed entries into a VPLS forwarding database — When a dynamic host's MAC address is placed in the DHCP lease state table, it will automatically be populated into the VPLS forwarding database associated with the SAP on which the host is learned. The dynamic host MAC address will override any static MAC entries using the same MAC and prevent learning of the MAC on another interface (implicit MAC Pinning). Existing static MAC entries with the same MAC address as the dynamic host are marked as inactive but not deleted. If all entries in the DHCP lease state associated with the MAC address are removed, the static MAC may be populated. New static MAC definitions for the VPLS instance can be created while a dynamic host exists associated with the static MAC address.
- To support enhanced Subscriber management, see section [RADIUS Authentication of Subscriber Sessions on page 735](#).

Note that if the system is unable to execute any of these tasks, the DHCP ACK message is discarded without adding a new lease state entry or updating an existing lease state entry; and an alarm is generated.

## DHCP and Layer 3 Aggregation

The default mode of operation for DHCP snooping is that the DHCP snooping agent instantiates a DHCP lease state based on information in the DHCP packet, the client IP address and the client hardware address.

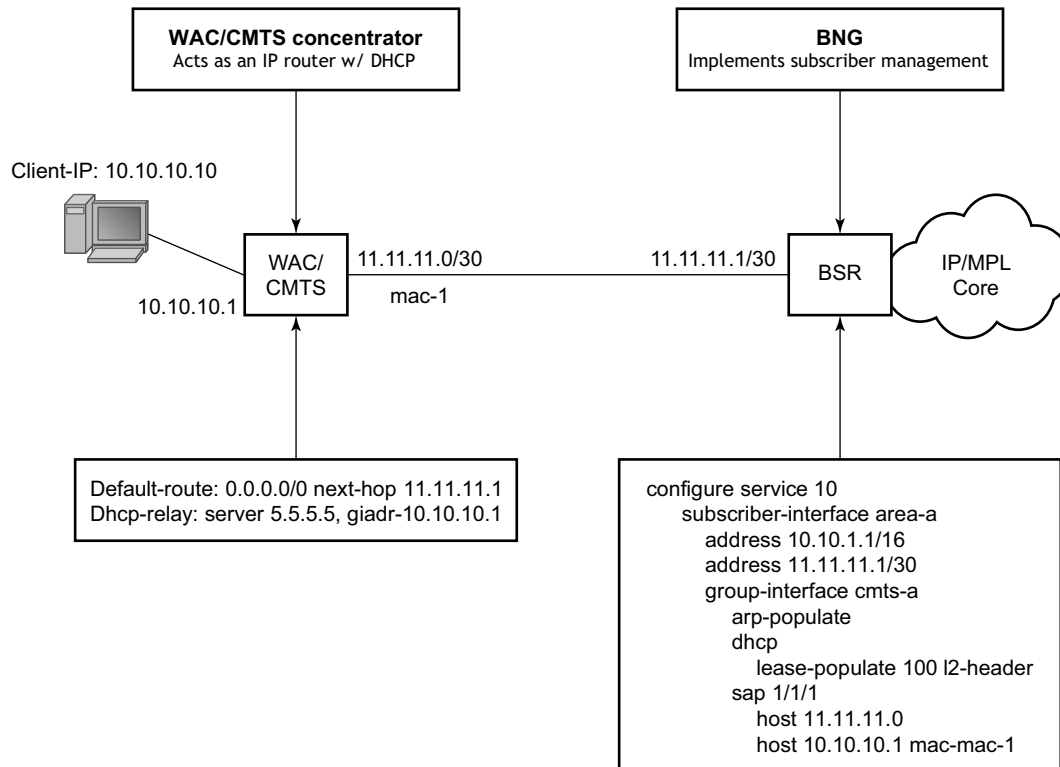
The mode of operation can be changed for DHCP snooping so that the Layer 2 header MAC address is used instead of the client hardware address from the DHCP packet for the DHCP lease state instantiation. This mode is selected by enabling l2-header in the **lease-populate** configuration command at the DHCP level. Because SR-Series routers do not have the ability to verify the DHCP information (both the **src-ip** and **src-mac** of the packet will be those of the previous relay point) anti-spoofing must be performed at the access node prior to the SR-Series routers. This mode provides compatibility with MAC concentrator devices, and cable modem termination system (CMTS) and WiMAX Access Controller (WAC).

A configuration example of a cable/wireless network together with subscriber management is shown in [Figure 18](#). The subnet used to connect to the CMTS/WAC must be defined as a subnet in the subscriber interface of the Layer 3 CO model under which the hosts will be defined. This means that all subscriber lease states instantiated on BSR must be from a “local” subscriber-subnet, even if those are behind the router, as there will be no additional layer 3 route installed pointing to them.

The important items to notice are static hosts at the subscriber interface side:

- IP-only static host pointing to CMTS/WAC WAN link is needed to allow BSR to reply to ARP requests originated from CMTS/WAC.
- IP-MAC static host pointing to CMTS/WAC access-facing interface is required to provide BSR with an arp entry for the DHCP relay address.

When dual-homing is used the CMTS/WAC may be configured with the same MAC for both upstream interfaces. If that is not possible the BSR can be configured with an optional MAC address. The BSR will then use the configured MAC address when instantiating the DHCP lease states.



Fig\_34

Figure 18: CMTS/WAC Network Configuration Example

## Local DHCP Servers

---

### Overview

The local-dhcp-server functions only if there is a relay (gateway) in front of it. Either a GI address is needed to find a subnet or Option 82, which is inserted by the relay, to perform authentication in the local-user-db.

A local DHCP server functions only if there is a relay agent (gateway) in front of it. The local DHCP server must be configured to assign addresses in one of the following ways:

1. Use a local user database authentication (user-db <local-user-db-name>)

The host is matched against the specified local user database. A successful user lookup should return information on one of the following valid addresses:

- A fixed IP address — The IP address should not overlap with the address ranges configured in the local dhcp server.
- A pool name — A free address of any subnet in that pool is offered.
- use-gi-address [scope <subnet | pool>] — The gi address is used to find a matching subnet. When scope is subnet, an address is allocated in the same subnet as the GI address only. When scope is pool, an address is allocated from any subnet within a local pool once that pool has been selected based on matching the “giaddr” field in the DHCP message with any of the configured subnets in the pool.
- use-pool-from-client — The pool name specified in the DHCP client message options and added by the DHCP relay agent is used. A free address of any subnet in that pool is offered.

When no valid address information is returned from the local user database lookup, no IP address is offered to the client.

2. without local user database authentication (no user-db)

One or both address assignment options must be configured:

- Use a pool name (use-pool-from-client)  
The pool name specified in the DHCP client message options and added by the DHCP relay agent is used. A free address of any subnet in that pool is offered.
- Use the gi address (use-gi-address [scope <subnet | pool>])  
The gi address is used to find a matching subnet. When scope is subnet, an address is allocated in the same subnet as the GI address only. When scope is pool, an address is allocated from any subnet within a local pool once that pool has been selected based on matching the “giaddr” field in the DHCP message with any of the configured subnets in the pool.

When both options are configured and a pool name is specified in the DHCP client message options, then the **use-pool-from-client** option has precedence over the **use-gi-address** option.

Note: The local dhcp server will not allocate any address if none of the above options are configured (no user-db, no use-gi-address, no use-pool-from-client).

Options and identification strings can be defined on several levels. In principle, these options are copied into the DHCP reply, but if the same option is defined several times, the following precedence is taken:

1. user-db host options
2. Subnet options
3. Pool options
4. From the client DHCP request.

A local DHCP server must be bound to a specified interface by referencing the server from that interface. The DHCP server will then be addressable by the IP address of that interface. A normal interface or a loopback interface can be used.

A DHCP client is defined by the MAC address and the circuit-id. This implies that for a certain combination of MAC and circuit-id, only one IP-address can be returned. If you ask 1 more, you get same address again. The same address will be returned if a re-request is made.

## Local DHCP Server Support

Local DHCP servers provide a standards-based full DHCP server implementation which allows a service provider the option of decentralizing the IP address management into the network. Local DHCP servers are supported on 7750 SR-7, and 7750 SR-12 models. Note that the 7450 ESS-Series routers use DHCP relay and proxy DHCP server functionality.

Three applications are targeted for the local DHCP server:

- Subscriber aggregation in either a single node (routed CO) or TPSDA.
- PPP clients — Either in a single node or a separate PPPoE server node and a second DHCP server node. The PPPoE server node may be configured to query the DHCP server node for an address and options to provide to the PPPoE client. The PPPoE server will provide the information in PPP format to the client.

DHCP server scenarios include:

- DHCP servers can be integrated with Enhanced Subscriber Management (ESM) for DHCP clients, DHCP relays and PPPoE clients.
- A stand-alone DHCP server can support DHCP clients and DHCP relays.
- IPv4 is supported. DHCP servers provide increased management over the IPv4 address space across its subscriber base, with support for public and private addressing in the same router.

DHCP servers are configurable and can be used in the bridged CO, routed CO, and dual-homing models.

## DHCPv6

In the stateful auto-configuration model, hosts obtain interface addresses and/or configuration information and parameters from a server. Servers maintain a database that keeps track of which addresses have been assigned to which hosts. The stateful auto-configuration protocol allows hosts to obtain addresses, other configuration information or both from a server.

- [DHCPv6 Relay Agent on page 331](#)
  - [DHCPv6 Prefix Options on page 331](#)
  - [Neighbor Resolution via DHCPv6 Relay on page 331](#)
  - [DHCPv6 Lease Persistency on page 332](#)
  - [Local Proxy Neighbor Discovery on page 332](#)
  - [IPv6oE Hosts Behind Bridged CPEs on page 333](#)
- 

### DHCPv6 Relay Agent

When the server unicast option is present in a DHCP message from the server, the option is stripped from the message before sending to the clients.

---

### DHCPv6 Prefix Options

The prefix delegation options described in RFC 3633, *IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6*, provide a mechanism for automated delegation of IPv6 prefixes using DHCP. This mechanism is intended for delegating a long-lived prefix from a delegating router to a requesting router, across an administrative boundary, where the delegating router does not require knowledge about the topology of the links in the network to which the prefixes will be assigned. For example, the delegating router can get a /48 prefix via DHCPv6 and assign /64 prefixes to multiple requesting routers. Prefix delegation is supported for the delegating router (not the requesting router).

---

### Neighbor Resolution via DHCPv6 Relay

This is similar to ARP populate for IPv4. When it is turned on, an SR router needs to populate the link-layer address to IPv6 address mapping into the neighbor database based on the DHCPv6 lease information received.

If the IPv6 address of the host doesn't belong to the subnets of the interface, such a neighbor record should not be created. This could happen when there is a downstream DHCPv6 relay router or prefix delegation requesting router.

## DHCPv6 Lease Persistency

DHCPv6 lease persistency is supported.

The following features are enabled:

- DHCPv6 lease information is reconciled to the standby.
  - DHCPv6 lease information can be stored on a flash card.
  - When rebooted, DHCPv6 lease information stored on a flash card can be used to repopulate the DHCPv6 table as well as the neighbor database if neighbor-resolution is enabled.
- 

## Local Proxy Neighbor Discovery

Local proxy neighbor discovery is similar to local proxy ARP. It is useful in the residential bridging environment where end users are not allowed to talk to each other directly.

When local proxy ND is turned on for an interface, the router:

- Responds to all neighbor solicitation messages received on the interface for IPv6 addresses in the subnet(s) unless disallowed by policy.
- Forwards traffic between hosts in the subnet(s) of the interface.
- Drops traffic between hosts if the link-layer address information for the IPv6 destination has not been learned.



## IPv6oE Hosts Behind Bridged CPEs

This feature adds support for dual-stack IPoE hosts behind bridged clients, receiving globally-routable address using SLAAC or DHCPv6. The feature also provides configurable support for handing out /128 addresses to bridged hosts from same /64 prefix or a unique /64 prefix per host. Bridged hosts that share the same /64 prefix are required to be all SLAAC hosts or DHCPv6 hosts, and are required to be associated with the same SAP. For SLAAC based assignment, downstream neighbor-discovery is automatically enabled to resolve the assigned address.

---

## IPv6 Link-Address Based Pool Selection

This feature provides the capability to select prefix pools for WAN or PD allocations based on configured Link-Address. The scope of selection is the pool or a prefix range within the pool.

---

## IPv6 Address/Prefix Stickiness

This feature extends lease identification criterion beyond DUID (default) for DHCPv6 leases held in the lease database for a configured period after the lease times out. DHCPv6 leases can be held in the lease database for a configurable period of time, after the lease time has expired. A large configured timeout value allows for address and prefix “stickiness”. When a subscriber requests a lease via DHCPv6 (IA\_NA or PD), existing lease is looked-up and handed out. The lease identification match criterion has been extended beyond DUID to also include interface-id or interface-id and link-local address.

---

## IPv4/v6 Linkage for Dual-Stack Hosts or Layer 3 RGs

In case of dual-stack Layer 3 RGs or dual-stack hosts behind Layer 2 CPEs, it is beneficial to have the ability to optionally link IPv6oE host management to DHCP state for v4. This feature provides configurable support to use circuit-id in DHCPv4 option-82 to authenticate DHCPv6. Also, a SLAAC host is created based on DHCPv4 authentication if RADIUS returns IPv6 framed-prefix. IPv6oE host is deleted when the linked IPv4oE host is deleted. In addition, with v4 and v6 linkage configured, sending of unsolicited unicast RA towards the client can be configured when v4 host state is created and IPv6 is configured for the client. The linkage between IPv4 and IPv6 is based on SAP and MAC address. The sharing of circuit-id from DHCPv4 for authentication of DHCPv6 (or SLAAC) allows the 7450 to work around lack of support for LDRA on Access-nodes.

## Host Connectivity Checks for IPv6

This feature provides support to perform SHCV checks on the global unicast address (assigned via SLAAC or DHCPv6 IA\_NA) and link-local address of a L3 RG or a bridged host. SHCV uses IPv6 NS and NA messages. The configuration is similar to IPv4 support in SHCV. The **host-connectivity-check** command is extended to be configured for IPv6 or both IPv4 and IPv6.

## DHCP Relay Enhancements

GRT-leaking can be used to relay DHCPv4 and DHCPv6 messages between a VRPN and the Global Routing Table (GRT) for network deployments where DHCP clients and server are in separate routing instances of which one is the Base routing instance.

In network deployments where DHCPv4 client subnets cannot be leaked in the DHCPv4 server routing instance, unicast renewals messages cannot be routed in the DHCPv4 server routing instance:

With the **relay-unicast-msg** command in the DHCPv4 relay on a regular interface or group-interface, it is possible to configure the gi-address of a DHCPv4 relayed message to any local address that is configured in the same routing instance. Unicast renewals are in this case relayed to the DHCPv4 server. In the upstream direction: update the source IP address and add the gateway IP address (gi-address) field before sending the message to the intended DHCP server (the message is not broadcasted to all configured DHCP servers). In the downstream direction: remove the gi-address and update the destination IP address to the value of the yiaddr (your IP address) field. By default, unicast DHCPv4 release messages are forwarded transparently. The optional **release-update-src-ip** flag, updates the source IP address with the value used for relayed DHCPv4 messages.

For retail subscriber interfaces, the “relay-unicast-msg” must be configured at the subscriber-interface dhcp CLI context.

The relay-unicast-msg function is not supported in combination with a double DHCPv4 relay (L3 DHCPv4 relay in front of a 7750 DHCPv4 relay with relay-unicast-msg enabled).

```
config>service>vprn>
interface "lo0" create
    address 192.168.0.1/32
    loopback
exit
subscriber-interface "sub-int-1" create
    address 10.1.0.254/24
    group-interface "group-int-1-1" create
        dhcp
            server 172.16.1.1
            relay-unicast-msg release-update-src-ip
            gi-address 192.168.0.1 src-ip-addr
            no shutdown
        exit
    exit
exit
```

## Flexible Host Identification in LUDB Based on DHCPv4/v6\_Options

Host identification plays a critical role during the assignment of the parameters to the host via LUDB. The parameters that can be assigned to the subscriber host can range from the IP addressing parameters and the subscriber identification string all the way to the parameters that define the service to which the subscriber is entitled.

LUDB access in the context of IPoE hosts is triggered by DHCP messages passing through the interface on which the LUDB access is configured. This is true regardless of the direction of the DHCP message flow (ingress/egress).

The parameters that define the characteristics of the host will be represented by an LUDB host *entry*. The parameters in the LUDB entry can be unique for each individual host, or they can be shared for a group of hosts. In the former case, the identification field for the LUDB host entry must be host specific while in the latter case the identification field for LUDB host entry could be derived from DHCP options that are common to a set of host.

The host identification in the LUDB can be based on a fixed set of predefined fields within 7x50. If this predefined set of fields is not flexible enough, a custom identification field can be constructed from the DHCP options that are processed by the Python script. Once this custom identifier is constructed, its value can be preserved for the duration of the DHCP transaction and it is used by the LUDB for the host identification.

An example of how this can be used is the following:

A Python script is installed in 7x50. This Python script will intercept incoming DHCP messages on the access side (Discover/Solicit/Request/Renew/Rebind) and it will consequently create a host identification string based on DHCP options in the packet. This string will then be cached and used for host identification in LUDB in both directions (access ingress and network ingress).

This functionality is supported for DHCPv4/DHCPv6 hosts.

---

## DHCP Caching

Subscriber host **identification** via LUDB is performed upon the arrival of the incoming DHCP messages on both, the access and the network side, while the host **instantiation and ESM string assignment** is performed only during the processing of the DHCP ACK/Reply messages. In other words, if Python without the caching is used for subscriber host identification and classification (into proper service class by means of deriving ESM strings), the DHCP options required for host identification must be present in all DHCP messages – even the ones sent by the DHCP servers. However, DHCP servers are not required to echo DHCP options sent by the clients and relay-agents. Consequently, the missing options from the server side would cause the subscriber host instantiation to fail.

To remedy this situation and cover all deployments models (even the ones where the DHCP options are not echoed back by the DHCP servers), a **caching** mechanism is introduced whereby the results of the Python processing on ingress access are locally stored in 7x50. This will ensure that the information about the subscriber host is readily available when the DHCP packet from the DHCP server arrives. Furthermore, since we already have the cached information, no additional Python processing on the network ingress is needed.

The caching is performed in a *DHCP Transaction Cache (DTC)*, which is accessible to Python and to the EMS module. Python will write the result of its processing to it and the Enhanced Subscriber Management (EMS) module within 7x50 will be able to access those results.

The cache entries are relatively short lived, with the lifetime of a *DHCP transaction*. *DHCP transaction* is defined as a pair of DHCP messages that have the same DHCP transaction-id number (<Discover, Offer>, <Request, Ack>, <Solicit, Advertize>, <Request, Reply>, <Renew, Ack>, etc).

---

## Flexible Creation of DHCPv4/6 Host Parameters Utilizing Python and Internal Caching

One of the facilities for flexible creation and assignment of subscriber host parameters is through Python scripting.

There are two models that allow assignment of the subscriber hosts parameters based on the Python processing, one without the utilizing the internal cache (DTC) and the other with the internal cache (DTC).

- Without utilizing the DTC, Python can process options in DHCP ACK message, derive the subscriber host parameters based on those options and consequently insert those parameters in a pre-configured DHCP option (defined in *sub-ident-policy*). The ESM module can be then instructed to extract those parameters and consequently instantiate the host with proper service levels. The drawback of this solution is that the DHCP server may not return all DHCP (v4 and v6) options that clients and relay-agents originally transmitted. Since those options are needed for subscriber parameter determination (but may be absent in the DHCP ACK/REPLY messages when the Python script is run), this solution falls short of covering all deployment cases. In addition, the range of parameters that can be assigned to a subscriber host in this fashion is smaller than the set of parameters utilizing the DTC.
- The internal cache (DTC) allows us to store the result of Python processing. The result is stored during the lifetime of the DHCP transaction. This method of string assignment does not rely on the DHCP server ability to return client's options, DHCPv4 and DHCPv6.

Parameters (ESM strings, IP addresses, etc.) present in the DTC will have priority over any other source that is providing overlapping parameters when it comes to ESM processing. In other words, if the same parameter is provided via DTC (Python), LUDB and Radius, the one provided via

DTC will be in effect. This prioritization will occur automatically without the need for any additional CLI.

For example, if the IPv4 address is provided via DTC during DHCP Discovery processing, then the mode of operation for this host will be proxy-to-dhcp (ESM will terminate DHCP, without going to the server), regardless of whether the IP address is also provided via LUDB or Radius.

This functionality is supported for DHCPv4/DHCPv6 hosts.

---

## Python DTC Variables and API

The following are the Python variables and APIs related to DTC:

### Subscriber Host Identification

**alc.dtc.derivedId** — A read/write (from the Python perspective) string to store the LUDB lookup key for subscriber host identification. This key is derived from the contents of the packet. This string will be used as a *match* criteria in LUDB. The derived-id can only be used when the lookup is performed in ESM. If the LUDB is attached to the local DHCP server, then the lookup based on the derived-id cannot be performed as the DHCP server has no means to derive such an ID from the DHCP message.

### Caching Any Data During the Lifetime of a Transaction

**alc.dtc.store(key,value)** => the operator can store any data he desires in one or more entries. The key can be any arbitrary string (printable ASCII characters), up to 32 bytes in length. The value part is 'unlimited' (memory permitting) in size.

**alc.dtc.retrieve(key)** => retrieve data from the DTC. The key must be an existing key, which is a string consisting of printable ASCII characters, up to 32 bytes in length.

For example, this can be used to cache the DHCP options that the client inserts but the server does not echo back. Those options can still be retrieved in 7x50 via cache in case that their presence is needed for any reason.

The lifespan of the cached data is tied to a DHCP transaction (a pair or corresponding DHCP messages flowing in opposite direction).

### ESM Related Parameters (ESM strings, routing context, etc.)

DTC provides an API to supply a subset of configuration parameters that can otherwise come from RADIUS and LUDB and are used by the ESM code to setup the subscriber host.

DTC parameters as defined below should NOT be considered as DHCP options that can be blindly returned to the DHCP client, but instead they should be considered as real configuration settings. For example, the lease-time option is used in LUDB to enforce the lease time for the client. As such, the ESM keeps state of the lease-time. The following parameters can be used to setup a subscriber host:

**alc.dtc.setESM** (key-from-below, value) => store data that is used by ESM. This data is write-only.

The keys will be predefined (**only** these can be used) and are shown in [Table 8](#). These keys are read-only static variables.

The LUDB column indicates the configuration option under the **config>subscr-mgmt>loc-user-db>ipoe>host** context in LUDB.

**Table 8: ESM-Related Python Variables**

DTC Variable	Type	LUDB	Radius Attribute	Comment
alc.dtc.subIdent	string	identification-strings >subscriber-id	Alc-Subsc-ID-Str	
alc.dtc.subProfileString	string	identification-strings >sub-profile-string	Alc-Subsc-Prof-Str	
alc.dtc.slaProfileString	string	identification-strings >sla-profile-string	Alc-SLA-Prof-Str	
alc.dtc.ancpString	string	identification-strings >ancp-string	Alc-ANCP-Str	
alc.dtc.appProfileString	string	identification-strings >app-profile-string	Alc-App-Prof-Str	
alc.dtc.intDestId	string	identification-strings >inter-dest-id	Alc-Int-Dest-Id-Str	
alc.dtc.catMapString	string	identification-strings >category-map-name	Alc-Credit-Control-CategoryMap	
alc.dtc.ipAddress	string	address	Framed-IPAddress	
alc.dtc.dhcp4DefaultGateway	string	options>default-router	Alc-Default-Router	
alc.dtc.subnetMask	string	address	Framed-IPNetmask	
alc.dtc.ipv4LeaseTime	integer	options>lease-time	Alc-Lease-Time	
alc.dtc.ipv4PrimDns	string	options>dns-server	Alc-Primary-Dns Client-DNS-Pri	
alc.dtc.ipv4SecDns	string	NA	Alc-Secondary-Dns Client-DNS-Sec	
alc.dtc.primNbns	string	options>netbios-name-server	Alc-Primary-Nbns RB-Client-NBNSPri	
alc.dtc.secNbns	string	NA	Alc-Secondary-Nbns RB-Client-NBNSSec	
alc.dtc.msapGroupInterface	string	msap-defaults>group-interface	Alc-MSAP-Interface	

**Table 8: ESM-Related Python Variables (Continued)**

alc.dtc.msapPolicy	string, integer	msap-defaults>policy	Alc-MSAP-Policy	
alc.ServiceId	string, integer	msap-defaults>service	Alc-MSAP-Serv-Id	
alc.dtc.retailServiceId	string	Retail-service-id	Alc-Retail-Serv-Id	
alc.dtc.ipv6Address	string	ipv6-address	Alc-Ipv6-Address	
alc.dtc.ipv6DelegatedP refix	string	ipv6-delegated-prefix	Delegated-IPv6- Prefix	
alc.dtc.ipv6SlaacPrefix	string	ipv6-slaac-prefix	Framed-IPv6- Prefix	
alc.dtc.ipv6WanPool	string	ipv6-slaac-prefix-pool	Framed-IPv6-Pool	
alc.dtc.ipv6PrefixPool	string	ipv6-delegated-prefix-pool	Alc-Delegated- IPv6-Pool	
alc.dtc.ipv6DelegatedP refixLength	integer	ipv6-delegated-prefix-len	Alc-Delegated- IPv6-Prefix- Length	
alc.dtc.accountingPolic y	string	acct-policy	N/A	
alc.dtc.dhcpv4GIAddr	string	gi-address	N/A	
alc.dtc.dhcv4ServerAd dress	string	server	N/A	
alc.dtc.dhcp4Pool	string	address>pool	Framed-Pool Ip-Address-Pool- Name	prim sec (“ ” – delimiter)
alc.dtc.linkAddress	string	link-address	N/A	
alc.dtc.dhcv6ServerAd dr	string	server6	N/A	
alc.dtc.ipv6PrimDns	string	options6>dns-server	Alc-Ipv6-Primary- Dns	
alc.dtc.ipv6SecDns	string	NA	Alc-Ipv6- Secondary-Dns	

For example, an IP address is assigned to a DTC variable as a string:



**alc.dtc.ipAddress = "192.168.0.10"** — This is performed through the following ALU API: **alc.dtc.setESM(alc.dtc.ipAddress,'192.168.0.10')**. The DTC logic then parses this variable and converts it into appropriate format for consumption by ESM code.

The values defined above are the ones that are mostly defined in the LUDB. Main use, however, is assigning ESM strings for the subscriber host instantiation phase during the processing of DHCP ACK/Reply messages. Consequently the Python script needs to be run only on DHCP Request messages (no need to run it on Discoveries for ESM string assignment, unless the LUDB derived-id is also needed).

DHCP options that are blindly returned to the DHCP client without the ESM code being aware of them cannot be configured via DTC. Such options should be configured via Radius (Alc-ToClient-Dhcp-Options – IPv4 only) or they can be inserted directly into DHCP messages via Python (bypassing DTC).

Other possible uses for DTC variables are:

- Assigning routing context information via Python (service-id, msap, msap-policy, retail service-id, etc). For example AN can insert certain 'hints' in various DHCP options that would suggest to us (via Python) in which service context to place the subscriber-host.
- IP address assignment via Python (DTC). This would address the DTC-to-DHCP-Proxy case where Python script is invoked on DHCP Discovery/Solicit. For example:
  - Discover arrives
  - Python generates an IP address, for example based on some DHCP options
  - The script stores the IP address by using `alc.dtc.setEsm(alc.dtc.ipAddress, "10.0.1.2");`
  - After the script is finished, ESM starts processing the packet (no LUDB/RADIUS authentication configured)
  - ESM finds the IP address already in DTC and decides to handle all DHCP and execute proxy function instead of relay
  - ESM sends an offer with the address that python generated
  - DHCP options should be provided as well in this case (lease-times, etc.)
  - The same applies to the DHCP Request

## DTC Debugging Facility

DTC debugging is part of the generic DHCP debugging facility that is enabled by the following commands:

**debug router** *<router-id>* **ip dhcp --** Enable DHCP debug on Layer 3 interfaces, including subscriber-interfaces.

**debug service id** *<service-id>* **dhcp --** Enable DHCP debugging on capture SAP.

If the DTC cache is populated via Python, the corresponding DTC entries will be shown as part of the matching DHCP message debug.

## Virtual Subnet for DHCPv4 Hosts

The **virtual-subnet** command in the **sub-if>dhcp** context allows the system to snoop and record the default router address in DHCP ACK messages for a DHCPv4 ESM host. The system can answer or traceroute request even if the default router address is not configured on the subscriber-interface.

This feature eliminates the need to configure every default-gw address on subscriber interface.

Beside default router address, the system will also calculate host's subnet by using an assigned address and the subnet mask option in ACK. Both recorded default router address and the subnet can be displayed with the **show service id virtual-subnet** command.

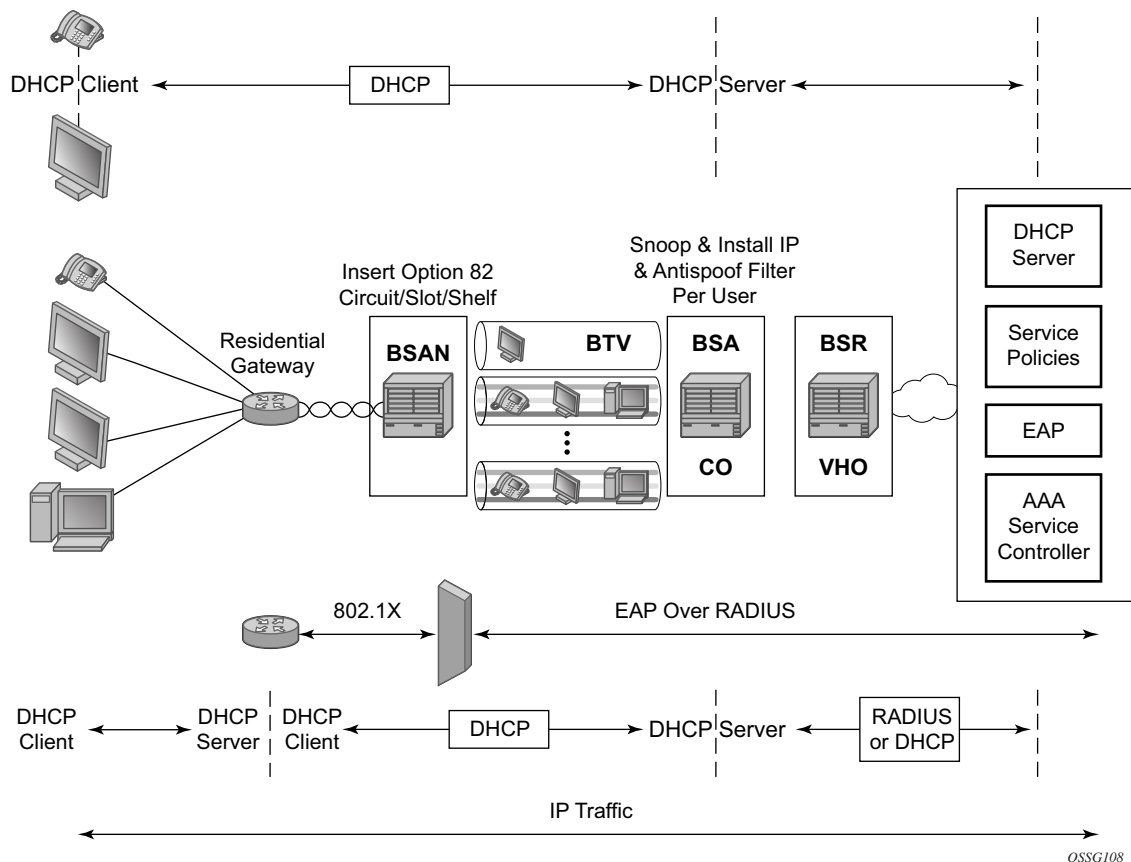
Every ESM subscriber only has one set of default router address and subnet.

## Proxy DHCP Server

This section describes the implementation of a proxy DHCP server capability provide a standards-based DHCP server which will front end to downstream DHCP client and DHCP relay enabled devices and interface with RADIUS to authenticate the IP host and subscriber and will obtain the IP configuration information for DHCP client devices.

The proxy DHCP server is located between an upstream DHCP server and downstream DHCP clients and relay agents when RADIUS is not used to provide client IP information.

Service providers can introduce DHCP into their networks without the need to change back-end subscriber management systems that are typically based around RADIUS (AAA). Service providers can support the use of DHCP servers and RADIUS AAA servers concurrently to provide IP information for subscriber IP devices.



### Figure 19: Typical DHCP Deployment Scenarios

DHCP is the predominant client-to-server based protocol used to request IP addressing and necessary information to allow an IP host device to connect to the network.

By implementing DHCP, the complexity of manually configuring every IP device that requires connectivity to the network is avoided. IP devices with DHCP can dynamically request the appropriate IP information to enable network access.

DHCP defines three components that are implemented in a variety of device types:

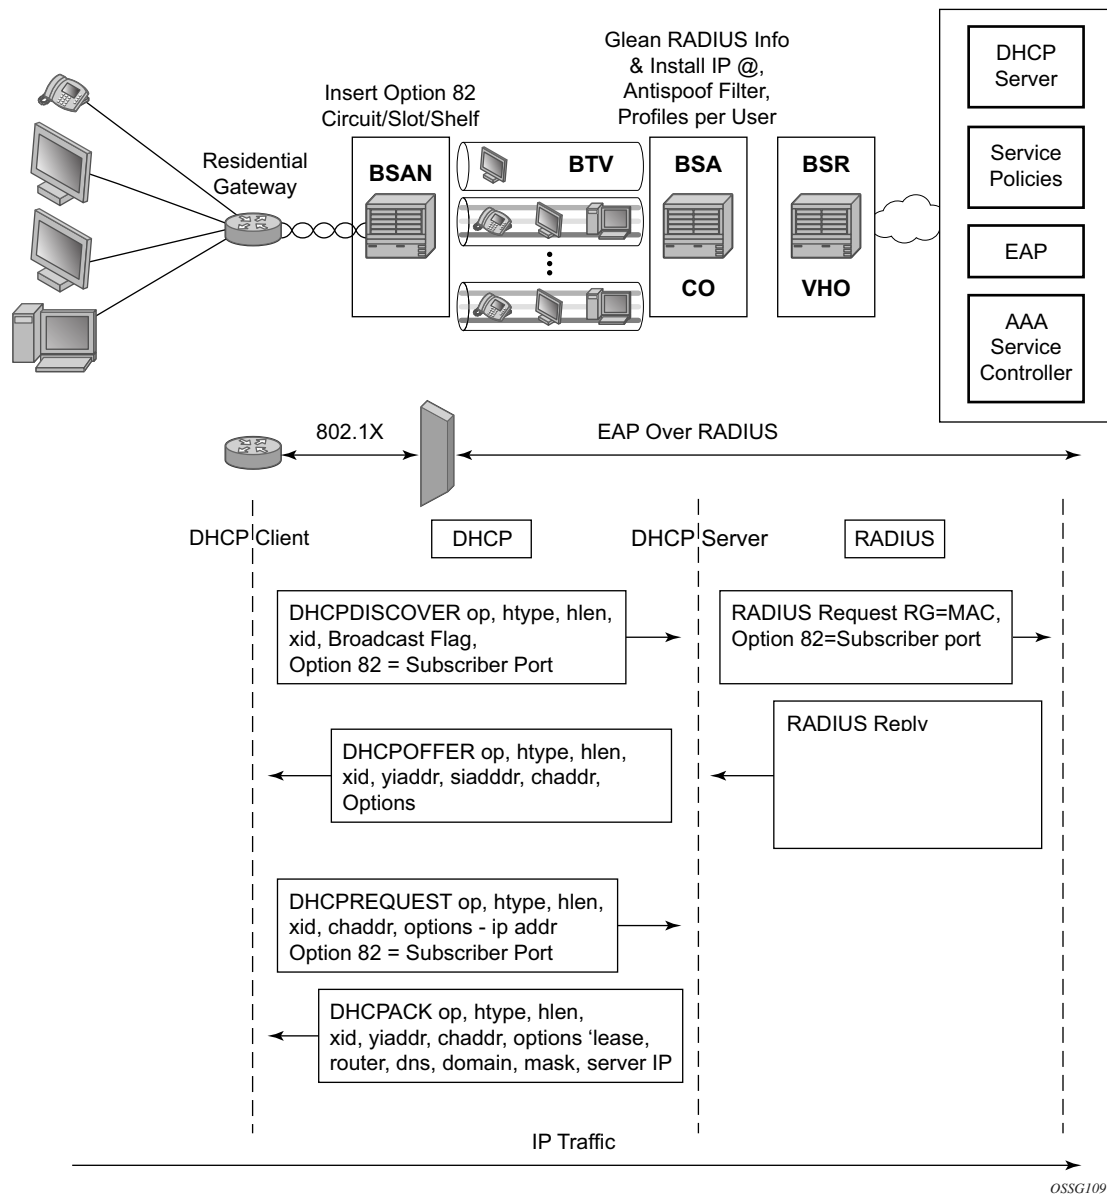
- The DHCP client allows an IP device (host) to request IP addressing information from a DHCP server to enable access to IP based networks. This is typically found in:
  - End user notebooks, desktops and servers
  - Residential gateways and CPE routers
  - IP phones
  - Set-top boxes
  - Wireless access points
- The DHCP relay agent passes (relays) DHCP client messages to pre-configured DHCP servers where a DHCP server is not on the same subnet as the IP host. This feature optionally adds information into DHCP messages (Option 82) which is typically used for identifying attaching IP devices and their location as part of subscriber management. This is typically found in:
  - Residential gateways and CPE routers
  - DSLAMs
  - Edge aggregation routers
- The DHCP server receives DHCP client messages and is responsible for inspecting the information within the messages and determining what IP information if any is to be provided to a DHCP client to allow network access. This is typically found in:
  - Dedicated stand alone servers
  - Residential gateways and CPE routers
  - Edge aggregation routers
  - Centralized management systems

DHCP is the predominant address management protocol in the enterprise community, however in the provider market PPP has traditionally been the means by which individual subscribers are identified, authenticated and provided IP addressing information.

The use of DHCP in the provider market is a growing trend for managing subscriber IP addressing, as well as supporting newer devices such as IP-enabled IP phones and set-top boxes. The majority of subscriber management systems rely heavily on RADIUS (RFC 2865, *Remote Authentication Dial In User Service (RADIUS)*) as the means for identifying and authorizing individual subscribers (and devices), deciding whether they will be allowed access to the network, and which policies should be put in place to control what the subscriber can do within network.

The proxy DHCP server capability enables the deployment of DHCP into a provider network, by acting as a proxy between the downstream DHCP devices and the upstream RADIUS based subscriber management system.

- Interact with downstream DHCP client devices and DHCP relay Agents in the path
- Interface with RADIUS to authenticate DHCP requests
- Receive all the necessary IP information to properly respond to a DHCP client
- Ability to override the allocate IP address lease time, for improved IP address management.



**Figure 20: Example of Triple Play Aggregation Network With DHCP to RADIUS Authentication**

Figure 20 displays a typical DHCP initial boot-up sequence with the addition of RADIUS authentication. The proxy DHCP server will interface with downstream DHCP client devices and then authenticate upstream using RADIUS to a providers subscriber management system.

In addition to granting the authentication of DHCP hosts, the RADIUS server can include RADIUS attributes (standard and vendor-specific attributes (VSAs)) which are then used by the edge router to:

- Provision objects related to a given DHCP host — Subscriber and SLA policy
- Provide IP addressing information to a DHCP client
- Support the features that leverage DHCP lease state
  - Dynamic ARP population
  - ARP reply agent
  - Anti-spoofing filters
  - MAC pinning
- Leverage host-connectivity-verify to determine the state of a downstream IP host

This feature offers the ability for a customer to integrate DHCP to the subscriber while maintaining their existing subscriber management system typically based around RADIUS. This provides the opportunity to control shifts to an all DHCP environment or to deploy a mixed DHCP and RADIUS subscriber management system.

In order to maximize its applicability VSAs of legacy BRAS vendors can be accepted so that a network provider is not forced to reconfigure its RADIUS databases (or at least with minimal changes).

To receive data from the RADIUS server the following are supported:

- Juniper (vendor-id 4874) attributes 4 (Primary DNS server) and 5 (Secondary DNS server)
- Redback (vendor-id 2352) attributes 1 (Primary DNS) and 2 (Secondary DNS).
- Juniper attributes 6 and 7 (Primary and Secondary NetBIOS nameserver).
- Redback attributes 99 and 100 (Primary and Secondary NetBIOS nameserver).

The following attributes can be sent to RADIUS:

- Sending authentication requests: (from the DSL Forum) (vendor-id 3561), attributes 1 (Circuit ID) and 2 (Remote ID).
- DSL Forum attributes 129 and 130 (Actual Data Rate Upstream and Downstream), 131 and 132 (Minimum Data Rate Upstream and Downstream) and 144 (Access Loop Encapsulation).

The complete list of TiMetra VSAs is available on a file included on the compact flash shipped with the image.

## Local DHCP Servers

---

### Terminology

- **Local 7x50 DHCP Server** – DHCP server instantiated on the local 7x50 node.
- **Remote 7x50 DHCP Server** – DHCP server instantiated on the remote 7x50 node (external to the local 7x50 node).
- **3rd party DHCP server** – DHCP server external to any 7x50 node and implemented outside of 7x50.
- **Intercommunication link** – the logical link between dual-homed 7x50 DHCP servers used for synchronizing DHCP lease states. Multi-chassis Synchronization (MCS) protocol runs over this link. Once this link is interrupted, synchronization of the leases between redundant DHCP servers is impaired. This link should be well protected with multiple underlying physical paths.
- **Local IP address-range/prefix** – refers to the local failover mode in which the IP address-range/prefix is configured in dual-homed DHCP environment. The local keyword does not refer to the locality (local vs remote) of the server on which the IP address-range/prefix is configured, but rather refers to the ownership of the IP address-range/prefix. The DHCP server on which the local IP address-range/prefix is configured, owns this IP address-range/prefix and consequently is allowed to delegate the IP addresses/prefixes from it at any time, regardless of the state of the intercommunication link.
- **Remote IP address-range/prefix** – refers to the remote failover mode in which the IP address-range/prefix is configured in dual-homed DHCP environment. The remote keyword does not refer to the locality of the server on which the IP address-range/prefix is configured, but rather refers to the ownership of the IP address-range/prefix. The DHCP server on which the remote IP address-range/prefix is configured, but does not own this IP address-range/prefix during normal operation and consequently is NOT allowed to delegate the IP addresses/prefixes from it. Only when the intercommunication link between the two nodes transition into particular (failed) state, the DHCP server is allowed to start delegating new IP addresses from the remote IP address-range/prefix.
- **IP address-range/prefix ownership** – 7x50 DHCP server can delegate new leases from an IP address-range/prefix that it owns. For example, an IP address-range/prefix designated as remote is not owned by the DHCP server on which it is configured unless certain conditions are met. Those conditions are governed by the state of the intercommunication link.
- **IP address-range/prefix takeover** – a 7x50 DHCP server that does not own an IP address-range/prefix can take over the ownership of this IP address-range/prefix under certain conditions. Once the ownership is taken, the new IP addresses can start being delegated from this IP address-range/prefix. Only the remote IP address-range/prefix can be taken over. Note that the takeover of an IP address-range/prefix has only local significance – in other words, the ownership is not taken away from some other 7x50



DHCP server that has the same IP address-range/prefix designated as local. It only means that IP address-range/prefix that is configured as remote is available to takeover for new IP address delegation.

- **Local PPPoX Address Pools** – this term refers to the method of accessing an IPv4/v6 address pool in 7x50 DHCP4/6 server. For PPPoX clients, the IPv4/v6 addresses are allocated from those pools without the need for an intermediate DHCP relay-agent (7x50 internal DHCP relay-agent). Although those pools are part of the local DHCP server in 7x50, the method of accessing them is substantially different than accessing local DHCP address pools for IPoE (DHCP) clients. IPoE (DHCP) and PPPoX hosts can share the same pool and yet each client type can access them in their own unique way:
  - IPoE client via DHCP messaging
  - PPPoE via internal API calls
- **Local PPPoX Pool Management** – IPv4 address allocation/management for PPPoX clients independent of DHCP process (DHCP lease state). An IPv4 address allocated via local PPPoX Pool Management is tied to the PPPoX session. It is without the need for an 7x50 internal DHCP relay-agent.

### Overview

7x50 DHCP server multi-homing will ensure continuity of IP address/prefix assignment/renewal process in case that an entire 7x50 DHCP server fails or in case of a failure of the active link that connects clients to one of the 7x50 DHCP servers in the access part of the network. DHCP server multi-homing is an integral part of the overall subscriber management multi-chassis protection scheme.

DHCP server multi-homing can be implemented outside of the BNG, without subscriber management enabled. However, in the following text, it is assumed that the subscriber management multi-homing (SRRP/MC-LAG, subscriber synchronization) is deployed along with DHCP server multi-homing.

Although the subscriber synchronization process and the DHCP lease states synchronization process use the same synchronization infrastructure within 7x50 (Multi Chassis Redundancy protocol), they are in essence two separate processes that are not aware of each other. As such, the mechanisms that drive their switchover are different. For example, the mechanism that drives subscriber switchover from one node to the other is driven by the access protection mechanism (SRRP/MC-LAG) while the switchover (or takeover) of the IP address-range/prefixes in a DHCP pool is driven by the state of the intercommunication link over which the leases are synchronized. The failure of an entire node makes those differences irrelevant since the access-link failure coincides with the intercommunication link failure and vice versa. However, link-only failures become critical when it comes to their interpretation by the protection mechanisms (SRRP/MC-LAG/DHCP server multi-homing). Regardless of nature of the failure, overall DHCP server multi-chassis protection scheme must be devised in such fashion that the two 7x50 DHCP servers never allocate the same IP address/prefix to two different clients. Otherwise IP address/prefix duplication will ensue. Unique IP address/prefix allocation is achieved by making only one 7x50 DHCP server responsible for IP address/prefix delegation out of the shared IP address-range/prefix.

There are two basic models for DHCP server dual-homing:

- Shared IP address-range/prefix is designated as local on one 7x50 DHCP server and as remote on the other.

In this case, the dhcp-relays must point to both DHCP servers; the one configured with the local IP address-range/prefix as well as the one with the remote IP address-range/prefix.

Under normal circumstances, the new IP addresses/prefixes can be only allocated from the DHCP server configured with the local IP address-range/prefix.

The DHCP server configured with the remote IP address-range/prefix will start delegating new lease from it only when it declares that the redundant peer with the local IP address-range/prefix becomes unavailable.

Detection of the peer unavailability is triggered by the failure of the intercommunication link which can be caused either by the nodal failure or simply by the loss of connectivity between the two nodes protecting each other. Thus, the loss of intercommunication link does not necessarily mean that the peering node is truly gone. It can simply mean that the

two nodes became isolated and unable to synchronize their DHCP leases between each other. In such environment, both nodes can potentially allocate the same IP address at the same time. To prevent this, additional intercommunication link states and associated timers are introduced to give the chance the operator ample time to fix the problem.

For example, the DHCP server will take over the remote IP address-range/prefix once the MCLT period expires while the intercommunication link is in PARTNER-DOWN state. The PARTNER-DOWN state is entered after a preconfigured timer (partner-down-delay) expires. The consequence of these two additional timers (partner-down-delay and MCLT) is that the new IP address delegation from the remote (shared) IP address-range/prefix will not be possible until the preset timers expire. This is needed and justified in case that the intercommunication link is interrupted, the nodes become isolated and consequently the DHCP lease state synchronization becomes impaired. On the other hand, if the DHCP server with local IP address-range/prefix becomes truly unavailable, those additional restoration times will cause interruption in service since the new IP addresses from the remote IP address-range/prefix will not be immediately available for delegation.

Note that only new IP address delegation from the remote IP address-range/prefix is affected by this behavior. The existing IP leases can be extended on both nodes at any time irrespective of whether the configured address-range/prefix is designated as local or remote.

To ensure uninterrupted service even for new lease delegation in this model (local-remote), two approaches can be adopted:

→ Segment the IP address/prefix space so that each node has an IP address-range/prefix designated as local.

For example, instead of designating IP address-range 10.10.10.0/24 as local on DHCP server A and as remote on DHCP server B, the 10.10.10.0/24 IP address will be split into two: 10.10.10.0/25 and 10.10.10.128/25.

The 10.10.10.0/25 would be designated as local on the DHCP server A and as remote on DHCP server B.

The 10.10.10.128/25 would be designated as remote on the DHCP server A and as local on DHCP server B.

In this fashion, one node will always be available to assign new leases without any overlap.

→ In case that only one shared IP address-range/prefix is deployed, the operator can bypass the timers (partner-down-delay and MCLT) that are put in place in case that DHCP server nodes become isolated. This bypass of the timers can be achieved via configuration. In this case, a safe operation is warranted only if the operator is certain that the intercommunication link failure will be caused by the nodal failure, and not the physical link failure between the two nodes.

- Shared IP address-range/prefix is designated as access-driven on both 7x50 DHCP servers.

In this scenario, the shared IP address-range/prefix is owned by both nodes and the ownership is not driven by the state of the intercommunication link.

To avoid IP address duplication, only one DHCP server at any given time must be responsible for IP address assignment from this shared IP address-range/prefix.

This will be ensured by the access protection mechanism (SRRP/MC-LAG) that will provide a single active path from the clients to the one of the DHCP servers.

In case that clients have access to the same IP address-range/prefix on both DHCP servers at the same time, the IP address duplication may occur.

Consider the following case:

Two DHCP clients send DHCP Discovers in the following fashion:

- DHCP client A sends DHCP Discover to the DHCP server A
- DHCP client B sends DHCP Discover to the DHCP server B
- DHCP server A assigns IP address 10.10.10.10 to the DHCP client A
- DHCP server B assigns IP address 10.10.10.10 to the DHCP client B
- This is a legitimate scenario since the DHCP lease states are not synchronized until the DHCP lease assignment is completed.
- Just before the DHCP ACK is sent to the respective clients from both nodes, the DHCP lease sync messages are exchanged between the peers.
- DHCP servers do not wait for the reply to the sync message before they send the DHCP Ack to the client.
- Once the DHCP syn message is received from the peer, the DHCP server will realize that the IP lease already exist. In this case, the newer IP lease will override the older.
- The result is that clients A and B use the same IP address and consequently the forwarding of the traffic will be impaired.

In access-driven model, the ESM subscriber host must be collocated with the DHCP server. In other words, the DHCP server must be instantiated in redundant BNGs. The dhcp-relays must point to the respective local 7x50 DHCP servers. There must be no cross-referencing of DHCP servers in this model. In addition, the IP address that the DHCP servers are associated with, must be the same on both DHCP servers. This is necessary to ensure uninterrupted service levels once the switchover in the access occurs.

For example:

- DHCP server A on BNG A is associated with the IP address 1.1.1.1 (for example loopback interface A on BNG A)
- DHCP server B on the peering BNG B is associated with the IP address 1.1.1.1 (configured in BNG B under loopback interface B)
- Dhcp-relay on BNG A points to the IP address 1.1.1.1 ? DHCP server in BNG A
- Dhcp-relay on BNG B points to the IP address 1.1.1.1 ? DHCP server in BNG B.

Consider the following when contemplating deployment of the two described models:

- Local-remote model is agnostic of the access protection mechanism. In fact, the access protection mechanism is not needed at all for safe operation.

Fast takeover of the single shared (remote) IP address-range/prefix can be provided only in cases where the operator can guarantee that the intercommunication link failure is caused by the nodal failure (entire DHCP server node becomes unavailable). Fast takeover is provided by bypassing the partner-down-timer and MCLT.

In case that multiple IP address-ranges/prefixes are deployed, bypass of the timers is not needed since the local IP address-ranges/prefixes are available on both nodes.

- Access-driven model allows a single IP address-range/prefix to be shared across the redundant DHCP server nodes. The access to the single DHCP server node from the client side is ensured by the protection mechanism deployed in the access part of the network (SRRP or MC-LAG).

---

## DHCP Lease Synchronization

DHCP server leases are synchronized over Multi-Chassis Synchronization (MCS) protocol. A DHCP lease synchronization message is sent to the peering node just before the DHCP ACK/Reply is sent to the client.

For example, the message flow for DHCPv4 lease establishment is the following:

DHCP Discover ; DHCP client — DHCP server


DHCP Offer ; DHCP server — DHCP client

DHCP Request ; DHCP client — DHCP server

DHCP Sync Message ; DHCP server — via MCS to the peering DHCP server

DHCP Ack ; DHCP server— DHCP client

DHCP server failover mechanism in the local-remote IP address-range/prefix model relies on the detection of the failure of the link over which DHCP states are synchronized (via MCS protocol). This link is normally disjoint from the access links towards the clients. MCS protocol normally runs over a direct link between the two redundant nodes (1) or over backbone links (2) in case that the direct link is not present.



**Figure 21: Redundancy Model**

In the access-driven IP address-range/prefix model, the DHCP server address-ranges/prefixes are not tied to the state of the intercommunication link. Instead, the DHCP server selection for IP address assignment is only governed by the path selected by the path protection mechanism (SRRP/MC-LAG) deployed in the access part of the network.

## Intercommunication Link Failure Detection

7x50 DHCP Server is a client of Multi-chassis Synchronization (MCS) application with 7x50. Once MCS transitions into out-of-sync state, the 7750 DHCP server redundancy assumes that there is a failure in the network. In essence, the DHCP server failure in dual-chassis configuration relies on the failure detection mechanism of MCS.

MCS runs over TCP, port 45067 and it is using either data traffic or keepalives to detect failure on the communication link between the two nodes. In the absence of any MCS data traffic for more than 0.5sec, MCS will send its own keepalive to the peer. If a reply is NOT received within 3sec,

MCS will declare its operational state as DOWN and the DB Sync state as out-of-sync. MCS will consequently notify its clients (DHCP Server being one of them) of this condition.

In essence, it can take up to 3 seconds before the DHCP client realizes that the interclasses communication link failed.

MCS Clients (applications) can optimally send their own proprietary keepalive messages to its partner over MCS to detect failure. DHCP Server does not utilize this method and it strictly relies on the failure notifications by MCS.

Note that the intercommunication link failure does not necessarily assume the same failed fate for the access links. In other words, it is perfectly possible (although unlikely) that both access links are operational while the inter-chassis communication link is broken.

The failure detection of the intercommunication link will lead to certain failover state transitions on the DHCP server. The DHCP lease handling in the local-remote model will depend on the particular failover state on the DHCP server and the duration of each failover state is determined by preconfigured timers.

---

## DHCP Server Failover States

DHCP Server when paired in redundant fashion can transition through several states:

- TRANSITION
- SHUTDOWN
- INIT
- STARTUP
- NORMAL – In this state, the 7x50 DHCP Server is serving all IP leases from the local and access-driven IP addresses-ranges/prefixes (assigning new leases and extending existing ones). Remote IP addresses-ranges/prefixes are not served (new or existing ones).
- COMMUNICATION INTERRUPTED – IP addresses/prefixes under the local and access-driven IP address-range/prefix are served in the same fashion as in the NORMAL state. The IP addresses/prefixes under the remote IP address-range/prefix will be renewed. However, the new address/prefix from the remote address-range/prefix will not be allocated until the partner-down timer expires, the failover state consequently transitions into PARTNER DOWN and the MCLT timer while in the PARTNER-DOWN state expires. This is necessary in case that the failure occurred only on the intercommunication link while the access link is still operational (DHCP server nodes become isolated).  
  
COMMUNICATION INTERRUPTED state indicates that there is a failure of some kind, but it cannot be determined whether an entire node failed or only the inter-chassis link. The access layer may still be fully operational, but the new leases (including RENEWS/REBINDS) cannot be synchronized between the two peers.
- PARTNER DOWN – Once the DHCP Server reaches this state, the remote IP address-

range/prefix is taken over and after an additional time period of MCLT (Maximum Client Lead Time), the new IP addresses/prefixes from it can be delegated to clients. PARTNER-DOWN state is an indication (and assumption at the same time) that the remote node is truly down.

Otherwise, the IP address duplication may occur in case that all of the following conditions are met:

- the DHCP server nodes become isolated
- the failover state is PARTNER-DOWN
- DHCP/PPPoX clients have simultaneous access to the same IP address-range/prefix on the both DHCP servers.

---

## Lease Time Synchronization

7x50 DHCP server state synchronization is different from the lease state synchronization of the subscriber host itself.

- Each subscriber host lease state is synchronized via sub-mgmt client application. The host lease state can be seen via the output of following CLI command:

**show service id *id* dhcp lease-state**

The output of this command represents the state of our internal DHCP relay (client).

- On the other hand, the DHCP server lease states can be observed via the following CLI command:

**show routed id dhcp local-dhcp-server *name* lease**

The concern is with the latter, the 7x50 DHCP server lease state synchronization. To ensure uninterrupted IP lease renewal process after a failure in the network, the DHCP server lease time that is synchronized between the 7x50 DHCP servers must always lead the currently assigned lease time for the period of the anticipated lease time in the next period.

For comparison purposes the two flow diagrams are juxtaposed in [Figure 22](#):

- The right side does not include any lead time during synchronization
- The left side includes the lead time during synchronization



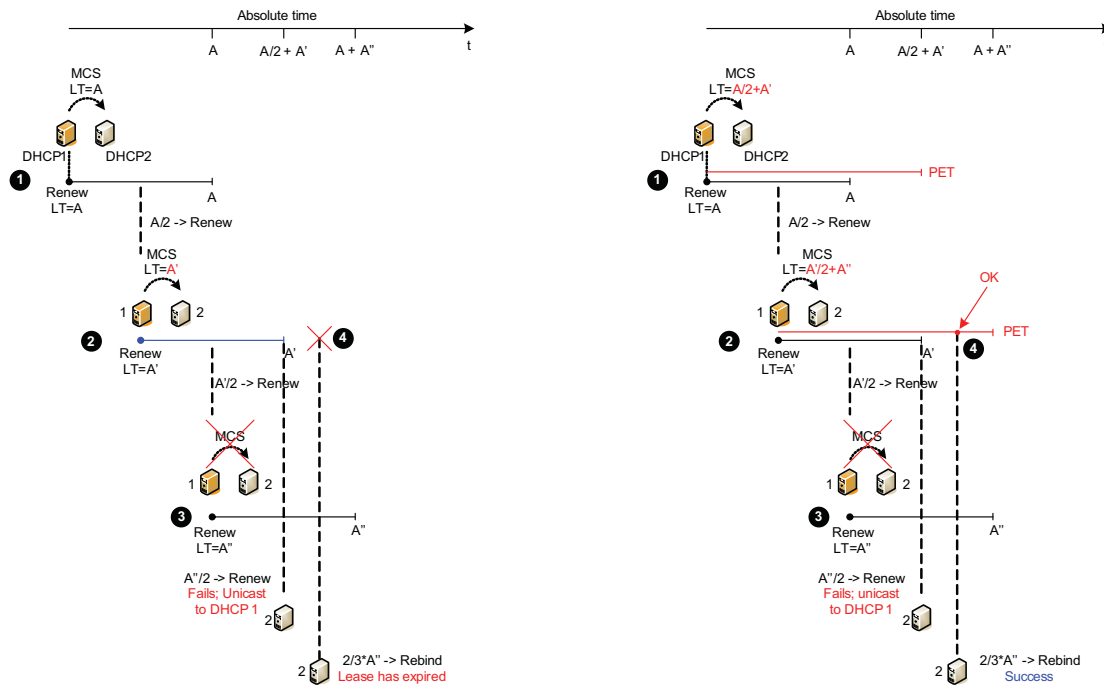


Figure 22: Potential Expiration Time

On the left side of the graph, the lease time is synchronized to the same value to which it was renewed.

In point 3, the primary DHCP server renews the lease time but fails to synchronize it due to its own failure (crash). The secondary server (2) has the old lease time  $A'$  in its database. The next time the client tries to REBIND its address/prefix, the lease in the secondary server will be expired (4). As a result, the IP address/prefix will not be renewed.

To remedy this situation, the primary server must synchronize the lease time as the current RENEW time + the next lease-time. In this fashion, as depicted in point 3, when the REBIND reaches server 2, the lease in its DB will still be active (4) and the server 2 will be able to extend the lease for the client.

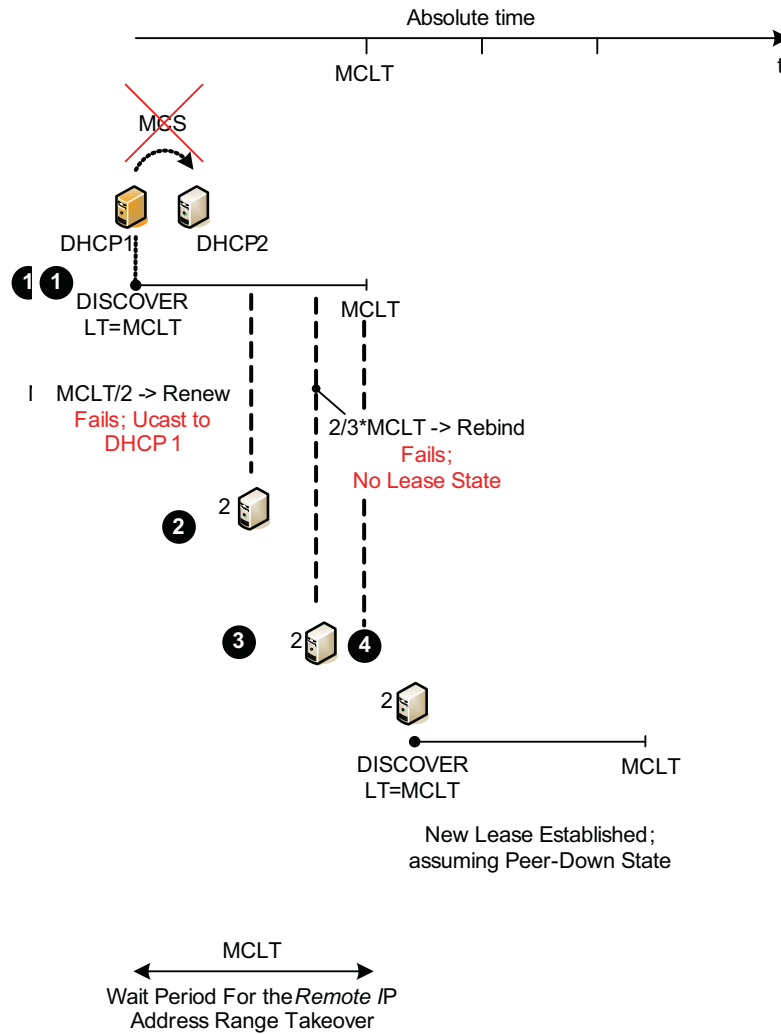
## Maximum Client Lead Time (MCLT)

Maximum Client Lead Time (MCLT) is the maximum time that the 7x50 DHCP Server can extend the lease time to its clients beyond the lease time currently known by the 7x50 DHCP partner node. By default this time is relatively short (10min).

The purpose of the MCLT is explained in the following scenario:

The local 7x50 DHCP server assigns a new IP lease to the client but it crashes before it sends a sync update to the partner server. Because of the local 7x50 DHCP server failure, the remote 7x50 DHCP server is not aware of the IP address/prefix that has been allocated on the local 7x50 DHCP server. This condition creates the possibility that the remote 7x50 DHCP server allocates the same address/prefix to another client. This would cause IP address/prefix duplication. MCLT is put in place to prevent this scenario.

MCLT based solution is shown in [Figure 23](#).



**Figure 23: Address/Prefix Allocation without Synchronization**

The sequence of events is the following:

1. 7x50 DHCP Server 1 is the local DHCP server (with the local address-range/prefix) that creates the IP lease state for a new client. The initial lease-time assigned to the client is MCLT which is normally shorter than the requested lease time.  
This 7x50 DHCP server fails before it gets a chance to synchronize the lease state with the 7x50 DHCP Server 2 (remote 7x50 DHCP server with the remote address-range/prefix).
2. The remote 7x50 DHCP server transitions into the PARTNER-DOWN state (assuming that the partner-down timer is 0). In this state the remote 7x50 DHCP server can extend the lease time to the existing clients but it can NOT assign a new lease for a period of MCLT. In MCLT/2 a new RENEW request is sent directly to the local 7x50 DHCP server. This server is DOWN and therefore it cannot reply.
3. The client broadcasts a REBIND request that reaches the remote 7x50 DHCP server. The remote 7x50 DHCP server has no knowledge of the requested lease and therefore it does not reply.
4. The lease for the client will expire and the client will have to reinitiate the IP address/prefix assignment process.

Since the remote 7x50 DHCP server is not aware of the lease state that was assigned by the local 7x50 DHCP server, there is a chance that the remote 7x50 DHCP server assigns to the new client the same IP address/prefix already allocated by the local 7x50 DHCP server just before it crashed. This is why the remote 7x50 DHCP server needs to wait for the MCLT time to expire so that the IP addresses/prefixes allocated (but never synchronized) by the local 7x50 DHCP server can time out.

When the communication channel between the chassis is interrupted, two scenarios are possible:

1. The entire node becomes unavailable. In this case the redundant node takes over and it starts reducing the lease time until the lease time reaches MCLT.
  - COMMUNICATION-INTERUPTED  
The remote 7x50 DHCP server only renews the leases but does not delegate new ones (primary is DOWN).  
  
The local 7x50 DHCP server renews the leases (which eventually trickle down to MCLT) and delegates the new ones with the lease time of MCLT (secondary is down).
  - PARTNER-DOWN  
The remote 7x50 DHCP server starts delegating new IP addresses/prefixes from the remote address-range/prefix after MCLT (primary is down). The lease time of the new clients is MCLT. A lease cannot be assigned for a period longer than what is agreed with the peer incremented with the MCLT. As for a “new” lease nothing is agreed yet so the sum falls back to the MCLT itself.

2. The communication channel is down but the remote 7x50 DHCP server is not (meaning that the clients have still access to both servers). The behavior in this case is following:

→ COMMUNICATION-INTERRUPTED

The remote 7x50 DHCP server keeps renewing existing leases but it does not delegate the new leases. The chances that this will happen are low as the clients will keep sending RENEWS via unicast to the local 7x50 DHCP server which is still active. The non-synched leases in the remote 7x50 DHCP server will time out. The local 7x50 DHCP server will start trickling down the lease time to MCLT.

The local 7x50 DHCP server will keep delegating the new leases, although with the MCLT lease time.

→ PARTNER-DOWN State

The local 7x50 DHCP server will keep extending the existing leases but it will also start delegating new IP leases once the initial MCLT elapses.

Note that both local and remote 7x50 DHCP server will delegate new leases in PARTNER-DOWN state (although the remote 7x50 DHCP server in PARTNER-DOWN state will have to wait an additional MCLT period before it can start delegating new leases).

---

## Sharing IPv4 Address-Range or IPv6 Prefix Between Redundant 7x50 DHCP Servers in Access-Driven Mode

Access-driven DHCP server redundancy model will ensure uninterrupted IP address assignment service from a single IP address-range/prefix in case that an access link forwards BNG fails. To avoid duplicate address allocation, there MUST be a single active path available from the clients to only one of the 7x50 DHCP servers in redundant configuration. This single active path is ensured via a protection mechanism in dual-homed environment in the access side of the network. The supported protection mechanisms in the access in 7x50 are SRRP or MC-LAG.

In access-driven DHCP redundancy model, the DHCP relay in each 7x50 node must point only to the IP address of the local DHCP server. In other words, the DHCP messages received on one DHCP server should never be relayed to the other. Since the IP address-ranges/prefixes are shared between the DHCP servers, accessing both DHCP servers with the same DHCP request can cause DHCP lease duplication. Moreover, the IP addresses of both DHCP servers must be the same in both nodes. Otherwise the DHCP renew process would fail.

Granting new leases out of the shared IP address-ranges or prefixes that are configured as access-driven is not dependent on the state of the inter-chassis communication link (MCS). Instead, the new leases can be granted from both nodes simultaneously and it is the role of the protection mechanism in the access to ensure that a single path to either server is always active.

This model will allow the newly active node, after a SRRP/MC-LAG switchover, to be able to serve new clients immediately from the same (shared) IP range or prefix. At the same time, upon the switchover, the corresponding subscriber-interface route is re-evaluated for advertisement with

a higher routing metric to the network side via SRRP awareness. The end result is that by aligning the subscriber-interface routes or prefixes with access-driven DHCP address ranges or prefixes in DHCP server, the IP address ranges or prefixes will be advertised to the network side only from the actively serving node (SRRP Master or active MC-LAG node). This will be performed indirectly via the corresponding subscriber-interface route that is aligned (by configuration) with the DHCP address range or prefix in access-driven mode.

In case that SRRP or MC-LAG is not deployed in conjunction with the access-driven configuration option, the IP address duplication could occur.

Note that Multi-chassis redundancy relies on MCS to synchronize various client applications. (subscriber states, DHCP states, IGMP states, etc) between the two chassis. Therefore the links over which MCS peering session is established must be highly redundant. Failed MCS peering session will render dual-chassis redundancy non-operational. Considering this fact, the DHCP failover scenario with SRRP/MC-LAG and shared IP address range should be evaluated considering the following cases:

**Table 9: DHCP Failover Scenarios**

Access Related Failure	Inter-Chassis Link State	Number of Failures	Action
None	None	None	Only the MASTER SRRP BNG will grant IP leases. The subnet tied to the SRRP instance is advertised to the network side on the MASTER SRRP node.
None	COMM-INT	Possibly multiple	Only the MASTER SRRP BNG will grant IP leases. Communication link between the two chassis has failed and the DHCP states cannot be synchronized. Operator is required to restore the links between the chassis.
None	PARTNER-DOWN	Possibly multiple	Same as above. The premise of this deployment model in general (SRRP/MC-LAG + access-driven) is that there is only one path leading to the DHCP server active. This path is governed by SRRP. Therefore there is no change in behavior between the PARTNER-DOWN and COMM-INT states in this particular scenario.


**Table 9: DHCP Failover Scenarios (Continued)**

<b>Access Related Failure</b>	<b>Inter-Chassis Link State</b>	<b>Number of Failures</b>	<b>Action</b>
Access Link Towards the Master SRRP	NORMAL	Single	SRRP will switch over. The new IP lease grants will continue from new SRRP Master using the same IP address range. IP leases will be synched OK. Subnets will be advertised from new Master via SRRP awareness.
Access Link Towards the Master SRRP	COMM-INT	Multiple	<p>SRRP will switch over. The new leases can be handed from the DHCP server on the newly SRRP Master node. However, this DHCP server may not have its lease state table up to date since the inter-chassis communication link is non-operational.</p> <p>Consequently, new SRRP Master node may start handing out leases that are already allocated on the node with the failed link. In this case, IP address duplication would ensue.</p> <p>This is why it of utmost importance that the intercommunication chassis link is well protected and the only event that causes it to go down is when the entire node goes down. Otherwise the nodes becomes isolated from each other and synchronization becomes non-operational.</p>
Access Link Towards the Master SRRP	PARTNER-DOWN	Multiple	Same as above.
Access Link Towards the Standby SRRP	NORMAL	Single	No effect on the operation since everything is active on the Master SRRP node anyway.
Access Link Towards the Standby SRRP	COMM-INT	Multiple	Intercommunication link is broken. The Master SRRP node will continue handing out the new leases and renewing the old ones. However, they will not be synchronized to the peering node.
Access Link Towards the Standby SRRP	PARTNER-DOWN	Multiple	Same as above.

**Table 9: DHCP Failover Scenarios (Continued)**

Access Related Failure	Inter-Chassis Link State	Number of Failures	Action
Entire Master Node	COMM-INT	Single	SRRP will switch over. However, lease duplication may occur on the newly Master since the intercommunication link is broken and this newly SRRP Master is not aware of DHCP leases that the peer (failed node) may have allocated to the clients while the intercommunication-link was broken.
Entire Master Node	PARTNER-DOWN	Single	Same as above.
Entire Standby Node	COMM-INT	Single	Operation continues OK but multi-chassis redundancy is lost.
Entire Standby Node	PARTNER-DOWN	Single	Same as above.

A possible deployment scenario is shown in [Figure 24](#).

**Figure 24: Failover Scenario with SRRP and DHCP in Access-Driven Mode**

## Fast-Switchover of IP Address/Prefix Delegation For Remote IP Address/Prefix Range

Some deployments require that the remote IP address/prefix range starts delegating new IP addresses/prefixes upon the failure of the intercommunication link, without waiting for the intercommunication link to transition from the COMM-INT state into the PARTNER-DOWN state and the MCLT to expire while in PARTNER-DOWN state.

In other words, the takeover of the remote IP address-range/prefix should follow the failure of the intercommunication link, without any significant delays.

This can be achieved by configuring both of the following two items under the dhcp failover CLI hierarchy:

- The partner-down-delay must be set to 0. This will cause the intercommunication link to bypass the COMM-INT state upon the failure and transition straight into the PARTNER-DOWN state. The remote IP address-range/prefix can be taken over only in PARTNER-DOWN state, once the MCLT expires.
- The ignore-mclt-on-takeover flag must be enabled. With this flag enabled, remote IP address/prefix can be taken over immediately upon entering the PARTNER-DOWN state of the intercommunication link, without having to wait for the MCLT to expire. Note that by setting this flag, the lease times of the existing DHCP clients, while the intercommunication link is in the PARTNER-DOWN state, will still be reduced to the MCLT over time and all new lease times will be set to MCLT. This behavior remains the same as originally intended for MCLT. This functionality must be exercised with caution. One needs to keep in mind that the partner-down-delay and MCLT timers were originally introduced to prevent IP address duplication in cases where DHCP redundant nodes transition out-of-sync due to the failure of intercommunication link. These timers (partner-down-delay and MCLT) would ensure that during their duration, the new IP addresses/prefixes are delegated only from one node – the one with local IP address-range/prefix. The drawback is of course that the new IP address delegation is delayed and thus service is impacted.

But if one could ensure that the intercommunication link is always available, then the DHCP nodes would stay in sync and the two timers would not be needed. This is why it is of utmost importance that in this mode of operation, the intercommunication link is well protected by providing multiple paths between the two DHCP nodes. The only event that should cause intercommunication link to fail is the entire nodal failure. This failure is acceptable since in this case only one DHCP node is available to provide new IP addresses/prefixes.



## DHCP Server Synchronization and Local PPPoX Pools

Since there is no classical DHCP lease state maintained for local PPPoX pools, the IP addresses will not be synchronized via DHCP Server. Instead they will be synchronized via PPPoX clients. Once the PPPoX subscriber is synchronized, the respective IP address lease will be updated in the respective local pool.

For example:

- A PPPoE client is created on 7x50 'A'.
- The IP address is assigned from the local pool on 7x50 'A'.
- The PPPoE client will be synchronized to the peering node 7x50 'B'.
- Once the client is synchronized in 7x50 'B', the IP address assignment will be synchronized by our internal PPPoE process on 7x50 'B' with the local pool.

One artifact of this behavior (IP address assignment in local DHCP pools is synchronized via PPPoX clients and not via DHCP server synchronization mechanism) is that during the node boot, the DHCP server must wait for the completion of PPPoX subscriber synchronization via MCS so that it learns which addresses/prefixes are already allocated on the peering node. Since the DHCP server can theoretically start assigning IP addresses before the PPPoX sync is completed, a duplicate address assignment may occur. For example an IP address lease can be granted via DHCP local pools while PPPoX sync is still in progress. Once the PPPoX sync is completed, the DHCP server may discover that the granted IP lease has already been allocated by the peering node. The most recent lease will be kept and the other will be removed from both systems. To prevent this scenario, a configurable timer can be set to an arbitrary value that will render sub-if non-operational until the timer expires. The purpose of this timer is to allow the PPPoX sync to complete before subscribers under the sub-intf can be served.

## Local address assignment

---

### Stateless Address Auto-configuration

In the stateless auto-configuration model, hosts can be assigned address statically or dynamically. For static prefix assignment, LUDB and Radius can be used. For dynamic assignment, a pool name returned from LUDB or Radius and the local DHCPv6 server is used for address management. Although, the DHCPv6 server is used there are no lease time associated with the SLAAC prefix assigned to hosts. To utilize the local pool for SLAAC prefix assignment, the command **local-address-assignment** is used under group-interface. The client-application type **ppp-slaac** and/or **ipoe-slaac** must first be specified. Afterwards, the server name of the local-address-server must also be provisioned.

## Configuring DHCP with CLI

This section provides information to configure DHCP using the command line interface.

- [Common Configuration Tasks on page 368](#)
  - [Enabling DHCP Snooping on page 368](#)
  - [Configuring Option 82 Handling on page 370](#)
  - [Enabling DHCP Relay on page 371](#)
  - [Configuring Local User Database Parameters on page 372](#)

## Common Configuration Tasks

Topics in this section are:

- [Enabling DHCP Snooping on page 368](#)
  - [Configuring Option 82 Handling on page 370](#)
  - [Enabling DHCP Relay on page 371](#)
  - [Configuring Local User Database Parameters on page 372](#)
- 

### Enabling DHCP Snooping

DHCP snooping is the process of copying DHCP packets and using the contained information for internal purposes. The BSA and BSR can use the snooped DHCP information to build anti-spoofing filters, populate the ARP table, send ARP replies, etc.

For VPLS, DHCP snooping must be explicitly enabled (using the **snoop** command) on the SAP or SDP where DHCP messages ingress the VPLS instance. It is recommended to enable snooping on both the interface to the DHCP server (to snoop ACK messages) and the interface to the subscriber (to snoop RELEASE messages).

For IES interfaces, lease-populate enables DHCP snooping for the subnets defined under the IP interface. The number of allowed simultaneous DHCP sessions on a SAP or interface can be limited using the lease-populate command with the parameter number-of-entries specified. Enabling lease-populate and snoop commands is effectively enabling “standard subscriber management” as described in [Standard and Enhanced Subscriber Management on page 773](#).

The following output displays an example of a partial BSA configuration with DHCP snooping enabled in a service:

```
*A:ALA-48>config>service# info
-----
...
    vpls 600 customer 701 create
        sap 1/1/4:100 split-horizon-group "DSL-group2" create
            description "SAP towards subscriber"
            dhcp
                lease-populate 1
                option
                    action replace
                    circuit-id
                    no remote-id
                exit
            no shutdown
        exit
    exit
    mesh-sdp 2:800 create
        dhcp
            snoop
        exit
    exit
    no shutdown
exit
...
-----
*A:ALA-48>config>service#
```

## Configuring Option 82 Handling

Option 82, or “Relay Information Option” is a field in DHCP messages used to identify the subscriber. The Option 82 field can already be filled in when a DHCP message is received at the router, or it can be empty. If the field is empty, the router should add identifying information (circuit ID, remote ID or both). If the field is not empty, the router can decide to replace it.

The following example displays an example of a partial BSA configuration with Option 82 adding on a VPLS service. Note that snooping must be enabled explicitly on a SAP.

```
A:ALA-1>config>service>vpls#
-----
      no shutdown
      description "Default tls description for service id 1"
      sap 1/1/11 split-horizon-group "2dslam" create
        dhcp
          no description
          snoop
          no lease-populate
          option
            action replace
            circuit-id ascii-tuple
            no remote-id
          exit
        no shutdown
      exit
    exit
  -----
A:ALA-1>config>service>vpls#
```

## Enabling DHCP Relay

Note that lease populate and DHCP relay are different features in which are not both required to be enabled at the same time. DHCP relay can be performed without populating lease tables.

The following example displays DHCP relay configured on an IES interface:

```
A:ALA-48>config>service>ies>if# info
-----
        address 10.10.42.41/24
        local-proxy-arp
        proxy-arp
            policy-statement "ProxyARP"
        exit
        sap 1/1/7:0 create
            anti-spoof ip
        exit
        arp-populate
        dhcp
            description "relay_ISP1"
            server 10.200.10.10 10.200.10.20
            lease-populate 1
            no shutdown
        exit
-----
A:ALA-48>config>service>ies>if#
```

## Configuring Local User Database Parameters

A local user data base defines a collection of hosts. There are 2 types of hosts: PPPoE and DHCP. A local user database can be used for the following:

- Perform authentication for PPPoE clients. For this only the hosts declared under PPPoE are used.
- Perform authentication and address management for the local DHCP server. For this both PPPoE and DHCP sections can be used depending on the client type indicated by a vendor-specific suboption inside Option 82 of the DHCP message.

Each host can be identified by a set of values. However, at any point in time only four of these values are taken into account for DHCP as defined by the **dhcp match-list** option and only three are considered for PPPoE as defined in the **pppoe match-list** option.

When trying to find a matching host, attempts are made to match as many items as possible. If several hosts match an incoming DHCP packet, the one with most match criteria is taken.

One host entry can map on several physical clients. For instance, when using a circuit ID, by masking when the interface-id is used, the host-entry is used for all the clients on that same interface.

DHCP host identification, called from the local DHCP server, includes:

- Circuit ID from OPTION 82. Note that for this field there is the possibility to mask the circuit ID (the **mask** command) before looking for the host.
- MAC address
- Remote ID from Option 82
- Option 60 from DHCP message, note that only first 32 bytes are looked at
- SAP ID from vendor-specific suboption of Option 82
- Service ID from vendor-specific suboption of Option 82
- String from vendor-specific suboption of Option 82
- System ID from vendor-specific suboption of Option 82

PPPOE host identification, called from the local DHCP server or from PPPoE host identification includes:

- Circuit ID
- MAC address
- Remote id
- User name, either complete user name, domain part only, or host part only

When a host cannot be inserted in the lookup database, it will be placed in an unmatched-hosts list. This can occur due to:



- Another host with the same host-identification exists. Note that only the host-identification that is specified in the match-list is taken into account for this.
- A host has no host-identification specified in the match-list.

When used for PPPOE-authentication, the fields are used as follows:

- password — Verifies the PPPoE user password. This is mandatory. If no password is required then it must be explicitly set to **ignore**.
- address:
  - no address — No address information. The address must be obtained by other means, either radius or DHCP-server.
  - gi-address — No meaning in this context. The address must be obtained by other means, either RADIUS or DHCP-server.
  - use-pool-from-client — No meaning in this context, address must be obtained by other means, either RADIUS or DHCP-server.
  - pool-name — The address must be obtained by other means, either RADIUS or a DHCP-server. When a DHCP server is used, this pool-name will be included in Option 82 vendor-specific suboption.
  - ip-address — This ip-address will be offered to the client.
- Identification-strings — Returns the strings used for enhanced subscriber management (ESM).
- Options — Only DNS servers and NBNS server are used, others are ignored.

When used from DHCP-server following applies:

- password — not used.
- address — Defines how the address should be allocated for this host.
  - no address — The host is not allowed. The clients mapping to this host will not get an IP address.
  - gi-address — Finds the matching subnet and an IP address is taken from that subnet.
  - pool-name — A free IP address is taken from that pool.
  - ip-address — This ip-address will be offered to the client.
  - use-pool-from-client — Use the poolname in the Option 82 vendor-specific suboption. If no poolname is provided there, falls back to the DHCP server default (none or use-gi-address).
- identification-strings — The operator can specify subscriber management strings and in which option the strings are sent back in dhcp-offer and dhcp-ack messages.
- options — The operator defines which options specific to this host should be sent back in the dhcp-offer and dhcp-ack messages. Note that the options defined here override options defined on the pool-level and subnet-level inside the local DHCP server.

The circuit ID from PPPoE or from Option 82 in DHCP messages can be masked in following ways:

## Configuring Local User Database Parameters

- **prefix-length** — Drop a fixed number of bytes at the beginning of the circuit-id.
- **suffix-length** — Drop a fixed number of bytes at the end of the circuit-id.
- **prefix-string** — The matching string will be dropped from the beginning of the circuit-id. The matching string can contain wildcards (\*). For example: incoming circuit-id "mybox|3|my\_interface|1/1/1:22" masked with "\*|\*|" will leave "my\_interface|1/1/1:22".
- **suffix-string** — The matching string will be dropped at the end of the circuit-id. For example: incoming circuit-id "mybox|3|my\_interface|1/1/1:22" masked with "|\*" will result in "mybox|3|my\_interface".

The following is an example of a local user database used for PPPoE authentication:

```
*A:ALA-48>config>subscr-mgmt# info
-----
...
    local-user-db "pppoe user db"
        description "pppoe authentication data base"
        ppp
            match-list username circuit-id
            mask prefix-string "*|*|" suffix-string "|*"
            host "john" create
                host-identification
                    username "john" no-domain
                exit
                password pap "23T8yPoe0w1R.BPGHB98i0qhJf7ZlZGCtXBKGnjrIrA" hash2
                no shutdown
            exit
            host "test.com" create
                host-identification
                    username "test.com" domain-only
                exit
                password ignore
                no shutdown
            exit
            host "john@test.com" create
                host-identification
                    username "john@test.com"
                exit
                password pap "23T8yPoe0w0TlflyCb4hskknvTYLqA2avvBB567g3eQ" hash2
                identification-strings 122 create
                    subscriber-id "john@test.com"
                    sla-profile-string "sla prof1"
                    sub-profile-string "subscr profile 1"
                    ancp-string "ancp string"
                    inter-dest-id "inter dest"
                exit
                no shutdown
            exit
            host "john@test.com on interface group-if"
                host-identification
                    circuit-id string "group-if"
                    username "john@test.com"
                exit
                password pap "23T8yPoe0w1R.BPGHB98i0qhJf7ZlZGCtXBKGnjrIrA" hash2
                address 10.1.2.3
```

```

        no shutdown
    exit
exit
no shutdown
exit
...
-----
*A:ALA-48>config>subscr-mgmt#

```

The following are some examples when a user tries to set up PPPoE:

- john@test.com tries to setup PPPoE with circuit-id "pe\_23|3|group-if|1/1/1": host "john@test.com on interface group-if" will match, the PAP password is checked and the IP address 10.1.2.3 is given to PPPoE to use for this host.
- john@test.com (on another interface): host "john@test.com" will match, the PAP password is checked, and identification strings are returned to PPPoE.
- alcatel@test.com: host "test.com" will match, no password check, the user is allowed.
- john@alcatel.com: host "john" will match and the password will be checked.
- anybody@anydomain: will not match and will not be allowed.

The following is an example of a local user database used for DHCP server for DHCP clients:

```

*A:ALA-50>config>subscr-mgmt# info
-----
...
    local-user-db "dhcp server user db"
        description "dhcp server user data base"
        dhcp
            match-list circuit-id mac
            mask prefix-string ".*|*" suffix-string ".*"
            host "mac 3 on interface" create
                host-identification
                    circuit-id string "group-if"
                    mac 00:00:00:00:00:03
                exit
                address 10.0.0.1
                no shutdown
            exit
            host "maskedCircId" create
                host-identification
                    circuit-id string "group-if"
                exit
                address pool "pool 1"
                identification-strings 122 create
                    subscriber-id "subscriber 1234"
                    sla-profile-string "sla prof 1"
                    sub-profile-string "sub prof 1"
                    ancp-string "ancpstring"
                    inter-dest-id "inter dest id 123"
                exit
                options
                    netbios-name-server 1.2.3.4
                    lease-time min 2
                exit

```

## Configuring Local User Database Parameters

```
                no shutdown
            exit
        exit
        no shutdown
    exit
...
-----
*A:ALA-50>config>subscr-mgmt#
```

The following is an access example:

- MAC 00:00:00:00:00:03 on circuit-id "pe5|3|group-if|1/1/1": host "mac 3 on interface" is matched and address 10.0.0.1 is offered to the DHCP client.
- Another MAC on circuit-id "pe5|3|group-if|2/2/2": host "maskedCircId" is matched and an address is taken from "pool1" (defined in the DHCP server). The identification-strings will be copied to Option 122 in the dhcp-offer and dhcp-ack messages. The options defined here will also be copied into dhcp-offer and dhcp-ack messages.
- The circuit-id "pe5|3|other\_group\_if|1/1/3": no host is matched. The client will only get an IP address if on DHCP server level you defined the *use-gi-address* parameter and the gi-address matches a subnet.
- 

The following is an example of a local user database used for a DHCP server, only for PPPoE clients:

If PPPoE does not get an IP address from RADIUS or the local-user-db used for authentication, the internal dhcp-client will be used to access a DHCP server which can be in the same node or in another node. These request are identified by inserting Option 82 suboption client-id in the dhcp-discover and dhcp-request messages. When the DHCP server receives this request and has a user-db connected to it, then the PPPoE section of that user-db is accessed.

```
*A:ALA-60>config>subscr-mgmt# info
-----
...
    local-user-db "pppoe user db"
        description "pppoe authentication data base"
        ppp
            match-list username
            host "internet.be" create
                host-identification
                    username "internet.com" domain-only
                exit
                address "pool_1"
                no shutdown
            exit
            host "john@internet.com" create
                host-identification
                    username "john@internet.com"
                exit
                identification-strings 122 create
                    subscriber-id "john@test.com"
                    sla-profile-string "sla prof1"
                    sub-profile-string "subscr profile 1"
```

```

        ancp-string "ancp string"
        inter-dest-id "inter dest"
    exit
    address use-gi
    no shutdown
exit
host "malicious@internet.com"
    host-identification
        circuit-id string "group-if"
        username "internet@test.com"
    exit
    no shutdown
exit
exit
    exit
    no shutdown
exit
...
-----
*A:ALA-60>config>subscr-mgmt#

```

The following is an access example:

- john@internet.com: GI is used to find a subnet and a free address will be allocated from that subnet. Identification strings are returned in Option 122.
- anybody@internet.com: pool\_1 will be used to find a free IP address.
- malicious@internet.com: no address is defined. This user will not get an IP address.

The following is an example of associating a local user database to a local DHCP server.

```

A:pe7>config>router>dhcp#
-----
        local-dhcp-server my_server
        description "my dhcp server"
        user-db "data base 1"
        ...
    exit
-----
A:pe7>config>router>dhcp#

```



---

# Triple Play DHCP Command Reference

---

## Configuration Commands

- [Router DHCP Commands on page 379](#)
- [VPLS DHCP Commands on page 382](#)
- [IES DHCP Commands on page 384](#)
- [VPRN DHCP Commands on page 387](#)
- [IES/VPRN IPv6-DHCP6 Commands on page 391](#)
- [Local User Database Commands on page 393](#)
- [Show Commands on page 400](#)
- [Clear Commands on page 402](#)
- [DHCP Debug Commands on page 403](#)

## Global DHCP Commands

```

config
  — system
    — dhcp6
      — adv-noaddrs-global [esm-proxy] [esm-relay] [relay] [server]
      — no adv-noaddrs-global

```

## Router DHCP Commands

```

config
  — router
    — dhcp
    — dhcp6
      — local-dhcp-server server-name [create]
      — no local-dhcp-server server-name
      — description description-string
      — no description
      — failover
        — [no] ignore-mclt-on-takeover
        — maximum-client-lead-time [hrs hours] [min minutes] [sec seconds]
        — no maximum-client-lead-time
        — partner-down-delay [hrs hours] [min minutes] [sec seconds]
        — no partner-down-delay
        — peer ip-address tag sync-tag
        — no peer
        — [no] shutdown
        — [no] startup-wait-time [min minutes] [sec seconds] [days days]
          [hrs hours] [min minutes] [sec seconds]

```

- [no] **force-renews**
- [no] **ignore-rapid-commit**
- [no] **interface-id-mapping**
- **lease-hold-time** [days *days*] [hrs *hours*] [min *minutes*] [sec *seconds*]
- **no lease-hold-time**
- [no] **lease-hold-time-for**
  - [no] **internal-lease-ipsec**
  - [no] **solicited-release**
- **pool** *pool-name* [create]
- **no pool** *pool-name*
  - **delegated-prefix-length** *bits*
  - **no delegated-prefix-length**
  - **description** *description-string*
  - **no description**
  - [no] **exclude-prefix** *ipv6-prefix/prefix-length*
  - **failover**
    - [no] **ignore-mclt-on-takeover**
    - **maximum-client-lead-time** [hrs *hours*] [min *minutes*] [sec *seconds*]
    - **no maximum-client-lead-time**
    - **partner-down-delay** [hrs *hours*] [min *minutes*] [sec *seconds*]
    - **no partner-down-delay**
    - **peer** *ip-address* **tag** *sync-tag*
    - **no peer**
    - [no] **shutdown**
    - [no] **startup-wait-time** [min *minutes*] [sec *seconds*] [days *days*] [hrs *hours*] [min *minutes*] [sec *seconds*]
- **max-lease-time** [days *days*] [hrs *hours*] [min *minutes*] [sec *seconds*]
- **no max-lease-time**
- **min-lease-time** [days *days*] [hrs *hours*] [min *minutes*] [sec *seconds*]
- **no min-lease-time**
- **minimum-free** *minimum-free* [percent] [event-when-depleted]
- **no minimum-free**
- **offer-time** [min *minutes*] [sec *seconds*]
- **no offer-time**
- **options**
  - **custom-option** *option-number* **address** [*ip-address*...(up to 4 max)]
  - **custom-option** *option-number* **hex** *hex-string*
  - **custom-option** *option-number* **string** *ascii-string*
  - **no custom-option** *option-number*
  - **dns-server** [*ip-address*...(up to 4 max)]
  - **domain-name** *domain-name*
  - **no domain-name**
  - **lease-rebind-time** [days *days*] [hrs *hours*] [min *minutes*] [sec *seconds*]
  - **no lease-rebind-time**
  - **lease-renew-time** [days *days*] [hrs *hours*] [min *minutes*] [sec *seconds*]
  - **no lease-renew-time**
  - **lease-time** [days *days*] [hrs *hours*] [min *minutes*] [sec *seconds*]
  - **no lease-time**



```

— netbios-name-server ip-address [ip-address...(up to 4
max)]
— no netbios-name-server
— netbios-node-type netbios-node-type
— no netbios-node-type
— [no] nak-non-matching-subnet
— prefix ipv6-addr/prefix-len [failover {local | remote}] [pd]
[wan-host] [create]
— no prefix ipv6-addr/prefix-len
— thresholds
—
— thresholds
— [no] minimum-free prefix-length [1..128]
— [no] depleted-event
— minimum percent [0..100]
— no minimum
— subnet {ip-address/mask|ip-address netmask} [create]
— no subnet {ip-address/mask|ip-address netmask}
— [no] address-range start-ip-address end-ip-address
— [no] drain
— [no] exclude-addresses start-ip-address [end-ip-
address]
— maximum-declined maximum-declined
— no maximum-declined
— minimum-free minimum-free [percent] [event-when-
depleted]
— no minimum-free
— options
— 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
— default-router ip-address [ip-address...(up to 4
max)]
— no default-router
— subnet-mask ip-address
— no subnet-mask
— subnet-binding key [sys-id-svc-id | sys-id | string] unbind-
delay [hrs hours] [min mins] [sec secs]
— no subnet-binding key
— use-gi-address [scope scope]
— no use-gi-address
— use-pool-from-client delimiter delimiter
— use-pool-from-client
— no use-pool-from-client
— use-link-address [scope scope]
— no use-link-address
— user-db local-user-db-name
— no user-db
— user-ident user-ident
— no user-ident

```

## VPLS DHCP Commands

```

config
  — service
    — [no] vpls service-id [customer customer-id] [vpn vpn-id] [m-vpls]
      — sap sap-id [split-horizon-group group-name] [capture-sap]
        — dhcp
          — description description-string
          — no description
          — lease-populate [nbr-of-entries]
          — no lease-populate
          — [no] option
            — action {dhcp-action}
            — no action
            — [no] circuit-id [ascii-tuple | vlan-ascii-tuple]
            — remote-id [mac | string string]
            — no remote-id
            — [no] vendor-specific-option
              — [no] client-mac-address
              — [no] sap-id
              — [no] service-id
              — string text
              — no string
              — [no] system-id
            — proxy-server
              — emulated-server
              — lease-time
              — [no] shutdown
            — [no] shutdown
            — [no] snoop
          — dhcp-user-db local-user-db-name
          — no dhcp-user-db
          — dhcp-python-policy policy-name
          — no dhcp-python-policy
          — dhcp6-user-db local-user-db-name
          — no dhcp6-user-db
          — dhcp6
            — description description-string
            — no description
            — [no] option
              — interface-id
              — interface-id ascii-tuple
              — interface-id vlan-ascii-tuple
              — no interface-id
              — remote-id
              — remote-id mac
              — remote-id string [32 chars max]
              — no remote-id
            — [no] shutdown
            — [no] snoop
          — dhcp-python-policy policy-name
          — no dhcp-python-policy
          — dhcp6-user-db local-user-db-name
          — no dhcp6-user-db
    — mesh-sdp sdp-id[:vc-id] [vc-type {ether | vlan}]
      — dhcp
        — description description-string

```

```

— no description
— [no] snoop
— ppp-user-db local-user-db-name
— no ppp-user-db
— pppoe-user-db local-user-db-name
— no pppoe-user-db
— spoke-sdp sdp-id[:vc-id] [vc-type {ether | vlan}] [split-horizon-group group-name]
    — dhcp
        — description description-string
        — no description
        — [no] snoop

```

## IES DHCP Commands

```

config
  — service
    — ies service-id [customer customer-id] [vpn vpn-id]
      — [no] interface ip-int-name
        — dhcp
          — description description-string
          — no description
          — gi-address ip-address [src-ip-addr]
          — no gi-address
          — lease-populate nbr-of-leases
          — no lease-populate
          — [no] option
            — action {replace | drop | keep}
            — no action
            — circuit-id [ascii-tuple | ifindex | sap-id | vlan-ascii-tuple]
            — no circuit-id
            — remote-id [mac | string string]
            — no remote-id
            — [no] vendor-specific-option
              — [no] client-mac-address
              — [no] sap-id
              — [no] service-id
              — string text
              — no string
              — [no] system-id
          — proxy-server
            — emulated-server ip-address
            — no emulated-server
            — lease-time [days days] [hrs hours] [min minutes] [sec seconds] [override]
            — no lease-time
            — [no] shutdown
          — [no] relay-plain-bootp
          — relay-unicast-msg [release-update-src-ip]
          — no relay-unicast-msg
          — python-policy policy-name
          — no python-policy
          — server server1 [server2...(up to 8 max)]
          — no server
          — [no] shutdown
          — [no] trusted
        — [no] ipv6
          — address ipv6-address/prefix-length [eui-64]
          — no address ipv6-address/prefix-length
          — [no] dhcp6-relay
            — description description-string
            — no description
            — lease-populate [nbr-of-leases]
            — lease-populate [nbr-of-leases] route-populate [pd na] [ta]
            — lease-populate [nbr-of-leases] route-populate pd [na] [ta] [exclude]
            — lease-populate [nbr-of-leases] route-populate [pd] [na] ta

```

```

— no lease-populate
— [no] neighbor-resolution
— [no] option
  — interface-id
  — interface-id ascii-tuple
  — interface-id ifindex
  — interface-id sap-id
  — no interface-id
  — [no] remote-id
— server ipv6z-address [ipv6z-address...(up to 8 max)]
— no server [ipv6z-address...(up to 8 max)]
— description description-string
— no description
— source-address ipv6-address
— no source-address
— [no] dhcp6-server
  — max-nbr-of-leases max-nbr-of-leases
  — no max-nbr-of-leases
  — [no] prefix-delegation
    — [no] prefix ipv6-address/prefix-length
      — duid duid [iaid iaid]
      — no duid
      — preferred-lifetime seconds
      — preferred-lifetime infinite
      — no preferred-lifetime
      — valid-lifetime seconds
      — valid-lifetime infinite
      — no valid-lifetime
    — [no] shutdown
— icmp6
  — packet-too-big [number seconds]
  — no packet-too-big
  — param-problem [number seconds]
  — no param-problem
  — redirects [number seconds]
  — no redirects
  — time-exceeded [number seconds]
  — no time-exceeded
  — unreachables [number seconds]
  — no unreachables
— [no] local-proxy-nd
— neighbor ipv6-address mac-address
— no neighbor ipv6-address
— proxy-nd-policy policy-name [policy-name...(up to 5 max)]
— no proxy-nd-policy
— [no] subscriber-interface ip-int-name
  — dhcp
    — gi-address ip-address [src-ip-addr]
    — no gi-address
    — relay-unicast-msg [release-update-src-ip]
    — no relay-unicast-msg
    — [no] virtual-subnet
  — [no] group-interface ip-int-name
    — dhcp
      — client-applications {[dhcp] [pppoe]}

```

- **no client-applications**
- **description** *description-string*
- **no description**
- **gi-address** *ip-address* [*src-ip-addr*]
- **no gi-address**
- **gi-address** [*nbr-of-leases*]
- **lease-populate** [*nbr-of-leases*] **l2-header** [**mac** *ieee-address*]
- **no lease-populate**
- [**no**] **match-circuit-id**
- [**no**] **option**
  - **action** {**replace** | **drop** | **keep**}
  - **no action**
  - **circuit-id** [**ascii-tuple** | **ifindex** | **sap-id** | **vlan-ascii-tuple**]
  - **no circuit-id**
  - **remote-id** [**mac** | **string** *string*]
  - **no remote-id**
  - [**no**] **vendor-specific-option**
    - [**no**] **client-mac-address**
    - [**no**] **pool-name**
    - [**no**] **sap-id**
    - [**no**] **service-id**
    - **string** *text*
    - **no string**
    - [**no**] **system-id**
- **relay-unicast-msg** [*release-update-src-ip*]
- **no relay-unicast-msg**
- **server** *server1* [*server2...*(up to 8 max)]
- **no server**
- [**no**] **shutdown**
- [**no**] **trusted**
- [**no**] **pppoe**
  - **anti-spoof** *pppoe-anti-spoofing-type*
  - **no anti-spoof**
  - **description** *description-string*
  - **no description**
  - **dhcp-client**
    - [**no**] **ccag-use-origin-sap**
  - **policy** *ppp-policy-name*
  - **no policy**
  - **sap-session-limit** *sap-session-limit*
  - **no sap-session-limit**
  - **session-limit** *session-limit*
  - **no session-limit**
  - [**no**] **shutdown**
  - **user-db** *local-user-db-name*
  - **no user-db**

## VPRN DHCP Commands

```

config
  — service
    — vprn
      — dhcp
      — dhcp6
        — local-dhcp-server-server server-name [create]
        — no local-dhcp-server server-name
          — failover
            — [no] ignore-mclt-on-takeover
            — maximum-client-lead-time [hrs hours] [min minutes] [sec seconds]
            — no maximum-client-lead-time
            — partner-down-delay [hrs hours] [min minutes] [sec seconds]
            — no partner-down-delay
            — peer ip-address tag sync-tag
            — no peer
            — [no] shutdown
            — [no] startup-wait-time [min minutes] [sec seconds] [days days] [hrs hours] [min minutes] [sec seconds]
        — lease-hold-time [days days] [hrs hours] [min minutes] [sec seconds]
        — no lease-hold-time
        — [no] lease-hold-time-for
          — [no] internal-lease-ipsec
          — [no] solicited-release
        — pool pool-name [create]
        — no pool pool-name
          — failover
            — [no] ignore-mclt-on-takeover
            — maximum-client-lead-time [hrs hours] [min minutes] [sec seconds]
            — no maximum-client-lead-time
            — partner-down-delay [hrs hours] [min minutes] [sec seconds]
            — no partner-down-delay
            — peer ip-address tag sync-tag
            — no peer
            — [no] shutdown
            — [no] startup-wait-time [min minutes] [sec seconds] [days days] [hrs hours] [min minutes] [sec seconds]
          — prefix ipv6-addr/prefix-len [failover {local | remote}] [pd] [wan-host] [create]
          — no prefix ipv6-addr/prefix-len
            — thresholds
              — [no] minimum-free prefix-length [1..128]
              — [no] depleted-event
              — minimum [percent [0..100]] [number [0..4294967295]]
              — no minimum
      — [no] interface ip-int-name
        — dhcp

```

- **description** *description-string*
- **no description**
- **gi-address** *ip-address* [*src-ip-addr*]
- **no gi-address**
- **lease-populatelease-populate**[**no**] **option**
  - **action** {**replace** | **drop** | **keep**}
  - **no action**
  - **circuit-id** [**ascii-tuple** | **ifindex** | **sap-id** | **vlan-ascii-tup**]
  - **no circuit-id**
  - **remote-id** [**mac** | **string** *string*]
  - **no remote-id**
  - [**no**] **vendor-specific-option**
    - [**no**] **client-mac-address**
    - [**no**] **pool-name**
    - [**no**] **sap-id**
    - [**no**] **service-id**
    - **string** *text*
    - **no string**
    - [**no**] **system-id**
- [**no**] **nak-non-matching-subnet**
- **python-policy** *policy-name*
- **no python-policy**
- **proxy-server**
  - **emulated-server** *ip-address*
  - **no emulated-server**
  - **lease-time** [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*] [**override**]
  - **no lease-time**
  - [**no**] **shutdown**
- **server** *server1* [*server2...*(up to 8 max)]
- **no server**
- [**no**] **shutdown**
- [**no**] **trusted**
- [**no**] **use-arp**
- **user-ident** *user-ident*
- **no user-ident**



## VPRN DHCP Subscriber Interface Commands

```

config
  — service
    — vprn
      — subscriber-interface ip-int-name [fwd-service service-id fwd-subscriber-interface
        ip-int-name] [create]
      — no subscriber-interface ip-int-name
        — dhcp
          — client-applications dhcp
          — client-applications pppoe
          — client-applications dhcp pppoe
          — no client-applications
          — description description-string
          — no description
          — gi-address ip-address [src-ip-addr]
          — no gi-address
          — lease-populate nbr-of-leases
          — no lease-populate
          — [no] option
            — [no] vendor-specific-option
              — [no] client-mac-address
              — [no] sap-id
              — [no] service-id
              — string text
              — no string
              — [no] system-id
            — proxy-server
              — emulated-server ip-address
              — no emulated-server
              — lease-time [days days] [hrs hours] [min minutes] [sec
                seconds] [override]
              — no lease-time
              — [no] shutdown
            — relay-unicast-msg [release-update-src-ip]
            — no relay-unicast-msg
            — server server1 [server2...(up to 8 max)]
            — no server
            — [no] shutdown
            — [no] virtual-subnet
          — [no] group-interface ip-int-name
            — [no] arp-populate
            — arp-timeout seconds
            — no arp-timeout
            — authentication-policy name
            — no authentication-policy
            — description description-string
            — no description
            — dhcp
              — description description-string
              — no description
              — gi-address ip-address [src-ip-addr]
              — no gi-address
              — lease-populate [nbr-of-leases] l2-header [mac ieee-
                address]

```

- **no lease-populate**
- **[no] match-circuit-id**
- **[no] option**
  - **action** {replace | drop | keep}
  - **no action**
  - **circuit-id** [ascii-tuple|ifindex|sap-id|vlan-ascii-tuple]
  - **no circuit-id**
  - **remote-id** [mac | string *string*]
  - **no remote-id**
  - **[no] vendor-specific-option**
    - **[no] client-mac-address**
    - **[no] sap-id**
    - **[no] service-id**
    - **string** *text*
    - **no string**
    - **[no] system-id**
- **proxy-server**
  - **emulated-server** *ip-address*
  - **no emulated-server**
  - **lease-time** [days *days*] [hrs *hours*] [min *minutes*] [sec *seconds*] [override]
  - **no lease-time**
  - **[no] shutdown**
- **relay-unicast-msg** [*release-update-src-ip*]
- **no relay-unicast-msg**
- **server** *server1* [*server2*...(up to 8 max)]
- **no server**
- **[no] shutdown**
- **[no] trusted**

## IES/VP RN IPv6-DHCP6 Commands

```

configure
  — service
    — service-id/ service-id
      — subscriber-interface
        — group-interface
          — ipv6
            — [no] allow-multiple-wan-addresses
            — dhcp6
              — [no] option
                — interface-id ascii-tuple
                — interface-id ifindex
                — interface-id sap-id
                — interface-id string string
                — no interface-id
                — [no] remote-id
            — [no] override-slaac
            — [no] pd-managed-route
            — [no] proxy-server
              — client-applications {[dhcp] [ppp]}
              — no client-applications
              — preferred-lifetime [days days] [hrs hours] [min minutes] [sec seconds]
              — preferred-lifetime infinite
              — no preferred-lifetime
              — rebind-timer [days days] [hrs hours] [min minutes] [sec seconds]
              — no rebind-timer
              — renew-timer [days days] [hrs hours] [min minutes] [sec seconds]
              — no renew-timer
              — server-id duid-en hex hex-string
              — server-id duid-en string ascii-string
              — server-id duid-ll
              — no server-id
              — no shutdown
              — valid-lifetime [days days] [hrs hours] [min minutes] [sec seconds]
              — valid-lifetime infinite
              — no valid-lifetime
            — python-policy name
            — no python-policy
            — [no] relay
              — client-applications [dhcp] [ppp]
              — no client-applications
              — description description-string
              — no description
              — link-address ipv6-address
              — no link-address
              — [no] server [ipv6z-address...(upto 8 max)]
              — no server
              — [no] shutdown
              — source-address ipv6-address

```

```

— no source-address
— user-db local-user-db-name
— no user-db
— [no] qos-route-lookup
— [no] router-advertisements
— current-hop-limit limit
— no current-hop-limit
— [no] dns-options
— [no] include-dns
— rdnss-lifetime seconds
— rdnss-lifetime infinite
— no rdnss-lifetime
— force-mcast [ip] [mac]
— no force-mcast
— [no] managed-configuration
— max-advertisement seconds
— no max-advertisement
— min-advertisement seconds
— no min-advertisement
— mtu bytes
— no mtu
— [no] other-stateful-configuration
— [no] prefix-options
— [no] autonomous
— [no] on-link
— preferred-lifetime seconds
— preferred-lifetime infinite
— no preferred-lifetime
—
— reachable-time milli-seconds
— no reachable-time
— retransmit-time milli-seconds
— no retransmit-time
— router-lifetime seconds
— router-lifetime no-default-router
— no router-lifetime
— [no] shutdown
— router-solicit
— inactivity-timer [days days] [hrs hours] [min
   minutes] [sec seconds]
— inactivity-timer infinite
— no inactivity-timer
— min-auth-interval [days days] [hrs hours] [min
   minutes] [sec seconds]
— no min-auth-interval
— [no] shutdown
— user-db local-user-db-name
— no user-db

```

## Local User Database Commands

- [IPoE Commands on page 393](#)
- [PPP Commands on page 397](#)

## IPoE Commands

```

config
— subscriber-mgmt
    — local-user-db local-user-db-name [create]
    — no local-user-db local-user-db-name
        — description description-string
        — no description
    — ipoe
        — host host-name [create]
        — no host host-name
            — acct-policy acct-policy-name [duplicate acct-policy-name]
            — no acct-policy
            — address gi-address [scope scope]
            — address ip-address
            — address pool pool-name [secondary-pool sec-pool-name]
                [delimiter delimiter]
            — address use-pool-from-client [delimiter delimiter]
            — no address
            — auth-domain-name domain-name
            — no auth-domain-name
            — auth-policy policy-name
            — no auth-policy
            — diameter-application-policy policy-name
            — no diameter-application-policy
            — diameter-auth-policy name
            — no diameter-auth-policy
            — gi-address ip-address
            — no gi-address
            — host-identification
                — circuit-id string ascii-string
                — circuit-id hex hex-string
                — no circuit-id
                — derived-id derived-id-string
                — no derived-id
                — encap-tag-range start-tag start-tag end-tag end-tag
                — no encap-tag-range
                — mac ieee-address
                — no mac
                — option60 hex-string
                — no option60
                — remote-id hex hex-string
                — remote-id string ascii-string
                — no remote-id
                — sap-id sap-id
                — no sap-id
                — service-id service-id
                — no service-id
                — string string

```

- **no string**
- **system-id** *system-id*
- **no system-id**
- **identification-strings** *option-number* [**create**]
- **no identification-strings**
  - **ancp-string** *ancp-string*
  - **no ancp-string**
  - **app-profile-string** *app-profile-string*
  - **no app-profile-string**
  - **category-map** *category-map-name*
  - **no category-map**
  - **inter-dest-id** *intermediate-destination-id*
  - **no inter-dest-id**
  - **sla-profile-string** *sla-profile-string*
  - **no sla-profile-string**
  - **sub-profile-strings** *sub-profile-string*
  - **no sub-profile-string**
  - **subscriber-id** *sub-ident-string*
  - **no subscriber-id**
- **ipv6-address** *ipv6-address*
- **no ipv6-address**
- **ipv6-delegated-prefix** *ipv6-prefix/prefix-length*
- **no ipv6-delegated-prefix**
- **ipv6-delegated-prefix-length** *bits*
- **no ipv6-delegated-prefix-length**
- **ipv6-delegated-prefix-pool** *pool-name*
- **no ipv6-delegated-prefix-pool**
- **[no] ipv6-lease-times**
  - **preferred-lifetime** [**days** *days*] [**hrs** *hrs*] [**min** *min*] [**sec** *sec*]
  - **preferred-lifetime** *infinite*
  - **no preferred-lifetime**
  - **rebind-timer** [**days** *days*] [**hrs** *hrs*] [**min** *min*] [**sec** *sec*]
  - **no rebind-timer**
  - **renew-timer** [**days** *days*] [**hrs** *hrs*] [**min** *min*] [**sec** *sec*]
  - **no renew-timer**
  - **valid-lifetime** [**days** *days*] [**hrs** *hrs*] [**min** *min*] [**sec** *sec*]
  - **valid-lifetime** *infinite*
  - **no valid-lifetime**
- **ipv6-slaac-prefix** *ipv6-prefix/prefix-length*
- **no ipv6-slaac-prefix**
- **ipv6-slaac-prefix-pool** *pool*
- **no ipv6-slaac-prefix-pool**
- **ipv6-wan-address-pool** *pool-name*
- **no ipv6-wan-address-pool**
- **link-address** *ipv6-address*
- **no link-address**
- **match-radius-proxy-cache**
  - **delete-hold-time** *seconds*
  - **no delete-hold-time**
  - **fail-action** {**continue**|**drop**}
  - **no fail-action**
  - **mac-format** *mac-format*
  - **no mac-format**

```

— match {circuit-id|mac|remote-id}
— match option [1..254] [option6 [1..65535]]
— match option6 [1..65535]
— no match
— server [service service-id] name server-name
— no server
— msap-defaults
— group-interface ip-int-name [prefix {port-id}]
— group-interface ip-int-name [suffix {port-id}]
— no group-interface
— policy msap-policy-name
— no policy
— service service-id
— no service
— options
— custom-option option-number address [ip-
address...(up to 4 max)]
— custom-option option-number hex hex-string
— custom-option option-number string ascii-string
— no custom-option option-number
— default-router ip-address [ip-address...(up to 4
max)]
— no default-router
— dns-server address [ip-address...(upto 4 max)]
— no dns-server
— domain-name domain-name
— no domain-name
— lease-rebind-time [days days] [hrs hours] [min min-
utes] [sec seconds]
— no lease-rebind-time
— lease-renew-time [days days] [hrs hours] [min min-
utes] [sec seconds]
— no lease-renew-time
— lease-time [days days] [hrs hours] [min minutes] [sec
seconds]
— no lease-time
— netbios-name-server ip-address [ip-address...(up to 4
max)]
— no netbios-name-server
— netbios-node-type netbios-node-type
— no netbios-node-type
— subnet-mask ip-address
— no subnet-mask
— options6
— dns-server ipv6-address [ipv6-address...(upto 4 max)]
— no dns-server
— to-client-options
— dhcpv4
— option option-number address ipv4-address [ipv4-
address...(upto 4 max)]
— option option-number hex hex-string
— option option-number string ascii-string
— no option option-number
— dhcpv6

```

- **option** *option-number* **address** *ipv6-address*  
[*ipv6-address...*(upto 4 max)]
- **option** *option-number* **hex** *hex-string*
- **option** *option-number* **string** *ascii-string*
- **no option** *option-number*
- **ipv4**
  - **option** *option-number* **address** [*ip-address...*(up to 4 max)]
  - **option** *option-number* **hex** *hex-string*
  - **option** *option-number* **string** *ascii-string*
  - **no option** *option-number*
- **ipv6**
  - **option** *option-number* **address** [*ip-address...*(up to 4 max)]
  - **option** *option-number* **hex** *hex-string*
  - **option** *option-number* **string** *ascii-string*
  - **no option** *option-number*
- **retail-service-id** *service-id*
- **no retail-service-id**
- **server** *ip-address*
- **no server**
- **server6** *ipv6-address*
- **no server6**
- **[no] shutdown**
- **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*
- **match-list** *ipoe-match-type-1* [*ipoe-match-type-2...*(up to 4 max)]
- **no match-list**



## PPP Commands

```

config
— subscriber-mgmt
— local-user-db local-user-db-name [create]
— no local-user-db local-user-db-name
— description description-string
— no description
— ppp
— host host-name [create]
— no host host-name
— [no] access-loop-encapsulation
— encap-offset [type type]
— no encap-offset
— rate-down rate
— no rate-down
— access-loop-information
— circuit-id sap-id
— circuit-id ASCII string
— no circuit-id
— remote-id mac
— remote-id ASCII string
— no remote-id
— acct-policy acct-policy-name [duplicate acct-policy-name]
— no acct-policy
— address gi-address [scope scope]
— address ip-address[/prefix-length]
— address pool pool-name [secondary-pool sec-pool-name]
[delimiter delimiter]
— address use-pool-from-client [delimiter delimiter]
— no address
— auth-policy policy-name
— no auth-policy
— diameter-application-policy policy-name
— no diameter-application-policy
— diameter-auth-policy name
— no diameter-auth-policy
— [no] force-ipv6cp
— host-identification
— circuit-id string ascii-string
— circuit-id hex hex-string
— no circuit-id
— encap-tag-range start-tag start-tag end-tag end-tag
— no encap-tag-range
— mac ieee-address
— no mac
— remote-id hex hex-string
— remote-id string ascii-string
— no remote-id
— sap-id sap-id
— no sap-id
— service-name service-name
— no service-name
— username user-name
— username user-name [no-domain]

```

```

— username user-name domain-only
— no username
— identification-strings option-number [create]
— no identification-strings
— ancp-string ancp-string
— no ancp-string
— app-profile-string app-profile-string
— no app-profile-string
— category-map category-map-name
— no category-map
— inter-dest-id intermediate-destination-id
— no inter-dest-id
— sla-profile-string sla-profile-string
— no sla-profile-string
— sub-profile-string sub-profile-string
— no sub-profile-string
— subscriber-id sub-ident-string
— no subscriber-id
— interface ip-int-name service-id service-id
— no interface
— ipv6-address ipv6-address
— no ipv6-address
— ipv6-delegated-prefix ipv6-prefix/prefix-length
— no ipv6-delegated-prefix
— ipv6-delegated-prefix-length bits
— no ipv6-delegated-prefix-length
— ipv6-delegated-prefix-pool pool-name
— no ipv6-delegated-prefix-pool
— [no] ipv6-lease-times
— preferred-lifetime [days days] [hrs hrs] [min min] [sec sec]
— preferred-lifetime infinite
— no preferred-lifetime
— rebind-timer [days days] [hrs hrs] [min min] [sec sec]
— no rebind-timer
— renew-timer [days days] [hrs hrs] [min min] [sec sec]
— no renew-timer
— valid-lifetime [days days] [hrs hrs] [min min] [sec sec]
— valid-lifetime infinite
— no valid-lifetime
— ipv6-slaac-prefix ipv6-prefix/prefix-length
— no ipv6-slaac-prefix
— ipv6-slaac-prefix-pool pool
— no ipv6-slaac-prefix-pool
— ipv6-wan-address-pool pool-name
— no ipv6-wan-address-pool
— l2tp
— group tunnel-group-name [service-id service-id]
— no group
— msap-defaults
— group-interface ip-int-name [prefix {port-id}]
— group-interface ip-int-name [suffix {port-id}]
— no group-interface
— policy msap-policy-name

```

- **no policy**
- **service** *service-id*
- **no service**
- **options**
  - **custom-option** *option-number* **address** [*ip-address...*(up to 4 max)]
  - **custom-option** *option-number* **hex** *hex-string*
  - **custom-option** *option-number* **string** *ascii-string*
  - **no custom-option** *option-number*
  - **dns-server** **address** [*ip-address...*(upto 4 max)]
  - **no dns-server**
  - **netbios-name-server** **ip-address** [*ip-address...*(up to 4 max)]
  - **no netbios-name-server**
- **options6**
- **pado-delay** *deci-seconds*
- **no pado-delay**
- **password** **ignore**
- **password** {**chap** *password* | **pap** *password*} [**hash**|**hash2**]
- **no password**
- **pre-auth-policy** *policy-name*
- **no pre-auth-policy**
- **to-client-options**
  - **ipv4**
    - **option** *option-number* **address** [*ip-address...*(up to 4 max)]
    - **option** *option-number* **hex** *hex-string*
    - **option** *option-number* **string** *ascii-string*
    - **no option** *option-number*
  - **ipv6**
    - **option** *option-number* **address** [*ip-address...*(up to 4 max)]
    - **option** *option-number* **hex** *hex-string*
    - **option** *option-number* **string** *ascii-string*
    - **no option** *option-number*
- **retail-service-id** *service-id*
- **no retail-service-id**
- [**no**] **shutdown**

## Show Commands

```

show
  — router
    — dhcp
      — lease-state [sap sap-id]
      — servers
      — statistics [interface ip-int-name | ip-address]
      — summary
      — local-dhcp-server server-name
        — associations
        — declined-addresses ip-address[/mask] [detail]
        — declined-addresses pool pool-name
        — free-addresses ip-address[/mask]
        — free-addresses summary [subnet ip-address[/mask]]
        — free-addresses pool pool-name
        — leases ip-address[/mask] address-from-user-db [detail]
        — leases ip-address[/mask] dhcp-host dhcp-host-name [detail]
        — leases ip-address[/mask] ppp-host ppp-host-name [detail]
        — leases ip-address[/mask] [detail]
        — pool-ext-stats [pool-name]
        — server-stats
        — subnet-ext-stats ip-address[/mask]
        — subnet-ext-stats pool pool-name
        — subnet-stats ip-address[/mask]
        — subnet-stats pool pool-name
        — summary
      — servers
    — router
      — dhcp6
        — local-dhcp-server server-name
          — associations
          — interface-id-mapping
          — leases [ipv6-address/prefix-length] [type] [state] [detail]
          — pool-ext-stats [pool-name]
          — prefix-ext-stats ipv6-address/prefix-length
          — prefix-ext-stats pool pool-name
          — pool-threshold-stats [pool-name] detail [format {exact|scientific}]
          — pool-threshold-stats [pool-name]
          — prefix-threshold-stats pool pool-name detail [format {exact|scientific}]
          — prefix-threshold-stats pool pool-name
          — prefix-threshold-stats ipv6-address/prefix-length detail [format {exact|scientific}]
          — prefix-threshold-stats ipv6-address/prefix-length
          — server-stats
          — summary
        — statistics
        — summary
  show
    — service
      — id service-id
        — dhcp
          — 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]} ] [detail]
— statistics [sap sap-id] | [sdp [sdp-id[:vc-id]] | interface ip-int-name]
— summary
— virtual-subnet subscriber sub-ident
— virtual-subnet [sap sap-id]
— dhcp6
  — lease-state [detail]
  — lease-state [detail] interface interface-name
  — lease-state [detail] ipv6-address ipv6-prefix[/prefix-length]
  — lease-state [detail] mac ieee-address
  — statistics [interface ip-int-name]
  — summary

```

```

show
  — system
    — dhcp6

```

## Configuration Commands

## Tools Commands

- tools**
  - **subscriber-mgmt**
    - **remap-lease-state** **old-mac** *ieee-address* **mac** *ieee-address*
    - **remap-lease-state** **sap** *sap-id* [*mac ieee-address*]

## Clear Commands

- clear**
  - **router**
    - **dhcp**
      - **lease-state** [**interface** *ip-int-name* | *ip-addr* | **ip-address** *ip-address* | **mac** *ieee-address*]
      - **local-dhcp-server** *server-name*
        - **declined-addresses** *ip-address[/mask]*
        - **declined-addresses** **pool** *pool-name*
        - **leases** *ip-address[/mask]* [**offered**]
        - **pool-ext-stats** [*pool-name*]
        - **server-stats**
        - **subnet-ext-stats** *ip-address[/mask]*
        - **subnet-ext-stats** **pool** *pool-name*
      - **statistics** [*ip-int-name* | *ip-address*]
    - **dhcp6**
      - **local-dhcp-server** *server-name*
        - **leases** [*ipv6-address/prefix-length*] [*type*] [*state*]
        - **leases** **all** [*type*] [*state*]
        - **pool-ext-stats** [*pool-name*]
        - **pool-threshold-stats** [*pool-name*]
        - **prefix-ext-stats** *ipv6-address/prefix-length*
        - **prefix-ext-stats** **pool** *pool-name*
        - **prefix-threshold-stats** *ipv6-address/prefix-length*
        - **prefix-threshold-stats** **pool** *pool-name*
        - **server-stats**
      - **statistics**
  - **service**
    - **id** *service-id*
      - **dhcp**
        - **lease-state** [**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* [**no-dhcp-release**]
        - **lease-state** **mac** *ieee-address* **no-dhcp-release**
        - **lease-state** **sap** *sap-id* [**no-dhcp-release**]
        - **lease-state** **sdp** *sdp-id:vc-id* [**no-dhcp-release**]
        - **statistics** [**sap** *sap-id* | **sdp** [*sdp-id[:vc-id]*] | **interface** *ip-int-name* | *ip-address*]
      - **dhcp6**
        - **lease-state** [**ip-address** *ipv6-address/prefix-length*] [**mac** *ieee-address*]

## DHCP Debug Commands

```

debug
  — router
    — ip
      — [no] dhcp [ip-int-name]
        — detail-level {low | medium | high}
        — no detail-level
        — mode {dropped-only | ingr-and-dropped | egr-ingr-and-dropped}
        — no mode

```

## Tools Commands

```

tools
  — perform
    — subscriber-mgmt
      — local-user-db local-user-db-name
        — ppp
          — authentication pppoe-user-name [password password]
          — host-lookup [mac ieee-address] [remote-id remote-id] [user-name user-name] [circuit-id circuit-id | circuit-id-hex circuit-id-hex]
        — dhcp
          — host-lookup [mac ieee-address] [remote-id remote-id] [sap-id sap-id] [service-id service-id] [string vso-string] [system-id system-id] [option60 hex-string] [circuit-id circuit-id | circuit-id-hex circuit-id-hex]

```





## Triple Play DHCP Configuration Commands

Note: For the 7450 ESS configurations, the DHCP6 and IPv6 ESM commands apply only when in mixed-mode.

### Global Commands

#### shutdown

<b>Syntax</b>	<b>[no] shutdown</b>
<b>Context</b>	<pre> config&gt;service&gt;ies&gt;if&gt;dhcp config&gt;service&gt;vpls&gt;sap&gt;dhcp config&gt;service&gt;vpls&gt;sap&gt;dhcp6 config&gt;service&gt;vpls&gt;sap&gt;dhcp&gt;option&gt;vendor config&gt;service&gt;vpls&gt;sap&gt;ipoe-session config&gt;service&gt;vprn&gt;if&gt;dhcp config&gt;service&gt;vprn&gt;if&gt;dhcp&gt;proxy-server config&gt;subscr-mgmt&gt;loc-user-db config&gt;subscr-mgmt&gt;loc-user-db&gt;dhcp&gt;host config&gt;subscr-mgmt&gt;loc-user-db&gt;dhcp&gt;host&gt;options config&gt;subscr-mgmt&gt;loc-user-db&gt;ppp&gt;host config&gt;router&gt;dhcp6&gt;server&gt;failover config&gt;router&gt;dhcp&gt;server&gt;failover </pre>
<b>Description</b>	<p>This command administratively disables an entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics.</p> <p>The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they may be deleted.</p> <p>The <b>no</b> form of this command places the entity into an administratively enabled state.</p>

#### description

<b>Syntax</b>	<b>description</b> <i>description-string</i> <b>no description</b>
<b>Context</b>	<pre> config&gt;service&gt;vpls&gt;sap&gt;dhcp config&gt;service&gt;vpls&gt;sap&gt;dhcp6 config&gt;service&gt;vpls&gt;sap&gt;ipoe-session config&gt;service&gt;ies&gt;if&gt;dhcp config&gt;service&gt;ies&gt;if&gt;ipv6&gt;dhcp6-relay config&gt;service&gt;vprn&gt;if&gt;dhcp config&gt;router&gt;dhcp&gt;server config&gt;router&gt;dhcp&gt;server&gt;pool config&gt;subscr-mgmt&gt;loc-user-db config&gt;service&gt;vprn&gt;sub-if&gt;ipv6&gt;dhcp6&gt;relay </pre>

```
config>service>ies>sub-if>ipv6>dhcp6>relay
```

<b>Description</b>	<p>This command creates a text description stored in the configuration file for a configuration context.</p> <p>The <b>description</b> command associates a text string with a configuration context to help identify the content in the configuration file.</p> <p>The <b>no</b> form of this command removes the string from the configuration.</p>
<b>Default</b>	<p>No description associated with the configuration context.</p>
<b>Parameters</b>	<p><i>description-string</i> — The description character string. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.</p>

## System Commands

### adv-noaddrs-global

<b>Syntax</b>	<b>adv-noaddrs-global [esm-proxy] [esm-relay] [relay] [server]</b> <b>no adv-noaddrs-global</b>
<b>Context</b>	config>system>dhcp6
<b>Description</b>	<p>This command configures the different DHCPv6 applications to send the NoAddrsAvail Status-Code in DHCPv6 Advertise messages at the global DHCP message level.</p> <p>By default, all applications send the NoAddrsAvail Status-Code in DHCPv6 Advertise messages at the IA_NA Option level.</p>
<b>Default</b>	no adv-noaddrs-global. All applications send the NoAddrsAvail Status-Code in DHCPv6 Advertise messages at the IA_NA Option level.
<b>Parameters</b>	<p>Different applications for which NoAddrsAvail Status-Code in DHCPv6 Advertise messages can be configured at the global DHCP message level.</p> <p>The only valid combination in current SROS is “adv-noaddrs-global esm-relay server”.</p> <p><b>esm-proxy</b> — Specifies the DHCPv6 proxy server on subscriber group-interfaces. Not supported in current SR OS.</p> <p><b>esm-relay</b> — Specifies the DHCPv6 relay on subscriber group-interfaces. Must be enabled together with the DHCPv6 server (server) application.</p> <p><b>relay</b> — Specifies the DHCPv6 relay on regular IES/VP RN interfaces. Not supported in current SR OS.</p> <p><b>server</b> — Specifies the DHCPv6 server. Must be enabled together with the DHCPv6 relay on subscriber interfaces (esm-relay) application.</p>

---

## DHCP Configuration Commands

### local-dhcp-server

<b>Syntax</b>	<b>local-dhcp-server</b> <i>server-name</i> [ <b>create</b> ] <b>no local-dhcp-server</b> <i>server-name</i>
<b>Context</b>	config>router>dhcp config>service>vprn>dhcp
<b>Description</b>	This command instantiates a local DHCP server. A local DHCP server can serve multiple interfaces but is limited to the routing context it was which it was created.
<b>Default</b>	none
<b>Parameters</b>	<i>server-name</i> — Specifies the name of local DHCP server.

### delegated-prefix-length

<b>Syntax</b>	<b>delegated-prefix-length</b> <i>bits</i> <b>delegated-prefix-length</b> <i>variable</i> <b>no delegated-prefix-length</b>
<b>Context</b>	configure>router>local-dhcp-server>pool
<b>Description</b>	This command configures the subscriber-interface level setting for delegated prefix length. The delegated prefix length for a subscriber- interface can be either set to a fixed value that is explicitly configured under the subscriber-interface CLI hierarchy or a variable value that can be obtained from various sources. This command can be changed only when no IPv6 prefixes are configured under the subscriber-interface.
<b>Default</b>	no delegated-prefix-length This means that the delegated prefix length is 64.
<b>Parameters</b>	<i>bits</i> — The delegated prefix length in bits. This value will be applicable to the entire subscriber-interface. In case that the delegated prefix length is also supplied via other means (LUDB, RADIUS or DHCP Server), such supplied value must match the value configured under the subscriber-interface. Otherwise the prefix instantiation in 7x50 will fail.  <b>Values</b> 48 — 64  <i>variable</i> — The delegated prefix value can be of any length between 48..64. The value itself can vary between the prefixes and it will be provided at the time of prefix instantiation. The order of priority for the source of the delegated prefix length is: <ul style="list-style-type: none"> <li>• LUDB</li> <li>• RADIUS</li> <li>• DHCPv6 server</li> </ul>

## failover

<b>Syntax</b>	<b>failover</b>
<b>Context</b>	config>router>dhcp>server config>router>dhcp6>server
<b>Description</b>	This command enables the context to configure failover parameters.

## maximum-client-lead-time

<b>Syntax</b>	<b>maximum-client-lead-time [hrs <i>hours</i>] [min <i>minutes</i>] [sec <i>seconds</i>]</b> <b>no maximum-client-lead-time</b>
<b>Context</b>	configure>router>dhcp>server>failover configure>router>dhcp>server>pool>failover configure>service>vpn>dhcp>server>failover configure>service>vpn>dhcp>server>pool>failover configure>router>dhcp6>server>failover configure>router>dhcp6>server>pool>failover configure>service>vpn>dhcp6>server>failover configure>service>vpn>dhcp6>server>pool>failover
<b>Context</b>	<p>Maximum-client-lead-time (MCLT) is the maximum time that a DHCP server can extend client lease time beyond the lease time currently known by the DHCP partner node. In dual-homed environment, the initial lease time for all DHCP clients is strictly restricted to MCLT. Consecutive DHCP renewals are allowed to extend the lease time beyond the MCLT.</p> <p>The MCLT is a safeguard against IP address/prefix duplication in cases of a lease synchronization failure.</p> <p>Consider a case whereby the primary DHCP server assign a new lease to the client but it crashes before it sends a sync update to the partner (secondary DHCP server). Because of the primary DHCP server failure, the secondary server (whose partner-down-delay is set to 0) is not aware of the IP address/prefix that has been allocated on the primary server. This condition creates the possibility in which the secondary DHCP server allocates the same address/prefix to another client. This would cause IP address/prefix duplication. MCLT is put in place to prevent this scenario.</p> <p>Lease synchronization failure can be caused either by a node failure, or a failure of the link over which the DHCP leases are synchronized (Multi-Chassis Synchronization (MCS) link). Synchronization failure detection can take up to three seconds. Once the synchronization failure is detected, the minimum time required for a DHCP server to start delegating new addresses/prefixes from the prefix designated as remote is the sum of the maximum-client-lead-time and the partner-down-delay.</p> <p>During the failed state (DHCP peer is unreachable), the DHCP lease time for the new clients will be restricted to MCLT while for the existing clients the lease time will over time (by consecutive DHCP renewals) gradually be reduced to the MCLT.</p>
<b>Default</b>	10 minutes
<b>Parameters</b>	<p><b>hrs <i>hours</i></b> — Specifies the maximum amount of time, in hours, that one server can extend a lease for a client's binding beyond the time known by the partner server.</p> <p><b>Values</b> 1 — 23</p>

**min minutes** — Specifies the maximum amount of time, in minutes, that one server can extend a lease for a client's binding beyond the time known by the partner.

**Values** 1 — 59

**sec seconds** — Specifies the maximum amount of time, in seconds, that one server can extend a lease for a client's binding beyond the time known by the partner.

**Values** 1 — 59

### partner-down-delay

<b>Syntax</b>	<b>partner-down-delay</b> [ <i>hrs hours</i> ] [ <i>min minutes</i> ] [ <i>sec seconds</i> ] <b>no partner-down-delay</b>
<b>Context</b>	config>router>dhcp>server>failover config>router>dhcp6>server>failover
<b>Description</b>	Since the DHCP lease synchronization failure can be caused by the failure of the Multi-Chassis Synchronization (MCS) link (and not necessary the entire node), there is a possibility that both DHCP servers are operational during the failure. The purpose of the partner-down-delay is to allow the operator enough time to remedy the failed situation and to avoid duplication of IP addresses/prefixes during the failure. During the partner-down-delay time, the prefix designated as remote will be eligible only for renewals of the existing DHCP leases that have been synchronized by the peering node. Only after the sum of the partner-down-delay and the maximum-client-lead-time will the prefix designated as remote be eligible for assignment of the new DHCP leases.
<b>Default</b>	23 hours, 59minutes and 59 seconds
<b>Parameters</b>	<p><b>hrs hours</b> — Specifies the partner-down delay time in hours.</p> <p><b>Values</b> 1 — 23</p> <p><b>min minutes</b> — Specifies the partner-down delay time in minutes.</p> <p><b>Values</b> 1 — 59</p> <p><b>sec seconds</b> — Specifies the partner-down delay time in seconds.</p> <p><b>Values</b> 1 — 59</p>

### peer

<b>Syntax</b>	<b>peer</b> <i>ip-address</i> <b>tag</b> <i>sync-tag-name</i> <b>no peer</b> <i>ip-address</i>
<b>Context</b>	config>router>dhcp6>server>failover config>router>dhcp>server>failover
<b>Description</b>	<p>DHCP leases are synchronized per DHCP server. The pair of synchronizing servers (peers) is identified by a tag. The synchronization information is carried over the Multi-Chassis Synchronization (MCS) link between the two peers. MCS link is a logical link (IP or MPLS).</p> <p>MCS runs over TCP, port 45067 and it is using either data traffic or keepalives to detect failure on the communication link between the two nodes. In the absence of any MCS data traffic for more than</p>

0.5sec, MCS will send its own keepalive to the peer. If a reply is NOT received within 3sec, MCS will declare its operation state as DOWN and the DB Sync state as out-of-sync. MCS will consequently notify its clients (DHCP Server being one of them) of this. It can take up to 3 seconds before the DHCP client realizes that the inter-chassis communication link has failed.

Note that the inter-chassis communication link failure does not necessarily assume the same failed fate for the access links.

**Parameters** *ip-address* — Specifies the IPv4 address of the peer.

**sync-tag** *sync-tag* — Specifies a synchronization tag to be used while synchronizing with the multi-chassis peer.

## startup-wait-time

**Syntax** **[no] startup-wait-time [min *minutes*] [sec *seconds*]**

**Context** config>router>dhcp6>server>failover  
config>router>dhcp>server>failover

**Description** This command enables startup-wait-time during which each peer waits after the initialization process before assuming the active role for the prefix designated as local. This is to avoid transient issues during the initialization process.

**Default** 2 minutes

**Parameters** **min** *minutes* — Specifies the time in minutes that one server attempts to contact the partner server. During this time, the server is unresponsive to DHCP client requests.

**Values** 1 — 10

**sec** *seconds* — Specifies the time in seconds that one server attempts to contact the partner server. During this time, the server is unresponsive to DHCP client requests.

**Values** 1 — 59

## force-renews

**Syntax** **[no] force-renews**

**Context** config>router>dhcp>server

**Description** This command enables the sending of sending forcerenew messages.  
The **no** form of the command disables the sending of forcerenew messages.

**Default** no force-renews

## ignore-rapid-commit

**Syntax** **[no] ignore-rapid-commit**

**Context** config>router>dhcp6>server

**Description** This command enables the Rapid Commit Option.  
The **no** form of the command disables the Rapid Commit Option.

### interface-id-mapping

**Syntax** **[no] interface-id-mapping**

**Context** config>router>dhcp6>server

**Description** If enabled, this command enables the behavior where unique /64 prefix is allocated per interface-id, and all clients having the same interface-id get an address allocated out of this /64 prefix. This is relevant for bridged clients behind the same local-loop (and same SAP), where sharing the same prefix allows communication between bridged clients behind the same local-loop to stay local. For SLAAC based assignment, downstream neighbor-discovery is automatically enabled to resolve the assigned address.

**Default** no interface-id-mapping

### lease-hold-time

**Syntax** **lease-hold-time [days days] [hrs hours] [min minutes] [sec seconds]**  
**no lease-hold-time**

**Context** config>service>vprn>dhcp>server  
config>router>dhcp>server  
config>service>vprn>dhcp6>server  
config>router>dhcp6>server

**Description** This command configures the time to remember this lease. This lease-hold-time is for unsolicited release conditions such as lease timeout and normal solicited release from DHCP client.  
The no form of the command reverts to the default.

**Default** sec 0

**Parameters** [days days][hrs hours] [min minutes] [sec seconds] — Specifies the lease hold time.

Values	
days:	[0..3650]
hours:	[0..23]
minutes:	[0..59]
seconds:	[0..59]

### lease-hold-time-for

**Syntax** **[no] lease-hold-time-for**

**Context** config>service>vprn>dhcp6>server  
config>router>dhcp6>server  
config>service>vprn>dhcp>server  
config>router>dhcp>server



<b>Description</b>	This command enables the context to configure <b>lease-hold-time-for</b> parameters which defines additional types of lease or triggers that cause system to hold up leases. Use the <b>lease-hold-time</b> to enable or disable lease hold up on the server level.
<b>Default</b>	lease-hold-time-for

## internal-lease-ipsec

<b>Syntax</b>	<b>[no] internal-lease-ipsec</b>
<b>Context</b>	config>service>vpn>dhcp6>server>lease-hold-time-for config>router>dhcp6>server> lease-hold-time-for config>service>vpn>dhcp>server config>router>dhcp>server
<b>Description</b>	This command enables the server to hold up the lease of local IPSec clients. The no form of the command disables the server to hold up the lease of local IPSec clients.
<b>Default</b>	no internal-lease-ipsec

## solicited-release

<b>Syntax</b>	<b>[no] solicited-release</b>
<b>Context</b>	config>service>vpn>dhcp6>server>lease-hold-time-for config>router>dhcp6>server> lease-hold-time-for config>service>vpn>dhcp>server config>router>dhcp>server
<b>Description</b>	This command enables server to hold up lease even in case of solicited release. For example, the server receives normal DHCP release message
<b>Default</b>	no solicited-release

## pool

<b>Syntax</b>	<b>pool pool-name [create]</b> <b>no pool pool-name</b>
<b>Context</b>	config>router>dhcp>server
<b>Description</b>	This command configures a DHCP address pool on the router.
<b>Default</b>	none
<b>Parameters</b>	<i>pool name</i> — Specifies the name of this IP address pool. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters.

### exclude-prefix

<b>Syntax</b>	<b>[no] exclude-prefix</b> <i>ipv6-prefix/prefix-length</i>										
<b>Context</b>	config>service>vpn>dhcp6>server>pool config>router>dhcp6>server>pool										
<b>Description</b>	<p>This command defines a prefix that to be excluded from available prefix in the pool. The typical use case is to exclude the interface address.</p> <ul style="list-style-type: none"> <li>• A held lease will be deleted if it got excluded by an exclude prefix.</li> <li>• An exclude range can never exclude only a part of an existing lease. If for example a /63 PD is assigned, an exclude of /64 which belongs to this /63 can NOT be configured.</li> <li>• A single exclude prefix can never exclude a whole include prefix.</li> <li>• When applying or removing an exclude prefix, the threshold stats are adjusted to reflect the actual address space and its usage.</li> </ul>										
<b>Default</b>	none										
<b>Parameters</b>	<i>ipv6-prefix/prefix-length</i> — Specifies an IPv6 prefix and prefix length.										
<b>Values</b>	<table> <tr> <td>ipv6-prefix</td><td>x:x:x:x:x:x:x:x (eight 16-bit pieces)</td></tr> <tr> <td></td><td>x:x:x:x:x:d.d.d.d</td></tr> <tr> <td></td><td>x - [0..FFFF]H</td></tr> <tr> <td></td><td>d - [0..255]D</td></tr> <tr> <td>prefix-length</td><td>- [0..128]</td></tr> </table>	ipv6-prefix	x: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	prefix-length	- [0..128]
ipv6-prefix	x: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										
prefix-length	- [0..128]										

### failover

<b>Syntax</b>	<b>failover</b>
<b>Context</b>	config>router>dhcp>server configure>service>vpn>dhcp>server
<b>Description</b>	This command enables the context to configure failover paramters.

### ignore-mclt-on-takeover

<b>Syntax</b>	<b>[no] ignore-mclt-on-takeover</b>
<b>Context</b>	configure>router>dhcp>server>failover configure>router>dhcp>server>pool>failover configure>router>dhcp6>server>failover configure>router>dhcp6>server>pool>failover configure>service>vpn>dhcp>server>failover configure>service>vpn>dhcp>server>pool>failover configure>service>vpn>dhcp6>server>failover configure>service>vpn>dhcp6>server>pool>failover
<b>Description</b>	With this flag enabled, the remote IP address/prefix can be taken over immediately upon entering the PARTNER-DOWN state of the intercommunication link, without having to wait for the MCLT to

expire. Note that by setting this flag, the lease times of the existing DHCP clients, while the intercommunication link is in the PARTNER-DOWN state, will still be reduced to the MCLT over time and all new lease times will be set to MCLT. This behavior remains the same as originally intended for MCLT.

Some deployments require that the remote IP address/prefix range starts delegating new IP addresses/prefixes upon the failure of the intercommunication link, without waiting for the intercommunication link to transition from the COMM-INT state into the PARTNER-DOWN state and the MCLT to expire while in PARTNER-DOWN state.

This can be achieved by enabling the **ignore-mclt-on-takeover** flag and by configuring the **partner-down-delay** to 0.

Enabling this functionality must be exercised with caution. One needs to keep in mind that the partner-down-delay and MCLT timers were originally introduced to prevent IP address duplication in cases where DHCP redundant nodes transition out-of-sync due to the failure of intercommunication link. These timers (**partner-down-delay** and MCLT) would ensure that during their duration, the new IP addresses/prefixes are delegated only from one node, the one with local IP address-range/prefix. The drawback is of course that the new IP address delegation is delayed and thus service is impacted.

But if one could ensure that the intercommunication link is always available, then the DHCP nodes would stay in sync and the two timers would not be needed. This is why it is of utmost importance that in this mode of operation, the intercommunication link is well protected by providing multiple paths between the two DHCP nodes. The only event that should cause intercommunication link to fail is the entire nodal failure. This failure is acceptable since in this case only one DHCP node is available to provide new IP addresses/prefixes.

**Default** no ignore-mclt-on-takeover

## maximum-client-lead-time

**Syntax** **maximum-client-lead-time** [hrs *hours*] [min *minutes*] [sec *seconds*]  
no maximum-client-lead-time

**Context** configure>router>dhcp>server>failover  
configure>router>dhcp>server>pool>failover  
configure>service>vpn>dhcp>server>failover  
configure>service>vpn>dhcp>server>pool>failover  
configure>router>dhcp6>server>failover  
configure>router>dhcp6>server>pool>failover  
configure>service>vpn>dhcp6>server>failover  
configure>service>vpn>dhcp6>server>pool>failover

**Description** The **maximum-client-lead-time** (MCLT) is the maximum time that a DHCP server can extend client's lease time beyond the lease time currently known by the DHCP partner node. In dual-homed environment, the initial lease time for all DHCP clients is by default restricted to MCLT. Consecutive DHCP renewals are allowed to extend the lease time beyond the MCLT.

The MCLT is a safeguard against IP address/prefix duplication in cases of a lease synchronization failure when local-remote failover model is deployed

Once the intercommunication link failure between the redundant DHCP servers is detected, the DHCP IP address range configured as remote will not be allowed to start delegating new leases until the MCLT + partner-down-delay intervals expire. This is to ensure that the new lease that was delegated from the 'local' IP address-range/prefix on one node, but was never synchronized due to

the intercommunication link failure, will expire before the same IP address/prefix is allocated from the remote IP address-range/prefix on the other node.

However, the already existing (and synchronized) lease times can be renewed from the remote IP address range at any time, regardless of the state of the intercommunication link (operational or failed).

Lease synchronization failure can be caused either by a node failure, or a failure of the link over which the DHCP leases are synchronized (intercommunication link). Synchronization failure detection can take up to 3 seconds.

During the failure, the DHCP lease time for the new clients will be restricted to MCLT while for the existing clients the lease time will over time (by consecutive DHCP renewals) be gradually reduced to the MCLT.

<b>Default</b>	10 minutes
<b>Parameters</b>	<b>hrs</b> <i>hours</i> — Specifies the maximum client lead time in hours. <b>Values</b> 1 — 23 <b>min</b> <i>minutes</i> — Configure the maximum client lead time in minutes. <b>Values</b> 1 — 59 <b>sec</b> <i>seconds</i> — Configure the maximum client lead time in seconds. <b>Values</b> 1 — 59

### partner-down-delay

<b>Syntax</b>	<b>partner-down-delay</b> [ <b>hrs</b> <i>hours</i> ] [ <b>min</b> <i>minutes</i> ] [ <b>sec</b> <i>seconds</i> ] <b>no partner-down-delay</b>
<b>Context</b>	configure>router>dhcp>server>failover configure>router>dhcp>server>pool>failover configure>service>vpn>dhcp>server>failover configure>service>vpn>dhcp>server>pool>failover configure>router>dhcp6>server>failover configure>router>dhcp6>server>pool>failover configure>service>vpn>dhcp6>server>failover configure>service>vpn>dhcp6>server>pool>failover
<b>Description</b>	<p>Since the DHCP lease synchronization failure can be caused by the failure of the intercommunication link (and not necessary the entire node), there is a possibility the redundant DHCP servers become isolated in the network. In other words, they can serve DHCP clients but they cannot synchronize the lease. This can lead to duplicate assignment of IP addresses, since the servers have configured overlapping IP address ranges but they are not aware of each other's leases.</p> <p>The purpose of the partner-down-delay is to prevent the IP lease duplication during the intercommunication link failure by not allowing new IP addresses to be assigned from the remote IP address range. This timer is intended to provide the operator with enough time to remedy the failed situation and to avoid duplication of IP addresses/prefixes during the failure.</p> <p>During the partner-down-delay time, the prefix designated as remote will be eligible only for renewals of the existing DHCP leases that have been synchronized by the peering node. Only after the sum of the partner-down-delay and the maximum-client-lead-time will the prefix designated as</p>

remote be eligible for delegation of the new DHCP leases. When this occurs, we say that the remote IP address range has been taken over.

It is possible to expedite the takeover of a remote IP address range so that the new IP leases can start being delegated from that range shortly after the intercommunication failure is detected. This can be achieved by configuring the partner-down-delay timer to 0 seconds, along with enabling the ignore-melt-on-takeover CLI flag. Caution must be taken before enabling this functionality. It is safe to bypass safety timers (partner-down-delay + MCLT) only in cases where the operator is certain that the intercommunication between the nodes has failed due to the entire node failure and not due to the intercommunication (MCS) link failure. Failed intercommunication due to the nodal failure would ensure that only one node is present in the network for IP address delegation (as opposed to two isolated nodes with overlapping IP address ranges where address duplication can occur). For this reason, the operator must ensure that there are redundant paths between the nodes to ensure uninterrupted synchronization of DHCP leases.

In access-driven mode of operation, partner-down-delay has no effect.

<b>Default</b>	23 hours, 59minutes and 59 seconds
<b>Parameters</b>	<p><b>hrs</b> <i>hours</i> — Specifies the partner-down delay time in hours.</p> <p><b>Values</b> 1 — 23</p> <p><b>min</b> <i>minutes</i> — Configure the partner-down delay time in minutes.</p> <p><b>Values</b> 1 — 59</p> <p><b>sec</b> <i>seconds</i> — Configure the partner-down delay time in seconds.</p> <p><b>Values</b> 1 — 59</p>

## peer

<b>Syntax</b>	<b>peer</b> <i>ip-address</i> <b>tag</b> <i>sync-tag</i> <b>no peer</b>
<b>Context</b>	configure>router>dhcp>server>failover configure>router>dhcp>server>pool>failover configure>service>vprn>dhcp>server>failover configure>service>vprn>dhcp>server>pool>failover configure>router>dhcp6>server>failover configure>router>dhcp6>server>pool>failover configure>service>vprn>dhcp6>server>failover configure>service>vprn>dhcp6>server>pool>failover
<b>Description</b>	<p>DHCP leases can be synchronized per DHCP server of DHCP pool. The pair of synchronizing servers or pools is identified by a tag. The synchronization information is carried over the Multi-Chassis Synchronization (MCS) link between the two peers. MCS link is a logical link (IP, or MPLS).</p> <p>MCS runs over TCP, port 45067 and it is using either data traffic or keepalives to detect failure on the communication link between the two nodes. In the absence of any MCS data traffic for more than 0.5sec, MCS will send its own keepalive to the peer. If a reply is NOT received within 3sec, MCS will declare its operation state as DOWN and the DB Sync state as out-of-sync. MCS will consequently notify its clients (DHCP Server being one of them) of this. It can take up to 3 seconds before the DHCP client realizes that the inter-chassis communication link has failed.</p>

## DHCP Configuration Commands

Note that the inter-chassis communication link failure does not necessarily assume the same failed fate for the access links. In other words the two redundant nodes can become isolated from each other in the network. This would occur in cases where only the intercommunication (MCS) link fails. It is of utmost importance that this MCS link be highly redundant.

<b>Default</b>	none
<b>Parameters</b>	<i>ip-address</i> — Specifies the IPv4 address of the peer. <i>tag</i> — Specifies a tag that will identify synchronizing DHCP servers or pools.

### startup-wait-time

<b>Syntax</b>	<b>[no] startup-wait-time [min <i>minutes</i>] [sec <i>seconds</i>]</b>
<b>Context</b>	configure>router>dhcp>server>failover configure>router>dhcp>server>pool>failover configure>service>vprn>dhcp>server>failover configure>service>vprn>dhcp>server>pool>failover configure>router>dhcp6>server>failover configure>router>dhcp6>server>pool>failover configure>service>vprn>dhcp6>server>failover configure>service>vprn>dhcp6>server>pool>failover
<b>Description</b>	This command enables startup-wait-time during which each peer waits after the initialization process before assuming the active role for the prefix designated as local or access-driven. This is to avoid transient issues during the initialization process.
<b>Default</b>	2 minutes
<b>Parameters</b>	<b>min</b> — Specifies the the startup wait time in minutes. <b>Values</b> 1 — 10 <b>sec</b> — Specifies the the startup wait time in seconds. <b>Values</b> 1 — 59

### max-lease-time

<b>Syntax</b>	<b>max-lease-time [days <i>days</i>] [hrs <i>hours</i>] [min <i>minutes</i>] [sec <i>seconds</i>]</b> <b>no max-lease-time</b>
<b>Context</b>	config>router>dhcp>server>pool
<b>Description</b>	This command configures the maximum lease time. The <b>no</b> form of the command returns the value to the default.
<b>Default</b>	10 days
<b>Parameters</b>	<i>time</i> — Specifies the maximum lease time. <b>Values</b> days : 0 — 3650 hours 0 — 23

minutes:	0 — 59
seconds	0 — 59

## min-lease-time

**Syntax** **min-lease-time** [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]  
**no min-lease-time**

**Context** config>router>dhcp>server>pool

**Description** This command configures the minimum lease time.  
 The **no** form of the command returns the value to the default.

**Default** 10 minutes

**Parameters** *time* — Specifies the minimum lease time.

**Values**

days :	0 — 3650
hours	0 — 23
minutes:	0 — 59
seconds	0 — 59

## minimum-free

**Syntax** **minimum-free** *minimum-free* [**percent**] [**event-when-depleted**]  
**no minimum-free**

**Context** config>router>dhcp>server>pool

**Description** This command specifies the desired minimum number of free addresses in this pool.  
 The **no** form of the command reverts to the default.

**Default** 1

**Parameters** *minimum-free* — Specifies the minimum number of free addresses.  
 0 — 255

**percent** — Specifies that the value indicates a percentage.

**event-when-depleted** — This parameter enables a system-generate event when all available addresses in the pool/subnet of local DHCP server are depleted.

## nak-non-matching-subnet

**Syntax** [**no**] **nak-non-matching-subnet**

**Context** config>service>vprn>dhcp>server>pool  
 config>router>dhcp>server>pool

**Description** With this command, if the local DHCPv4 server receives a DHCP request with option 50 (means client try to request a previous allocated message as described in section 3.2 of RFC 2131, *Dynamic*

## DHCP Configuration Commands

*Host Configuration Protocol*) and the address allocation algorithm ends up using a pool and the address in option50 is not in pool, then system will return a DHCP NAK, otherwise system just drop the DHCP packet.

**Default** no nak-non-matching-subnet

### offer-time

**Syntax** **offer-time** [**min** *minutes*] [**sec** *seconds*]  
**no offer-time**

**Context** config>router>dhcp>server>pool

**Description** This command configures the offer time.  
The **no** form of the command returns the value to the default.

**Default** 1 minute

**Parameters** *time* — Specifies the offer time.

<b>Values</b>	minutes:	0 — 10
	seconds	0 — 59

### msap-defaults

**Syntax** **msap-default**

**Context** config>sub-mgmt>lu-db>dhcp>hos  
config>sub-mgmt>lu-db>ipoe>host  
config>sub-mgmt>lu-db>ppp>host

**Description** This command configures MSAP authentication defaults.

### group-interface

**Syntax** **group-interface** *ip-int-name* [**prefix** {*port-id*}]  
**group-interface** *ip-int-name* [**prefix** {*port-id*}]  
**group-interface** *ip-int-name* [**suffix** {*port-id*}]  
**no group-interface**

**Context** config>sub-mgmt>lu-db>dhcp>host  
config>subscr-mgmt>loc-user-db>ipoe>host>msap-defaults  
config>sub-mgmt>lu-db>ppp>host

**Description** This command configures the group interface.

**Parameters** *ip-int-name* — Specifies the IP interface name.

<b>Values</b>	32 chars max (must start with a letter)
---------------	---

**Parameters** **prefix** {*port-id*} — Specifies the port ID as the prefix to the specified ip-int-name.



**suffix** *{port-id}* — Specifies the port ID as the suffix to the specified ip-int-name.

## policy

<b>Syntax</b>	<b>policy</b> <i>msap-policy-name</i> <b>no policy</b>
<b>Context</b>	config>sub-mgmt>lu-db>dhcp>host config>subscr-mgmt>loc-user-db>ipoe>host>msap-defaults config>sub-mgmt>lu-db>ppp>host
<b>Description</b>	This command configures the MSAP policy.
<b>Parameters</b>	<i>msap-policy-name</i> — Specifies the policy name.

## service

<b>Syntax</b>	<b>service</b> <i>service-id</i> <b>no service</b>
<b>Context</b>	config>sub-mgmt>lu-db>dhcp>host config>subscr-mgmt>loc-user-db>ipoe>host>msap-defaults config>sub-mgmt>lu-db>ppp>host
<b>Description</b>	This command sets retail-service for a given subscriber host.
<b>Parameters</b>	<i>service-id</i> — Specifies the service ID as an interger.
<b>Values</b>	1-2147483648

## retail-service

<b>Syntax</b>	<b>[no] retail-service</b> <i>service-id</i>
<b>Context</b>	config>sub-mgmt>lu-db>dhcp>hos config>sub-mgmt>lu-db>ppp>host
<b>Description</b>	This command sets default service for all subscribers created based on trigger packets received on the given capture SAP in case the corresponding VSA is not included in the RADIUS authentication response. This command is applicable to capture SAP only.
<b>Default</b>	no retail-service

## options

<b>Syntax</b>	<b>options</b>
<b>Context</b>	config>router>dhcp>local-dhcp-serve>pool config>router>dhcp>local-dhcp-serve>pool>subnet

## DHCP Configuration Commands

```
config>subscr-mgmt>loc-user-db>dhcp>host
config>subscr-mgmt>loc-user-db>ppp>host
```

**Description** This command enables the context to configure pool options. The options defined here can be overruled by defining the same option in the local user database.

**Default** none

### custom-option

**Syntax** **custom-option** *option-number* **address** [*ip-address...*(up to 4 max)]  
**custom-option** *option-number* **hex** *hex-string*  
**custom-option** *option-number* **string** *ascii-string*  
**no custom-option** *option-number*

**Context** config>router>dhcp>local-dhcp-serve>pool>options  
config>router>dhcp>local-dhcp-serve>pool>subnet>options  
config>subscr-mgmt>loc-user-db>dhcp>host>options  
config>subscr-mgmt>loc-user-db>ppp>host>options

**Description** This command configures specific DHCP options. The options defined here can overrule options in the local user database.

The **no** form of the removes the option from the configuration.

**Default** none

**Parameters** *option-number* — specifies the option number that the DHCP server uses to send the identification strings to the DHCP client.

**Values** 1 — 254

**address** *ip-address* — Specifies the IP address of this host.

**hex** *hex-string* — Specifies the hex value of this option.

**Values** 0x0..0xFFFFFFFF...(maximum 254 hex nibbles)

**string** *ascii-string* — Specifies the value of this option.

**Values** Up to 127 characters maximum.

### dns-server

**Syntax** **dns-server** **address** [*ip-address...*(upto 4 max)]  
**no dns-server**

**Context** config>router>dhcp>server>pool>options  
config>subscr-mgmt>loc-user-db>dhcp>host>options  
config>subscr-mgmt>loc-user-db>ipoe>host>options  
config>subscr-mgmt>loc-user-db>ppp>host>options

**Description** This command configures the IP address of the DNS server.

**Default** none

**Parameters** *ipv6-address* — The IPv4 address of the DNS server. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).

## dns-server

**Syntax** **dns-server** *ipv6-address* [*ipv6-address*...(up to 4 max)]  
**no dns-server**

**Context** config>subscr-mgmt>loc-user-db>ppp>host>options6  
 config>subscr-mgmt>loc-user-db>dhcp>host>options6

**Description** Configure IPv6 DNS server addresses that can be used for name resolution

**Default** no dns-server

**Parameters** *ipv6-address* — - IPv6 address of the a DNS server.

## domain-name

**Syntax** **domain-name** *domain-name*  
**no domain-name**

**Context** config>router>dhcp>server>pool>options  
 config>subscr-mgmt>loc-user-db>dhcp>host>options  
 config>subscr-mgmt>loc-user-db>ipoe>host>options

**Description** This command configures the default domain for a DHCP client that the router uses to complete unqualified hostnames (without a dotted-decimal domain name).  
 The **no** form of the command removes the name from the configuration.

**Default** none

**Parameters** *domain-name* — Specifies the domain name for the client.

**Values** Up to 127 characters

## lease-rebind-time

**Syntax** **lease-rebind-time** [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]  
**no lease-rebind-time**

**Context** config>router>dhcp>server>pool>subnet>options  
 config>subscr-mgmt>loc-user-db>dhcp>host>options  
 config>subscr-mgmt>loc-user-db>ipoe>host>options

**Description** This command configures the time the client transitions to a rebinding state.  
 The **no** form of the command removes the time from the configuration.

**Default** none

### Parameters

<b>Parameters</b>	<i>time</i> — Specifies the lease rebind time.		
<b>Values</b>	days:	0 — 3650	
	hours:	0 — 23	
	minutes:	0 — 59	
	seconds	0 — 59	

## lease-renew-time

<b>Syntax</b>	<b>lease-renew-time</b> [ <b>days</b> <i>days</i> ] [ <b>hrs</b> <i>hours</i> ] [ <b>min</b> <i>minutes</i> ] [ <b>sec</b> <i>seconds</i> ] <b>no lease-renew-time</b>
<b>Context</b>	config>router>dhcp>server>pool>options config>subscr-mgmt>loc-user-db>dhcp>host>options config>subscr-mgmt>loc-user-db>ipoe>host>options
<b>Description</b>	This command configures the time the client transitions to a renew state. The <b>no</b> form of the command removes the time from the configuration.
<b>Default</b>	none
<b>Parameters</b>	<i>time</i> — Specifies the lease renew time.
<b>Values</b>	days: 0 — 3650 hours: 0 — 23 minutes: 0 — 59 seconds 0 — 59

## lease-time

<b>Syntax</b>	<b>lease-time</b> [ <b>days</b> <i>days</i> ] [ <b>hrs</b> <i>hours</i> ] [ <b>min</b> <i>minutes</i> ] [ <b>sec</b> <i>seconds</i> ] <b>no lease-time</b>
<b>Context</b>	config>router>dhcp>server>pool>options config>subscr-mgmt>loc-user-db>dhcp>host>options
<b>Description</b>	This command configures the amount of time that the DHCP server grants to the DHCP client permission to use a particular IP address. The <b>no</b> form of the command removes the lease time parameters from the configuration.
<b>Default</b>	none
<b>Parameters</b>	<i>time</i> — Specifies the lease time.
<b>Values</b>	days : 0 — 3650 hours 0 — 23 minutes: 0 — 59 seconds 0 — 59

## netbios-name-server

<b>Syntax</b>	<b>netbios-name-server ip-address</b> [ <i>ip-address...</i> (up to 4 max)] <b>no netbios-name-server</b>
<b>Context</b>	config>router>dhcp>server>pool>options config>subscr-mgmt>loc-user-db>dhcp>host>options config>subscr-mgmt>loc-user-db>ppp>host>options config>subscr-mgmt>loc-user-db>ipoe>host>options
<b>Description</b>	This command configures up to four Network Basic Input/Output System (NetBIOS) name server IP addresses.
<b>Default</b>	none
<b>Parameters</b>	<i>ip-address</i> — The IP address of the NetBIOS name server. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).

## netbios-node-type

<b>Syntax</b>	<b>netbios-node-type netbios-node-type</b> <b>no netbios-node-type</b>		
<b>Context</b>	config>router>dhcp>server>pool>options config>subscr-mgmt>loc-user-db>dhcp>host>options config>subscr-mgmt>loc-user-db>ipoe>host>options		
<b>Description</b>	This command configures the Network Basic Input/Output System (NetBIOS) node type.		
<b>Default</b>	none		
<b>Parameters</b>	<i>netbios-node-type</i> — Specifies the netbios node type.  <table> <tr> <td><b>Values</b></td><td>           B — Broadcast node uses broadcasting to query nodes on the network for the owner of a NetBIOS name.            P — Peer-to-peer node uses directed calls to communicate with a known NetBIOS name server for the IP address of a NetBIOS machine name.            M — Mixed node uses broadcasted queries to find a node, and if that fails, queries a known P-node name server for the address.            H — Hybrid node is the opposite of the M-node action so that a directed query is executed first, and if that fails, a broadcast is attempted.         </td></tr> </table>	<b>Values</b>	B — Broadcast node uses broadcasting to query nodes on the network for the owner of a NetBIOS name. P — Peer-to-peer node uses directed calls to communicate with a known NetBIOS name server for the IP address of a NetBIOS machine name. M — Mixed node uses broadcasted queries to find a node, and if that fails, queries a known P-node name server for the address. H — Hybrid node is the opposite of the M-node action so that a directed query is executed first, and if that fails, a broadcast is attempted.
<b>Values</b>	B — Broadcast node uses broadcasting to query nodes on the network for the owner of a NetBIOS name. P — Peer-to-peer node uses directed calls to communicate with a known NetBIOS name server for the IP address of a NetBIOS machine name. M — Mixed node uses broadcasted queries to find a node, and if that fails, queries a known P-node name server for the address. H — Hybrid node is the opposite of the M-node action so that a directed query is executed first, and if that fails, a broadcast is attempted.		

## prefix

<b>Syntax</b>	<b>prefix ipv6-addr/prefix-len</b> [failover {local   remote}] [pd] [wan-host] [create] <b>no prefix ipv6-addr/prefix-len</b>
<b>Context</b>	configure>router>dhcp6>server>pool configure>service>vpn>dhcp6>server>pool
<b>Description</b>	This is an existing command and we just need to add the failover option.

## DHCP Configuration Commands

<b>Default</b>	failover local
<b>Parameters</b>	<i>ipv6-addr/prefix-len</i> —
<b>Values</b>	ipv6-address    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 prefix-length [1..128]
	<b>failover {local   remote}</b> — This command designates a prefix as local or remote. This is used when multi-chassis synchronization is enabled.
<b>Values</b>	<b>local</b> — A prefix designated as local is always used to renew the existing addresses/prefixes or to assign a new one. <b>remote</b> — A prefix designated as remote is used only to renew the existing DHCP leases. The new leases will be assigned from it only after the maximum-client-lead-time + partner-down-delay time elapses.

## thresholds

<b>Syntax</b>	<b>thresholds</b>
<b>Context</b>	config>service>vprn>dhcp6>server>pool config>router>dhcp6>server>pool
<b>Description</b>	This command enables the context to configure pool level thresholds.
<b>Default</b>	thresholds

## thresholds

<b>Syntax</b>	<b>thresholds</b>
<b>Context</b>	config>service>vprn>dhcp6>server>pool>prefix config>router>dhcp6>server>pool>prefix
<b>Description</b>	This command enables the context to configure prefix level thresholds.
<b>Default</b>	thresholds

## minimum-free

<b>Syntax</b>	[no] <b>minimum-free prefix-length</b> [1..128]
<b>Context</b>	config>service>vprn>dhcp6>server>pool>thresholds config>router>dhcp6>server>pool>thresholds
<b>Description</b>	This command creates a threshold for a given prefix length on the pool level. Up to 128 thresholds could be created. For example, with <b>minimum-free prefix-length 64</b> , then the usage of /64 prefix in the pool is counted.

There are two types of thresholds could be defined on pool level:

- Depleted — The system sends out a warning when the prefix with the configured length is no long available in the pool.
- Minimum free — A percentage-based threshold which represents the minimal available percentage of prefix with the configured length in the pool. The system will send out warning if the actual percentage is lower than the configured percentage

Configuration of this command also enables the system stats collection for **configure prefix length**, which could be displayed via the **show router <router-id>dhcp6 local-dhcp-server "d6" pool-threshold-stats** command.

<b>Default</b>	none
<b>Parameters</b>	<b>1..128</b> — Specifies the IPv6 prefix length.

## minimum-free

<b>Syntax</b>	<b>[no] minimum-free prefix-length [1..128]</b>
<b>Context</b>	config>service>vprn>dhcp6>server>pool>prefix>thresholds config>router>dhcp6>server>pool>>prefix>thresholds
<b>Description</b>	This command creates a threshold for a given prefix length on the prefix level. Up to 128 thresholds could be created. For example, with <b>minimum-free prefix-length 64</b> , then the usage of /64 prefix in the prefix is counted.  There are two types of thresholds could be defined on pool level: <ul style="list-style-type: none"> <li>• Depleted — The system sends out a warning when the prefix with the configured length is no long available in the provisioned prefix.</li> <li>• Minimum free — A percentage or number based threshold which represent the minimal available percentage or number of the prefix with configured length in the provisioned prefix. The system will send out warning if the actual percentage is lower than the configured percentage</li> </ul> Configuration of this command also enables the system stats collection for <b>configure prefix length</b> , which can be displayed with the <b>show router &lt;router-id&gt;dhcp6 local-dhcp-server "d6" prefix-threshold-stats</b> command.
<b>Default</b>	none
<b>Parameters</b>	<b>1..128</b> — Specifies the IPv6 prefix length.

## depleted-event

<b>Syntax</b>	<b>[no] depleted-event</b>
<b>Context</b>	config>service>vprn>dhcp6>server>pool>thresholds>minimum-free config>router>dhcp6>server> pool>thresholds>minimum-free
<b>Description</b>	This command enables the system to send out warnings when the prefix with the configured length is no long available in the pool.
<b>Default</b>	none

## depleted-event

<b>Syntax</b>	<b>[no] depleted-event</b> config>service>vprn>dhcp6>server>pool>prefix>thresholds>minimum-free config>router>dhcp6>server> pool>prefix>thresholds>minimum-free
<b>Description</b>	This command enables the system to send out a warning when the prefix with a configured length is no long available in the provisioned prefix.  For example:  <pre>prefix 2001:0:0:ffe0::/50 pd wan-host create   thresholds     minimum-free prefix-length 64     depleted-event</pre> <p>With the above configuration, the system will send out a warning when there is no available /64 that can be allocated out of 2001:0:0:ffe0::/50.</p>
<b>Default</b>	none

## minimum

<b>Syntax</b>	<b>minimum percent [0..100]</b> <b>no minimum</b>
<b>Context</b>	config>service>vprn>dhcp6>server>pool>thresholds>minimum-free config>router>dhcp6>server> pool>thresholds>minimum-free
<b>Description</b>	This command specifies a percentage based threshold which represent the minimal available percentage of the prefix with configured length in the pool. The system will send out a warning if the actual percentage is lower than the configured percentage.
<b>Default</b>	none
<b>Parameters</b>	<b>percent [0..100]</b> — Specifies the percentage of used prefixes with the minimum free threshold length in the pool compared to the number of provisioned prefixes.

## minimum

<b>Syntax</b>	<b>minimum [percent [0..100]] [number [0..4294967295]]</b> <b>no minimum</b>
<b>Context</b>	config>service>vprn>dhcp6>server>pool>prefix>thresholds>minimum-free config>router>dhcp6>server> pool>prefix>thresholds>minimum-free
<b>Description</b>	This command configures a percentage-based or number-based threshold which represents the minimal available percentage or number of the prefix with a configured length in the provisioned prefix. The system will send out a warning if the actual percentage or number is lower than the configured threshold.  For example:  <pre>prefix 2001:0:0:ffe0::/50 pd wan-host create   thresholds</pre>



```

minimum-free prefix-length 64
minimum number 3

```

With the above configuration, the system will send a warning when the number of available /64 in prefix 2001:0:0:ffe0::/50 is less than 3.

<b>Default</b>	none
<b>Parameters</b>	<p><b>percent</b> [0..100] — Specifies the percentage of used prefixes with the minimum free threshold length in the pool compared to the number of provisioned prefixes.</p> <p><b>number</b> [0..4294967295] — Specifies the number of prefixes.</p>

## to-client-options

<b>Syntax</b>	<b>to-client-options</b>
<b>Context</b>	config>subscr-mgmt>loc-user-db>ipoe>host
<b>Description</b>	This command configures the DHCP options to send to the client.

## option

<b>Syntax</b>	<p><b>option</b> <i>option-number</i> <b>address</b> [<i>ip-address...</i>(up to 4 max)]</p> <p><b>option</b> <i>option-number</i> <b>hex</b> <i>hex-string</i></p> <p><b>option</b> <i>option-number</i> <b>string</b> <i>ascii-string</i></p> <p><b>no option</b> <i>option-number</i></p>
<b>Context</b>	<p>config&gt;router&gt;dhcp&gt;local-dhcp-serve&gt;pool&gt;options</p> <p>config&gt;router&gt;dhcp&gt;local-dhcp-serve&gt;pool&gt;subnet&gt;options</p> <p>config&gt;subscr-mgmt&gt;loc-user-db&gt;dhcp&gt;host&gt;options</p> <p>config&gt;subscr-mgmt&gt;loc-user-db&gt;ppp&gt;host&gt;options</p> <p>config&gt;subscr-mgmt&gt;loc-user-db&gt;ipoe&gt;host&gt;to-client-options&gt;ipv4</p> <p>config&gt;subscr-mgmt&gt;loc-user-db&gt;ipoe&gt;host&gt;to-client-options&gt;ipv6</p>
<b>Description</b>	<p>This command configures specific DHCP options. The options defined here can overrule options in the local user database.</p> <p>The <b>no</b> form of the removes the option from the configuration.</p>
<b>Default</b>	none
<b>Parameters</b>	<p><i>option-number</i> — specifies the option number that the DHCP server uses to send the identification strings to the DHCP client.</p> <p><b>Values</b> 1 — 254</p> <p><b>address</b> <i>ip-address</i> — Specifies the IP address of this host.</p> <p><b>hex</b> <i>hex-string</i> — Specifies the hex value of this option.</p> <p><b>Values</b> 0x0..0xFFFFFFFF...(maximum 254 hex nibbles)</p> <p><b>string</b> <i>ascii-string</i> — Specifies the value of this option.</p> <p><b>Values</b> Up to 127 characters maximum.</p>

option

<b>Syntax</b>	<b>option</b> <i>option-number</i> <b>address</b> <i>ipv6-address</i> [ <i>ipv6-address...</i> (upto 4 max)] <b>option</b> <i>option-number</i> <b>hex</b> <i>hex-string</i> <b>option</b> <i>option-number</i> <b>string</b> <i>ascii-string</i> <b>no option</b> <i>option-number</i>
<b>Context</b>	configure>subscr-mgmt>loc-user-db>ipoe>host>to-client-options>dhcpv6 configure>subscr-mgmt>loc-user-db>ppp>host>to-client-options>dhcpv6
<b>Description</b>	This command configures DHCPv6 options via LUDB that will be passed in all DHCP messages to the client. The options will be blindly appended to any options already present in the DHCP message. In other words, there is no intelligent merging of the options where overlapping options from different sources are further evaluated to determine whether only one option or multiple options should be returned to the client.

Multiple DHCP options can be configured at the same time although each option requires its own option statement. Those options are equivalent to RADIUS VSAs **Ale-ToClient-Dhcp6-Options**.


DHCPv6 options can be provided via DHCPv6 server in the relay case. In addition, DHCPv6 options provided via LUDB or RADIUS can be supplied and consequently appended to the already existing options. In case that DHCPv6 options are provided simultaneously via LUDB and RADIUS, the RADIUS as a source of DHCPv6 option will be blocked and the options supplied via LUDB will be passed to the client. This is valid for the relay and proxy case.

Any DHCP option may be encoded in the option statement. An example of the option statement for DHCPv6 DNS servers is given below:

```
option 23 2001:db8::1 2001:db8::2.
```


Options are stored serially in the options field of DHCP message header, with no padding between the options. Options are byte-aligned but are not aligned in any other way such as on 2 or 4 byte boundaries.

The format of DHCPv6 options is:



- option-code — An unsigned integer identifying the specific option type carried in this option.
- option-len — An unsigned integer giving the length of the option-data field in this option in octets.
- option-data —The data for the option; the format of this data depends on the definition of the option.

DHCPv6 options are scoped by using encapsulation. Some options apply generally to the client, some are carried with other options, such as IA-NA:



option-code — OPTION\_IA\_NA (3).

option-len — 12 + length of IA\_NA-options field.

IAID — The unique identifier for this IA\_NA; the IAID must be unique among the identifiers for all of this client's IA\_NAs. The number space for IA\_NA IAIDs is separate from the number space for IA\_TA IAIDs.

T1 — The time at which the client contacts the server from which the addresses in the IA\_NA were obtained to extend the lifetimes of the addresses assigned to the IA\_NA; T1 is a time duration relative to the current time expressed in units of seconds.

T2 — The time at which the client contacts any available server to extend the lifetimes of the addresses assigned to the IA\_NA; T2 is a time duration relative to the current time expressed in units of seconds.

IA\_NA-options — Options associated with this IA\_NA.

**Default** no option

**Parameters**

*option-number* — Specifies the number of the option. This can be a well known option (some of which are defined in RFC 3315, *Dynamic Host Configuration Protocol for IPv6 (DHCPv6)*), or an anonymous option.

**address** *ipv6-address* — Specifies IPv6 address as an option.

**hex** *hex-string* — Specifies options coded as Hex characters.

**string** *ascii-string* — Specifies options coded as string.

## option

**Syntax**

**option** *option-number* **address** *ipv4-address* [*ipv4-address*...(upto 4 max)]

**option** *option-number* **hex** *hex-string*

**option** *option-number* **string** *ascii-string*

**no option** *option-number*

**Context**

configure>subscr-mgmt>loc-user-db>ipoe>host>to-client-options>dhcpv4

configure>subscr-mgmt>loc-user-db>ppp>host>to-client-options>dhcpv4

**Description**

This command configures DHCPv4 options via LUDB that will be passed in all DHCP messages to the client. The options will be blindly appended to any options already present in the DHCP message.

In other words, there is no intelligent merging of the options where overlapping options from different sources are further evaluated to determine whether only one option or multiple options should be returned to the client.

Multiple DHCP options can be configured at the same time although each option requires its own option statement. Those options are equivalent to RADIUS VSAs **Aic-ToClient-Dhcp4-Options**.

DHCPv4 options can be provided via DHCPv4 server in the relay case. In addition, DHCPv4 options provided via LUDB or RADIUS can be supplied and consequently appended to the already existing options. In case that DHCPv4 options are provided simultaneously via LUDB and RADIUS, the RADIUS as a source of DHCPv4 option will be blocked and the options supplied via LUDB will be passed to the client. This is valid for the relay and proxy case.

Any DHCP option may be encoded in the option statement. An example of the option statement for DHCPv4 default-gateway is given below:

```
option 3 192.168.1.254
```

DHCPv4 options may be fixed length or variable length. They are appended at the end of DHCPv4 messages. All options begin with a tag octet, which uniquely identifies the option. Fixed-length options without data consist of only a tag octet. Only options 0 and 255 are fixed length. All other options are variable-length.

<b>Default</b>	no option
<b>Parameters</b>	<i>option-number</i> — Number of the option. This can be a well known option, or a an anonymous option. <i>address ipv4-address</i> — Specifies IPv4 address as an option. <i>hex hex-string</i> — Specifies options coded as Hex characters. <i>string ascii-string</i> — Specifies options coded as string.

## subnet

<b>Syntax</b>	<b>subnet</b> { <i>ip-address/mask</i>   <i>ip-address netmask</i> } [ <b>create</b> ] <b>no subnet</b> { <i>ip-address/mask</i>   <i>ip-address netmask</i> }
<b>Context</b>	config>router>dhcp>server>pool
<b>Description</b>	This command creates a subnet of IP addresses to be served from the pool. The subnet cannot include any addresses that were assigned to subscribers without those addresses specifically excluded. When the subnet is created no IP addresses are made available until a range is defined.
<b>Default</b>	none
<b>Parameters</b>	<i>ip-address</i> — Specifies the base IP address of the subnet. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets). <i>mask</i> — The subnet mask in dotted decimal notation. Allowed values are dotted decimal addresses in the range 128.0.0.0 – 255.255.255.252. Note that a mask of 255.255.255.255 is reserved for system IP addresses. <i>netmask</i> — Specifies a string of 0s and 1s that mask or screen out the network part of an IP address so that only the host computer part of the address remains.

## address-range

<b>Syntax</b>	<b>[no] address-range</b> <i>start-ip-address end-ip-address</i> [ <b>failover</b> { <b>local</b>   <b>remote</b> }]
<b>Context</b>	config>router>dhcp>server>pool>subnet
<b>Description</b>	This command configures a range of IP addresses to be served from the pool. All IP addresses between the start and end IP addresses will be included (other than specific excluded addresses).
<b>Default</b>	none
<b>Parameters</b>	<p><i>start-ip-address</i> — Specifies the start address of this range to include. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).</p> <p><i>end-ip-address</i> — Specifies the end address of this range to include. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).</p> <p><b>failover local</b> — Specifies that the DHCP server failover control type is in control under normal operation.</p> <p><b>failover remote</b> — Specifies that the remote DHCP server failover system is in control under normal operation.</p>

## drain

<b>Syntax</b>	<b>[no] drain</b>
<b>Context</b>	config>service>vprn>dhcp>server>pool>subnet
<b>Description</b>	<p>This command subnet draining which means no new leases can be assigned from this subnet and existing leases are cleaned up upon renew/rebind.</p> <p>The <b>no</b> form of the command means the subnet is active and new leases can be assigned from it.</p>

## exclude-addresses

<b>Syntax</b>	<b>[no] exclude-addresses</b> <i>start-ip-address</i> [ <i>end-ip-address</i> ]
<b>Context</b>	config>router>dhcp>server>pool>subnet
<b>Description</b>	This command specifies a range of IP addresses that excluded from the pool of IP addresses in this subnet.
<b>Default</b>	none
<b>Parameters</b>	<p><i>start-ip-address</i> — Specifies the start address of this range to exclude. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).</p> <p><i>end-ip-address</i> — Specifies the end address of this range to exclude. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).</p>

### maximum-declined

<b>Syntax</b>	<b>maximum-declined</b> <i>maximum-declined</i> <b>no maximum-declined</b>
<b>Context</b>	config>router>dhcp>server>pool>subnet
<b>Description</b>	This command configures the maximum number of declined addresses allowed.
<b>Default</b>	64
<b>Parameters</b>	<i>maximum-declined</i> — Specifies the maximum number of declined addresses allowed. <b>Values</b> 0 — 4294967295

### minimum-free

<b>Syntax</b>	<b>minimum-free</b> <i>minimum-free</i> [ <b>percent</b> ] [ <b>event-when-depleted</b> ] <b>no minimum-free</b>
<b>Context</b>	config>router>dhcp>server>pool>subnet
<b>Description</b>	This command configures the minimum number of free addresses in this subnet. If the actual number of free addresses in this subnet falls below this configured minimum, a notification is generated.
<b>Default</b>	1
<b>Parameters</b>	<i>minimum-free</i> — Specifies the minimum number of free addresses in this subnet. <b>Values</b> 0 — 255 <b>percent</b> — Specifies that the value indicates a percentage. <b>event-when-depleted</b> — This parameter enables a system-generate event when all available addresses in the pool/subnet of local DHCP server are depleted.

### default-router

<b>Syntax</b>	<b>default-router</b> <i>ip-address</i> [ <i>ip-address...</i> (up to 4 max)] <b>no default-router</b>
<b>Context</b>	config>router>dhcp>server>pool>subnet>options config>subscr-mgmt>loc-user-db>ipoe>host>options
<b>Description</b>	This command configures the IP address of the default router for a DHCP client. Up to four IP addresses can be specified. The <b>no</b> form of the command removes the address(es) from the configuration.
<b>Default</b>	none
<b>Parameters</b>	<i>ip-address</i> — Specifies the IP address of the default router. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).

## subnet-mask

<b>Syntax</b>	<b>subnet-mask</b> <i>ip-address</i> <b>no subnet-mask</b>
<b>Context</b>	config>router>dhcp>local-dhcp-serve>pool>subnet>options config>subscr-mgmt>loc-user-db>dhcp>host>options config>subscr-mgmt>loc-user-db>ipoe>host>options
<b>Description</b>	This command specifies the subnet-mask option to the client. The mask can either be defined (for supernetting) or taken from the pool address.  The <b>no</b> form of the command removes the address from the configuration.
<b>Default</b>	none
<b>Parameters</b>	<i>ip-address</i> — Specifies the IP address of the subnet mask. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).

## subnet-binding key

<b>Syntax</b>	<b>subnet-binding key</b> [ <b>sys-id-svc-id</b>   <b>sys-id</b>   <b>string</b> ] <b>unbind-delay</b> [ <b>hrs</b> <i>hours</i> ] [ <b>min</b> <i>mins</i> ] [ <b>sec</b> <i>secs</i> ] <b>no subnet-binding key</b>
<b>Context</b>	config>router>dhcp>local-dhcp-server>pool config>service>vprn>dhcp>local-dhcp-server>pool
<b>Description</b>	The command enables the pool to bind three selectable parameters, <b>sys-id-svc-id</b> , <b>sys-id</b> , or a <b>string</b> to a subnet. These parameters are retrieved from DHCP relay Option 82 vendor specific option (VSO). The intent of this feature is to allow multiple BNG to share a DHCP pool. When a subnet is bound to a VSO, only DHCP discoveries with matching VSO are allowed to allocate additional DHCP addresses. For example, if <b>sys-id</b> is the chosen VSO, a DHCP discovery will bind the <b>sys-id</b> to a subnet. Only DHCP discoveries with matching <b>sys-id</b> are allowed to allocate additional addresses from the same subnet. If a DHCP discovery fails to match any bindings, and if a new subnet is still available, it will first bind the VSO to the new subnet and offer the subscriber an IP address.  Once all addresses are released back to the pool, the subnet is once again available for binding after the unbind-delay has expired. The unbind-delay expiration is to hold the subnet for a small period of time until the subnet has successful remove itself from the routing table. The delay is configurable to allow enough time for routing update to occur. By default, the delay is 5 minute with a minimal required value of 1 second.
<b>Default</b>	<b>unbind-delay min 5</b>
<b>Parameters</b>	<i>key</i> — The desire key to which the subnet to bind: <b>sys-id-svc-id</b>   <b>sys-id</b>   <b>string</b> <i>hours</i> — [0 — 24] the delay for the subnet to unbind in hours. <i>minutes</i> — [0 — 59] the delay for the subnet to unbind in minutes. <i>seconds</i> — [0 — 59] the delay for the subnet to unbind in seconds.

### use-gi-address

<b>Syntax</b>	<b>use-gi-address</b> [ <b>scope</b> <i>scope</i> ]
<b>Context</b>	config>router>dhcp>server
<b>Description</b>	This command enables the use of gi-address matching. By default, the scope is subnet and addresses are allocated only from the subnet where the gi-address belongs, even if the pool contains multiple other subnets. When the scope is pool, addresses are allocated from any subnet in the pool that contains the subnet where the gi-address belongs.
<b>Default</b>	no use-gi-address
<b>Parameters</b>	<p><b>scope</b> <i>scope</i> — Specifies if addresses are handed out for a certain subnet where the gi-address belongs to only or for all subnets part of the pool.</p> <p><b>Values</b></p> <p><b>subnet</b> — Addresses are only handed out for the subnet where the gi-address is part of</p> <p><b>pool</b> — All subnets part of the pool which contain subnet where the gi-address is part of can hand out addresses.</p>

### use-pool-from-client

<b>Syntax</b>	<b>use-pool-from-client</b> <i>delimiter delimiter</i> <b>use-pool-from-client</b> <b>no use-pool-from-client</b>
<b>Context</b>	config>router>dhcp>server
<b>Description</b>	This command enables the use of the pool indicated by DHCP client. When enabled, the IP address pool to be used by this server is the pool is indicated by the vendor-specific sub-option 13 of the DHCP option 82. When disabled or if there is no sub-option 13 in the DHCP message, the pool selection falls back to the “use-gi-address” configuration.
<b>Default</b>	no use-pool-from-client
<b>Parameters</b>	<b>delimiter</b> <i>delimiter</i> — A single ASCII character specifies the delimiter of separating primary and secondary pool names in Option82 VSO.

### user-ident

<b>Syntax</b>	<b>user-ident</b> <i>user-ident</i> <b>no user-ident</b>
<b>Context</b>	config>router>dhcp>local-dhcp-server config>service>vprn>dhcp>server
<b>Description</b>	This command configures the keys for identification of the DHCPv4 lease being held in the lease-database (for configured period after lease timeout). Subscriber requesting a lease via DHCPv4 that matches an existing lease based on this configured key is handed the matched prefix or address. This allows address and prefix “stickiness” for DHCPv4 assigned prefixes (IA_NA or PD).
<b>Default</b>	duid



<b>Parameters</b>	<i>user-ident</i> — Specifies the the user identification method
<b>Values</b>	<b>duid</b> — Specifies the IPv4 DHCP unique identifier from DHCPv4. <b>interface-id</b> — Specifies the IPv4 interface-id option. <b>interface-id-link-local</b> — Specifies the interface-id and link-local address.

## user-ident

<b>Syntax</b>	<b>user-ident</b> <i>user-ident</i> <b>no user-ident</b>
<b>Context</b>	config>router>dhcp6>local-dhcp-server config>service>vprn>dhcp6>server
<b>Description</b>	This command configures the keys for identification of the DHCPv6 lease being held in the lease-database (for configured period after lease timeout). Subscriber requesting a lease via DHCPv6 that matches an existing lease based on this configured key is handed the matched prefix or address. This allows address and prefix “stickiness” for DHCPv6 assigned prefixes (IA_NA or PD).
<b>Default</b>	duid
<b>Parameters</b>	<i>user-ident</i> — Specifies the the user identification method
<b>Values</b>	<b>duid</b> — Specifies the IPv6 DHCP unique identifier from DHCPv6. <b>interface-id</b> — Specifies the IPv6 interface-id option. <b>interface-id-link-local</b> — Specifies the interface-id and link-local address.

## use-link-address

<b>Syntax</b>	<b>use-link-address</b> [ <i>scope scope</i> ] <b>no use-link-address</b>
<b>Context</b>	config>router>dhcp6>local-dhcp-server
<b>Description</b>	If configured, local pool selection for v6 address or prefix assignment will use the configured link-address under relay configuration. The selected pool will contain a prefix covering the link-address. The scope option defines the scope for the match. With scope <b>subnet</b> , the prefix or address selection is limited to the prefix in the pool that covers the link-address. With scope <b>pool</b> , all the prefixes in the selected pool are eligible for assignment.
<b>Default</b>	scope subnet
<b>Parameters</b>	<b>scope</b> <i>scope</i> — Specifies the scope of the IP address selection.
<b>Values</b>	<b>subnet</b> — Specifies that the prefix or address selection is limited to the prefix in the pool that covers the link address. <b>pool</b> — Specifies that all prefixes in the selected pool are eligible for assignment.

## user-db

<b>Syntax</b>	<b>user-db</b> <i>local-user-db-name</i>
---------------	--

## DHCP Configuration Commands

### **no user-db**

<b>Context</b>	config>router>dhcp>server
<b>Description</b>	This command configures a local user database for authentication.
<b>Default</b>	not enabled
<b>Parameters</b>	<i>local-user-db-name</i> — Specifies the name of a local user database.

## Service Commands

### dhcp

<b>Syntax</b>	<b>dhcp</b>
<b>Context</b>	config>service>vpls>sap config>service>vpls>spoke-sdp config>service>vpls>mesh-sdp config>service>ies>interface config>service>vprn config>service>vprn>interface config>service>vprn>sub-if config>service>vprn>sub-if>grp-if config>service>ies>sub-if>grp-if config>service>ies>sub-if config>service>ies>sub-if>grp-if
<b>Description</b>	This command enables the context to configure DHCP parameters.

### dhcp6

<b>Syntax</b>	<b>dhcp6</b>
<b>Context</b>	config>service>vpls>sap config>service>vpls>spoke-sdp config>service>vpls>mesh-sdp config>service>ies>interface config>system config>service>vprn config>service>vprn>interface config>service>vprn>sub-if config>service>vprn>sub-if>grp-if config>service>ies>sub-if>group-grp-if config>service>vprn>sub-if>grp-if>ipv6 config>service>ies>sub-if>grp-if>ipv6 config>service>ies>sub-if config>service>ies>sub-if>grp-if
<b>Description</b>	This command enables the context to configure DHCP6 parameters.

### relay

<b>Syntax</b>	<b>[no] relay</b>
<b>Context</b>	config>service>vprn>sub-if>grp-if>ipv6>dhcp6 config>service>ies>sub-if>grp-if>ipv6>dhcp6

```
config>service>vprn>sub-if>ipv6>dhcp6
config>service>ies>sub-if>ipv6>dhcp6
```

**Description** This command enables the context to configure DHCPv6 relay parameters for this interface.

## client-applications

**Syntax** **client-applications dhcp**  
**client-applications pppoe**  
**client-applications dhcp pppoe**  
**no client-applications**

**Context** config>service>vprn>sub-if>dhcp  
config>service>ies>sub-if>dhcp  
config>service>vprn>sub-if>grp-if>dhcp  
config>service>ies>sub-if>grp-if>dhcp  
config>service>vprn>sub-if>grp-if>ipv6>dhcp6>proxy  
config>service>vprn>sub-if>grp-if>ipv6>dhcp6>relay  
config>service>ies>sub-if>grp-if>ipv6>dhcp6>proxy  
config>service>ies>sub-if>grp-if>ipv6>dhcp6>relay  
config>service>vprn>sub-if>ipv6>dhcp6>proxy  
config>service>vprn>sub-if>ipv6>dhcp6>relay  
config>service>ies>sub-if>ipv6>dhcp6>proxy  
config>service>ies>sub-if>ipv6>dhcp6>relay

**Description** This command enables DHCP relay and proxy-server for the configured client types.  
The **no** form of the command resets the default client application (dhcp).

**Default** client-applications dhcp

**Parameters** **dhcp** — Enables IPoE clients to use the DHCP relay or proxy-server  
**pppoe** — Enables PPPoE clients to use the DHCP relay or proxy-server that PPPoE will attempt to request an IP address for a PPPoE client from the DHCP server(s)ly assigned to PPPoE node.

## match-circuit-id

**Syntax** **[no] match-circuit-id**

**Context** config>service>ies>sub-if>grp-if>dhcp  
config>service>vprn>sub-if>dhcp  
config>service>vprn>sub-if>grp-if>dhcp

**Description** This command enables matching Option 82 circuit ID on relayed DHCP packet matching.  
For Routed CO, the group interface DHCP relay process is stateful. When packets are relayed to the server the virtual router ID, transaction ID, SAP ID, and client hardware MAC address of the relayed packet are tracked. When a response is received from the server the virtual router ID, transaction ID, and client HW MAC address must be matched to determine the SAP on which to send the packet out. In some cases, the virtual router ID, transaction ID, and client HW MAC address are not guaranteed to be unique.

When the **match-circuit-id** command is enabled this part of the key is used to guarantee correctness in the lookup. This is only needed when are dealing with an IP aware DSLAM that proxies the client HW mac address.

**Default** no match-circuit-id

## lease-populate

**Syntax** **lease-populate** [*nbt-of-entries*]  
**no lease-populate**

**Context** config>service>vpls>if>dhcp>option  
config>service>ies>if>dhcp>option

**Description** This command enables dynamic host lease state management for SAPs.  
For VPLS, DHCP snooping must be explicitly enabled (using the **snoop** command) at all points where DHCP messages requiring snooping enter the VPLS instance (both from the DHCP server and from the subscribers). Lease state information is extracted from snooped DHCP ACK messages to populate lease state table entries for the MSAP.

The optional number-of-entries parameter is used to define the number lease state table entries allowed for an MSAP or IP interface. If number-of-entries is omitted, only a single entry is allowed. Once the maximum number of entries has been reached, subsequent lease state entries are not allowed and subsequent DHCP ACK messages are discarded.

The retained lease state information representing dynamic hosts may be used to:

- Populate an MSAP based anti-spoof filter table to provide dynamic anti-spoof filtering. If the system is unable to populate the dynamic host information in the anti-spoof filter table on the SAP, the DHCP ACK message must be discarded without adding new lease state entry or updating an existing lease state entry.
- Generate dynamic ARP replies if **arp-reply-agent** is enabled.

The **no** form of the command disables dynamic host lease state management for the MSAP.

**Default** no lease-populate

## lease-populate

**Syntax** **lease-populate** [*nbr-of-leases*]  
**lease-populate** [*nbr-of-leases*] **l2-header** [*mac ieee-address*]  
**no lease-populate**

**Context** config>subscr-mgmt>msap-policy>vpls-only>dhcp  
config>service>vpls>sap>dhcp  
config>service>ies>interface>dhcp  
config>service>vprn>interface>dhcp  
config>service>ies>sub-if>grp-if>dhcp  
config>service>vprn>sub-if>grp-if>dhcp  
config>service>vprn>sub-if>dhcp

**Description** This command enables and disables dynamic host DHCPv4 lease state management for SAPs.

For VPLS, DHCP snooping must be explicitly enabled (using the **snoop** command) at all points where DHCP messages requiring snooping enter the VPLS instance (both from the DHCP server and from the subscribers). Lease state information is extracted from snooped DHCP ACK messages to populate lease state table entries for the SAP.

The optional number-of-entries parameter defines the number lease state table entries allowed.

- for this SAP in case of a VPLS service
- for this interface in case of an IES or VPRN interface
- for each SAP in case of an IES or VPRN group-interface
- for this interface in case of an IES or VPRN retail subscriber-interface

If number-of-entries is omitted, only a single entry is allowed. Once the maximum number of entries has been reached, subsequent lease state entries are not allowed and subsequent DHCP ACK messages are discarded.

The retained lease state information representing dynamic hosts may be used to:

- Populate a SAP based anti-spoof filter table to provide dynamic anti-spoof filtering. If the system is unable to populate the dynamic host information in the anti-spoof filter table on the SAP, the DHCP ACK message must be discarded without adding new lease state entry or updating an existing lease state entry.
- Populate the system's ARP cache based on the arp-populate configuration. Applicable to IES and VPRN interfaces or group-interfaces.
- Populate managed entries into a VPLS forwarding database. VPLS forwarding database population is an implicit feature that automatically places the dynamic host's MAC address into the VPLS FDB. When a dynamic host's MAC address is placed in the lease state table, it will automatically be populated into the VPLS forwarding database associated with the SAP on which the host is learned. The dynamic host MAC address will override any static MAC entries using the same MAC and prevent dynamic learning of the MAC on another interface. Existing static MAC entries with the same MAC address as the dynamic host are marked as inactive but not deleted. If all entries in the lease state table associated with the MAC address are removed, the static MAC may be populated. New static MAC definitions for the VPLS instance may be created while a dynamic host exists associated with the static MAC address
- Generate dynamic ARP replies if **arp-reply-agent** is enabled. Applicable to VPLS service SAPs

<b>Default</b>	no lease-populate
<b>Parameters</b>	<i>nbr-of-leases</i> — Specifies the number of DHCPv4 leases allowed.
<b>Values</b>	1 — 32767 1 — 65535 (chassis-mode d, SF/CPM-4 or later) 1 — 262143 (chassis-mode d, SF/CPM-4 or later, retail subscriber interfaces only)
<b>l2-header</b>	— Indicates a mode of operation where anti-spoof entry associated with the given DHCP state is created based on the MAC address from the Layer 2 header. The Layer 2 header flag is not set by default. This parameter is only applicable for group-interfaces.
<b>mac</b>	— Specifies that the provisioned ieee-address will be used in the anti-spoofing entries for this SAP. The parameter may be changed mid-session. Existing sessions will not be re-programmed unless a tools perform command is issues for the lease. This parameter is only applicable for group-interfaces.

## option

<b>Syntax</b>	<b>[no] option</b>
<b>Context</b>	<pre>config&gt;service&gt;vpls&gt;sap&gt;dhcp config&gt;service&gt;vpls&gt;sap&gt;dhcp6 config&gt;service&gt;ies&gt;interface&gt;dhcp config&gt;service&gt;vprn&gt;interface&gt;dhcp config&gt;service&gt;vprn&gt;sub-if&gt;dhcp config&gt;service&gt;vprn&gt;sub-if&gt;grp-if&gt;dhcp config&gt;service&gt;ies&gt;sub-if&gt;grp-if&gt;dhcp</pre>
<b>Description</b>	<p>This command enables DHCP Option 82 (Relay Agent Information Option) parameters processing and enters the context for configuring Option 82 sub-options.</p> <p>The <b>no</b> form of this command returns the system to the default.</p>
<b>Default</b>	no option

## action

<b>Syntax</b>	<b>action {replace   drop   keep}</b> <b>no action</b>
<b>Context</b>	<pre>config&gt;service&gt;vpls&gt;sap&gt;dhcp&gt;option config&gt;service&gt;ies&gt;interface&gt;dhcp&gt;option config&gt;service&gt;vprn&gt;interface&gt;dhcp&gt;option config&gt;service&gt;vprn&gt;sub-if&gt;grp-if&gt;dhcp&gt;option config&gt;service&gt;ies&gt;sub-if&gt;grp-if&gt;dhcp</pre>
<b>Description</b>	<p>This command configures the Relay Agent Information Option (Option 82) processing.</p> <p>The <b>no</b> form of this command returns the system to the default value.</p>
<b>Default</b>	The default is to keep the existing information intact.
<b>Parameters</b>	<p><b>replace</b> — In the upstream direction (from the user), the Option 82 field from the router is inserted in the packet (overwriting any existing Option 82 field). In the downstream direction (towards the user) the Option 82 field is stripped (in accordance with RFC 3046).</p> <p><b>drop</b> — The DHCP packet is dropped if an Option 82 field is present, and a counter is incremented.</p> <p><b>keep</b> — The existing information is kept in the packet and the router does not add any additional information. In the downstream direction the Option 82 field is not stripped and is forwarded towards the client.</p> <p>In Vendor-Specific Options (VSOs) scenarios, the behavior is slightly different. Even with the action=keep, the router will insert his own vso into the Option 82 field. This will only be done when the incoming message has already an Option 82 field.</p> <p>If no Option 82 field is present, the router will not create the Option 82 field - in that case, no VSO will be added to the message.</p>

## circuit-id

<b>Syntax</b>	<b>circuit-id</b> [ascii-tuple   vlan-ascii-tuple] <b>no circuit-id</b>
<b>Context</b>	config>service>vpls>sap>dhcp>option
<b>Description</b>	<p>When enabled, the router sends an ASCII-encoded tuple in the <b>circuit-id</b> sub-option of the DHCP packet. This ASCII-tuple consists of the access-node-identifier, service-id, and SAP-ID, separated by “ ”.</p> <p>In order to send a tuple in the circuit ID, the <b>action replace</b> command must be configured in the same context.</p> <p>If disabled, the <b>circuit-id</b> sub-option of the DHCP packet will be left empty.</p> <p>The <b>no</b> form of this command returns the system to the default.</p>
<b>Default</b>	circuit-id
<b>Parameters</b>	<p><b>ascii-tuple</b> — Specifies that the ASCII-encoded concatenated tuple consisting of the access-node-identifier, service-id, and interface-name is used.</p> <p><b>vlan-ascii-tuple</b> — Specifies that the format will include VLAN-id and dot1p bits in addition to what is included in ascii-tuple already. The format is supported on dot1q and qinq ports only. Thus, when the Option 82 bits are stripped, dot1p bits will be copied to the Ethernet header of an outgoing packet.</p>

## circuit-id

<b>Syntax</b>	<b>circuit-id</b> [ascii-tuple   ifindex   sap-id   vlan-ascii-tuple] <b>no circuit-id</b>
<b>Context</b>	config>service>ies>if>dhcp>option config>service>ies>sub-if>grp-if>dhcp>option config>service>vprn>if>dhcp>option config>service>vprn>sub-if>grp-if>dhcp>option
<b>Description</b>	<p>When enabled, the router sends an ASCII-encoded tuple in the <b>circuit-id</b> sub-option of the DHCP packet. This ASCII-tuple consists of the access-node-identifier, service-id, and SAP-ID, separated by “ ”.</p> <p>In order to send a tuple in the circuit ID, the <b>action replace</b> command must be configured in the same context.</p> <p>If disabled, the <b>circuit-id</b> sub-option of the DHCP packet will be left empty.</p> <p>The <b>no</b> form of this command returns the system to the default.</p>
<b>Default</b>	circuit-id
	<p><b>ascii-tuple</b> — Specifies that the ASCII-encoded concatenated tuple will be used which consists of the access-node-identifier, service-id, and interface-name, separated by “ ”.</p> <p><b>ifindex</b> — Specifies that the interface index will be used. (The If Index of a router interface can be displayed using the command <b>show&gt;router&gt;interface&gt;detail</b>)</p> <p><b>sap-id</b> — Specifies that the SAP identifier will be used.</p>



**vlan-ascii-tuple** — Specifies that the format will include VLAN-id and dot1p bits in addition to what is included in ascii-tuple already. The format is supported on dot1q-encapsulated ports only. Thus, when the option 82 bits are stripped, dot1p bits will be copied to the Ethernet header of an outgoing packet.

## remote-id

<b>Syntax</b>	<b>remote-id</b> [ <b>mac</b>   <b>string</b> <i>string</i> ] <b>no remote-id</b>
<b>Context</b>	config>service>vpls>sap>dhcp>option config>service>ies>if>dhcp>option config>service>vprn>if>dhcp>option config>service>ies>sub-if>grp-if>dhcp>option config>service>ies>sub-if>grp-if>dhcp>option
<b>Description</b>	This command specifies what information goes into the remote-id sub-option in the DHCP relay packet.  If disabled, the <b>remote-id</b> sub-option of the DHCP packet will be left empty.  The <b>no</b> form of this command returns the system to the default.
<b>Default</b>	remote-id
<b>Parameters</b>	<b>mac</b> — This keyword specifies the MAC address of the remote end is encoded in the sub-option. <b>string</b> <i>string</i> — Specifies the remote-id.

## vendor-specific-option

<b>Syntax</b>	<b>[no] vendor-specific-option</b>
<b>Context</b>	config>service>vpls>sap>dhcp>option config>service>ies>if>dhcp>option config>service>vprn>if>dhcp>option config>service>ies>sub-if>grp-if>dhcp>option config>service>vprn>sub-if>grp-if>dhcp>option
<b>Description</b>	This command configures the Alcatel-Lucent vendor specific sub-option of the DHCP relay packet.

## client-mac-address

<b>Syntax</b>	<b>[no] client-mac-address</b>
<b>Context</b>	config>service> ies>if>dhcp>option>vendor config>service>vprn>if>dhcp>option>vendor config>service>vprn>sub-if>grp-if>dhcp>option>vendor config>service>ies>sub-if>grp-if>dhcp>option>vendor
<b>Description</b>	This command enables the sending of the MAC address in the Alcatel-Lucent vendor specific sub-option of the DHCP relay packet.

The **no** form of the command disables the sending of the MAC address in the Alcatel-Lucent vendor specific sub-option of the DHCP relay packet.

### pool-name

<b>Syntax</b>	<b>[no] pool-name</b>
<b>Context</b>	config>service>vprn>if>dhcp>option>vendor config>service>vprn>sub-if>grp-if>dhcp>option>vendor config>service>ies>if>dhcp>option>vendor config>service>ies>sub-if>grp-if>dhcp>option>vendor
<b>Description</b>	This command sends the pool name in the Alcatel vendor specific suboption of the DHCP relay packet.  The <b>no</b> form of the command disables the sending.

### sap-id

<b>Syntax</b>	<b>[no] sap-id</b>
<b>Context</b>	config>service>vpls>sap>dhcp>option>vendor config>service>vprn>if>dhcp>option>vendor config>service>vprn>sub-if>grp-if>dhcp>option>vendor config>service>ies>sub-if>grp-if>dhcp>option>vendor
<b>Description</b>	This command enables the sending of the SAP ID in the Alcatel-Lucent vendor specific sub-option of the DHCP relay packet.  The <b>no</b> form of the command disables the sending of the SAP ID in the Alcatel-Lucent vendor specific sub-option of the DHCP relay packet.

### service-id

<b>Syntax</b>	<b>[no] service-id</b>
<b>Context</b>	config>service>vpls>sap>dhcp>option>vendor config>service>vprn>if>dhcp>option>vendor config>service>vprn>sub-if>grp-if>dhcp>option>vendor config>service>ies>sub-if>grp-if>dhcp>option>vendor
<b>Description</b>	This command enables the sending of the service ID in the Alcatel-Lucent vendor specific sub-option of the DHCP relay packet.  The <b>no</b> form of the command disables the sending of the service ID in the Alcatel-Lucent vendor specific sub-option of the DHCP relay packet.

## string

<b>Syntax</b>	<b>[no] string <i>text</i></b>
<b>Context</b>	config>service>vpls>sap>dhcp>option>vendor config>service>vprn>if>dhcp>option>vendor config>service>vprn>sub-if>grp-if>dhcp>option>vendor config>service>ies>sub-if>grp-if>dhcp>option>vendor
<b>Description</b>	This command specifies the string in the Alcatel-Lucent vendor specific sub-option of the DHCP relay packet.  The <b>no</b> form of the command returns the default value.
<b>Parameters</b>	<i>text</i> — The string can be any combination of ASCII characters up to 32 characters in length. If spaces are used in the string, enclose the entire string in quotation marks (“ ”).

## system-id

<b>Syntax</b>	<b>[no] system-id</b>
<b>Context</b>	config>service>vpls>sap>dhcp>option>vendor config>service>vprn>if>dhcp>option>vendor config>service>vprn>sub-if>grp-if>dhcp>option>vendor config>service>ies>sub-if>grp-if>dhcp>option>vendor
<b>Description</b>	This command specifies whether the system-id is encoded in the Alcatel-Lucent vendor specific sub-option of Option 82.

## override-slaac

<b>Syntax</b>	<b>[no] override-slaac</b>
<b>Context</b>	config>service>vprn>sub-if>>ipv6>dhcp6 config>service>ies>sub-if>ipv6>dhcp6 config>service>vprn>sub-if>grp-if>ipv6>dhcp6 config>service>ies>sub-if>grp-if>ipv6>dhcp6
<b>Description</b>	This command allows a DHCP IA_NA address to override and replace a host existing SLAAC address. When this feature is enabled, a subscriber SLAAC address is removed once the DHCP IA_NA address assignment is completed. If used with conjunction with the <b>allow-multiple-wan-address</b> command, the DHCP IA_NA address will also override the SLAAC address.

## pd-managed-route

<b>Syntax</b>	<b>[no] pd-managed-route</b>
<b>Context</b>	config>service>vprn>sub-if>ipv6>dhcp6 config>service>ies>sub-if>ipv6>dhcp6 config>service>vprn>sub-if>grp-if>ipv6>dhcp6

```
config>service>ies>sub-if>grp-if>ipv6>dhcp6
```

- Description** This command enables DHCP IA-PD (delegated prefix) to be modeled as managed (framed) route instead of as a subscriber-host. Antispoof filtering for the subscriber host associated with the IA-PD route must be set to nh-mac. The subscriber specific parameters (sla-profile, sla-profile, etc) will be ignored during the authentication phase since IA-PD is not modeled as a subscriber host.
- The next-hop for PD managed route must be an IPv6 sub-host (DHCP IA-NA or SLAAC) with the same mac address as the one in the DHCP lease state for the managed IA-PD. DHCP IA-NA next-hop host will always override SLAAC next-hop host if both are available. In case that the v6 next-hop is not present at the time when the framed IA-PD is instantiated, the IA-PD will be setup but the traffic destined to the IA-PD managed route will be black-holed.
- The typical subscriber host information for DHCP IA-PD modeled as a route is removed from the operational show commands related to the subscriber host state (i.e. show service active-subscribert or show service id X subscriber-host). However, DHCP IA-PD route is displayed as s managed route for the corresponding IPv6 subscriber host (DHCP IA-NA or SLAAC).
- DHCP IA-PD information for managed IA-PD route is still maintained in the DHCPv6 lease state and as such it can be displayed with the appropriate show command.

**Default** no pd-managed-route

## enable-ingress-stats

- Syntax** [no] enable-ingress-stats
- Context** config>service>ies>sub-if>grp-if  
config>service>vprn>sub-if>grp-if
- Description** This command enables the collection of ingress interface IP stats. This command is only applicable to IP statistics, and not to uRPF statistics.
- If enabled, then the following statistics are collected:
- IPv4 offered packets
  - IPv4 offered octets
  - IPv6 offered packets
  - IPv6 offered octets
- Note that octet statistics for IPv4 and IPv6 bytes at IP interfaces include the Layer 2 frame overhead.
- Default** no enable-ingress-stats

## duid

<b>Syntax</b>	<b>duid</b> <i>duid</i> [ <b>iaid</b> <i>iaid</i> ] <b>no duid</b>
<b>Context</b>	config>service>ies>if>ipv6>dhcp6>pfx-delegate>prefix
<b>Description</b>	This command configures the DHCP Unique Identifier (DUID) of the DHCP client.
<b>Parameters</b>	<p><i>duid</i> — Specifies the ID of the requesting router. If set to a non zero value the prefix defined will only be delegated to this router. If set to zero, the prefix will be delegated to any requesting router.</p> <p><b>iaid</b> <i>iaid</i> — Specifies the identity association identification (IAID) from the requesting router that needs to match in order to delegate the prefix defined in this row. If set to 0 no match on the received IAID is done.</p>

## preferred-lifetime

<b>Syntax</b>	<b>preferred-lifetime</b> [ <b>days</b> <i>days</i> ] [ <b>hrs</b> <i>hours</i> ] [ <b>min</b> <i>minutes</i> ] [ <b>sec</b> <i>seconds</i> ] <b>preferred-lifetime infinite</b> <b>no preferred-lifetime</b>								
<b>Context</b>	config>service>ies>if>ipv6>dhcp6>pfx-delegate>prefix config>service>vprn>if>ipv6>dhcp6>pfx-delegate>prefix config>service>vprn>sub-if>ipv6>dhcp6>proxy config>service>ies>sub-if>ipv6>dhcp6>proxy config>service>vprn>sub-if>grp-if>ipv6>dhcp6>proxy config>service>ies>sub-if>grp-if>ipv6>dhcp6>proxy								
<b>Description</b>	<p>This command configures the IPv6 prefix/mask preferred life time. The preferred-lifetime value cannot be bigger than the valid-lifetime value.</p> <p>The <b>no</b> form of the command reverts to the default value.</p>								
<b>Default</b>	604800 seconds (7 days)								
<b>Parameters</b>	[ <b>days</b> <i>days</i> ][ <b>hrs</b> <i>hours</i> ] [ <b>min</b> <i>minutes</i> ] [ <b>sec</b> <i>seconds</i> ] — Specifies the preferred lifetime.								
<b>Values</b>	<table> <tr><td>days:</td><td>[0..3650]</td></tr> <tr><td>hours:</td><td>[0..23]</td></tr> <tr><td>minutes:</td><td>[0..59]</td></tr> <tr><td>seconds:</td><td>[0..59]</td></tr> </table>	days:	[0..3650]	hours:	[0..23]	minutes:	[0..59]	seconds:	[0..59]
days:	[0..3650]								
hours:	[0..23]								
minutes:	[0..59]								
seconds:	[0..59]								

## preferred-lifetime

<b>Syntax</b>	<b>preferred-lifetime seconds</b> <b>preferred-lifetime infinite</b> <b>no preferred-lifetime</b>
<b>Context</b>	config>service>ies>sub-if>grp-if>ipv6>rtr-adv>pfx-opt config>service>vprn>sub-if>grp-if>ipv6>rtr-adv>pfx-opt config>service>ies>sub-if>ipv6>rtr-adv>pfx-opt config>service>vprn>sub-if>ipv6>rtr-adv>pfx-opt

<b>Description</b>	This command specifies the remaining time for this prefix to be preferred, thus time until deprecation.
<b>Default</b>	3600 seconds
<b>Parameters</b>	<i>seconds</i> — Specifies the time for the prefix to remain preferred on this group-interface in seconds. <b>Values</b> 0 — 4294967295 <b>infinite</b> — Specifies that the remaining time will never expire. Note that the value <b>4294967295</b> seconds is interpreted as infinite.

## valid-lifetime

<b>Syntax</b>	<b>valid-lifetime <i>seconds</i></b> <b>valid-lifetime infinite</b> <b>no valid-lifetime</b>
<b>Context</b>	config>service>ies>if>ipv6>dhcp6>pfx-delegate>prefix config>service>vpn>sub-if>grp-if>ipv6>dhcp6 config>service>ies>sub-if>grp-if>ipv6>dhcp6
<b>Description</b>	This command configures the time, in seconds, that the prefix is valid. The maximum value 4294967295 is considered equal to infinity. The <b>no</b> form of the command reverts to the default value.
<b>Default</b>	2592000 seconds (30 days)
<b>Parameters</b>	<i>seconds</i> — Specifies the time, in seconds, that this prefix remains valid. <b>Values</b> 1 — 4294967294 <b>infinite</b> — Specifies that this prefix remains valid infinitely.

## valid-lifetime

<b>Syntax</b>	<b>valid-lifetime <i>seconds</i></b> <b>valid-lifetime infinite</b> <b>no valid-lifetime</b>
<b>Context</b>	config>service>ies>sub-if>grp-if>ipv6>rtr-adv config>service>vpn>sub-if>grp-if>ipv6>rtr-adv
<b>Description</b>	This command specifies the remaining time for this prefix to be valid for the purpose of on-link determination.
<b>Default</b>	86400
<b>Parameters</b>	<i>seconds</i> — Specifies the time for the prefix to remain valid on this group-interface in seconds. <b>Values</b> 0 — 4294967295 <b>infinite</b> — Specifies that the remaining time will never expire.

## python-policy

<b>Syntax</b>	<b>python-policy</b> <i>name</i> <b>no python-policy</b>
<b>Context</b>	config>service>ies>if>dhcp config>service>vprn>if>dhcp>
<b>Description</b>	This command specifies the python-policy to be used for DHCPv4 relay.
<b>Parameters</b>	<i>name</i> — Specifies the name of an existing python script up to 32 characters in length.

## python-policy

<b>Syntax</b>	<b>python-policy</b> <i>name</i> <b>no python-policy</b>
<b>Context</b>	config>service>vprn>sub-if config>service>vprn>sub-if>ipv6>dhcp6 config>service>ies>sub-if>ipv6>dhcp6 config>service>ies>if>ipv6>dhcp6-relay config>service>vprn>if>ipv6>dhcp6-relay
<b>Description</b>	This command specifies the python-policy to be used for DHCPv6 relay.
<b>Parameters</b>	<i>name</i> — Specifies the name of an existing python script up to 32 characters in length.

## emulated-server

<b>Syntax</b>	<b>emulated-server</b> <i>ip-address</i> <b>no emulated-server</b>
<b>Context</b>	config>service>ies>if>dhcp>proxy-server config>service>ies>sub-if>grp-if>dhcp>proxy-server config>service>vpls>sap>dhcp>proxy-server config>service>vprn>sub-if>grp-if>dhcp
<b>Description</b>	<p>This command configures the IP address which will be used as the DHCP server address in the context of the SAP. Typically, the configured address should be in the context of the subnet represented by the service.</p> <p>The <b>no</b> form of this command reverts to the default setting. The local proxy server will not become operational without the emulated-server address being specified.</p>
<b>Parameters</b>	<i>ip-address</i> — Specifies the emulated server's IP address. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).

## emulated-server

<b>Syntax</b>	<b>emulated-server</b> <i>ip-address</i> <b>no emulated-server</b>
<b>Context</b>	config>service>vprn>if>dhcp>proxy config>service>vprn>sub-if>grp-if>dhcp>proxy-server
<b>Description</b>	This command configures IP address which will be used as DHCP server address in context of the SAP. Typically, configured address should be in context of the subnet represented by VPRN. No version of these commands reverts to default setting. The local proxy server will not become operational without emulated-server address being specified.
<b>Parameters</b>	<i>ip-address</i> — Specifies the emulated server's IP address.

## lease-time

<b>Syntax</b>	<b>lease-time</b> [ <b>days</b> <i>days</i> ] [ <b>hrs</b> <i>hours</i> ] [ <b>min</b> <i>minutes</i> ] [ <b>sec</b> <i>seconds</i> ] [ <b>override</b> ] <b>no lease-time</b>
<b>Context</b>	config>service>vpls>sap>dhcp>proxy-server config>service>ies>if>dhcp>proxy-server config>service>ies>sub-if>grp-if>dhcp>proxy-server config>service>vprn>if>dhcp>proxy config>service>vprn>sub-if>grp-if>dhcp>proxy-server
<b>Description</b>	This command defines the length of lease-time that will be provided to DHCP clients. By default the local-proxy-server will always make use of the lease-time information provide by either a RADIUS or DHCP server.  The <b>no</b> form of this command disables the use of the lease-time command. The local-proxy-server will use the lease-time offered by either a RADIUS or DHCP server.
<b>Default</b>	7 days 0 hours 0 seconds
<b>Parameters</b>	<b>override</b> — Specifies that the local-proxy-server will use the configured lease-time information to provide DHCP clients.  <b>radius-override</b> — Supported only in the <b>config&gt;service&gt;vpls&gt;sap&gt;dhcp&gt;proxy-server</b> context, specifies that the local-proxy-server will use the configured lease-time information to provide DHCP clients.  <i>days</i> — Specifies the number of days that the given IP address is valid.  <b>Values</b> 0 — 3650  <i>hours</i> — Specifies the number of hours that the given IP address is valid.  <b>Values</b> 0 — 23  <i>minutes</i> — Specifies the number of minutes that the given IP address is valid.  <b>Values</b> 0 — 59  <i>seconds</i> — Specifies the number of seconds that the given IP address is valid.  <b>Values</b> 0 — 59



## snoop

<b>Syntax</b>	<b>snoop</b> <b>no snoop</b>
<b>Context</b>	config>service>vpls>sap>dhcp config>service>vpls>sap>dhcp6 config>service>vpls>spoke-sdp>dhcp config>service>vpls>mesh-sdp>dhcp config>service>vprn>if>dhcp>option
<b>Description</b>	<p>This command enables DHCP snooping of DHCP messages on the SAP or SDP. Enabling DHCP snooping on interfaces (SAPs and SDP bindings) is required where DHCP messages important to lease state table population are received, or where Option 82 information is to be inserted. This includes interfaces that are in the path to receive messages from either DHCP servers or from subscribers.</p> <p>Use the <b>no</b> form of the command to disable DHCP snooping on the specified SAP or SDP binding.</p>
<b>Default</b>	no snoop dhcp6

## dhcp-user-db

<b>Syntax</b>	<b>dhcp-user-db</b> <i>local-user-db</i> <b>no dhcp-user-db</b>
<b>Context</b>	configure>service>vpls>sap
<b>Description</b>	This command enabled access to LUDB for DHCPv4 hosts under the capture SAP. The name of this ludb must match the name of ludb configured under the <b>configure&gt;service&gt;vprn/ies&gt;subscr-intf&gt;group-intf&gt;dhcp</b> hierarchy.
<b>Default</b>	no dhcp-user-db
<b>Parameters</b>	<i>local-user-db</i> — Specifies the name of the local-user-database up to 32 characters max.

## dhcp-python-policy

<b>Syntax</b>	<b>dhcp-python-policy</b> <i>policy-name</i> <b>no dhcp-python-policy</b>
<b>Context</b>	config>service>vpls>sap
<b>Description</b>	<p>This command specifies the name of the Python policy. The Python policy is created in the <b>config&gt;python&gt;python-policy</b> <i>name</i> context.</p> <p>The <b>no</b> form of the command reverts to the default.</p>
<b>Default</b>	none
<b>Parameters</b>	<i>policy-name</i> — Specifies a Python policy name up to 32 characters in length.

## dhcp6

<b>Syntax</b>	<b>dhcp6</b>
<b>Context</b>	config>service>vpls>sap config>service>ies>sub-if>grp-if>ipv6 config>service>vprn>sub-if>grp-if>ipv6
<b>Description</b>	This command configures DHCP6 parameters for this SAP.

## interface-id

<b>Syntax</b>	<b>interface-id</b> <b>interface-id ascii-tuple</b> <b>interface-id vlan-ascii-tuple</b> <b>no interface-id</b>
<b>Context</b>	config>service>ies>if>ipv6>dhcp6>option config>service>vprn>if>ipv6>dhcp6>option config>service>vpls>sap>dhcp6>option
<b>Description</b>	This command configure the interface-id suboption of the DHCP6 Relay packet The <b>no</b> form of the command disables the sending of interface ID options in the DHCPv6 relay packet
<b>Parameters</b>	<b>ascii-tuple</b> — Specifies that the ASCII-encoded concatenated tuple will be used which consists of the access-node-identifier, service-id, and interface-name, separated by “ ”. <b>vlan-ascii-tuple</b> — Specifies that the format will include VLAN-id and dot1p bits in addition to what is included in ascii-tuple already. The format is supported on dot1q-encapsulated ports only. Thus, when the option 82 bits are stripped, dot1p bits will be copied to the Ethernet header of an outgoing packet. <b>mac</b> — This keyword specifies the MAC address of the remote end is encoded in the sub-option.

## remote-id

<b>Syntax</b>	<b>remote-id</b> <b>remote-id mac</b> <b>remote-id string</b> [32 chars max] <b>no remote-id</b>
<b>Context</b>	config>service>ies>if>ipv6>dhcp6>option config>service>vprn>if>ipv6>dhcp6>option config>service>vpls>sap>dhcp6>option
<b>Description</b>	This command enables the sending of remote ID option in the DHCPv6 relay packet. The client DHCP Unique Identifier (DUID) is used as the remote ID. The <b>no</b> form of the command disables the sending of remote ID option in the DHCPv6 relay packet.

## interface-id

<b>Syntax</b>	<b>interface-id</b> <b>interface-id ascii-tuple</b> <b>interface-id vlan-ascii-tuple</b> <b>no interface-id</b>
<b>Context</b>	config>service>vpls>sap>dhcp6>option
<b>Description</b>	This command configures the interface-id suboption of the DHCP6 relay packet. The no form of the command reverts to the default.
<b>Default</b>	none
<b>Parameters</b>	<b>ascii-tuple</b> — Specifies that the ASCII-encoded concatenated tuple consisting of the access-node-identifier, service-id, and interface-name is used.  <b>vlan-ascii-tuple</b> — Specifies that the format will include VLAN-id and dot1p bits in addition to what is included in ascii-tuple already. The format is supported on dot1q and qinq ports only. Thus, when the option 82 bits are stripped, dot1p bits will be copied to the Ethernet header of an outgoing packet.

## dhcp6-user-db

<b>Syntax</b>	<b>dhcp6-user-db local-user-db</b> <b>no dhcp6-user-db</b>
<b>Context</b>	configure>service>vpls>sap
<b>Description</b>	This command enabled access to LUDB for DHCPv6 hosts under the capture SAP. The name of this ludb must match the name of ludb configured under the <b>configure&gt;service&gt;vprn/ies&gt;subscr-intf&gt;group-intf&gt;dhcp</b> hierarchy.
<b>Default</b>	no dhcp6-user-db
<b>Parameters</b>	<i>local-user-db</i> — Specifies the name of the local-user-database up to 32 characters max.

## ppp-user-db

<b>Syntax</b>	<b>ppp-user-db local-user-db-name</b> <b>no ppp-user-db</b>
<b>Context</b>	configure>service>vpls
<b>Description</b>	This command enabled access to LUDB for PPPoE and PPPoEoA v4and v6 hosts under the capture SAP. The name of this ludb must match the name of ludb configured under the <b>configure&gt;service&gt;vprn/ies&gt;subscr-intf&gt;group-intf&gt;pppoe</b> hierarchy.
<b>Default</b>	no pppoe-user-db
<b>Parameters</b>	<i>local-user-db</i> — Specifies the name of the local-user-database up to 256 characters max.

## pppoe-user-db

<b>Syntax</b>	<b>pppoe-user-db</b> <i>local-user-db-name</i> <b>no pppoe-user-db</b>
<b>Context</b>	configure>service>vpls
<b>Description</b>	This command enabled access to LUDB for PPPoE and PPPoEoA v4and v6 hosts under the capture SAP. The name of this ludb must match the name of ludb configured under the <b>configure&gt;service&gt;vprn/ies&gt;subscr-intf&gt;group-intf&gt;pppoe</b> hierarchy.
<b>Default</b>	no pppoe-user-db
<b>Parameters</b>	<i>local-user-db</i> — Specifies the name of the local-user-database up to 256 characters max.

## filter

<b>Syntax</b>	<b>filter</b> <i>filter-id</i> <b>no filter</b>
<b>Context</b>	config>service>ies>sub-if>grp-if>dhcp
<b>Description</b>	This command configures the DHCP filter for this interface

## gi-address

<b>Syntax</b>	<b>gi-address</b> <i>ip-address</i> [ <i>src-ip-addr</i> ] <b>no gi-address</b>
<b>Context</b>	config>service>ies>if>dhcp config>service>vprn>interface>dhcp config>service>vprn>sub-if>dhcp config>service>ies>sub-if>grp-if>dhcp config>service>ies>sub-if>dhcp
<b>Description</b>	<p>This command configures the gateway interface address for the DHCP relay. A subscriber interface can include multiple group interfaces with multiple SAPs. The GI address is needed, when the router functions as a DHCP relay, to distinguish between the different subscriber interfaces and potentially between the group interfaces defined.</p> <p>By default, the GI address used in the relayed DHCP packet is the primary IP address of a normal IES interface. Specifying the GI address allows the user to choose a secondary address. For group interfaces a GI address must be specified under the group interface DHCP context or subscriber-interface DHCP context in order for DHCP to function.</p>
<b>Default</b>	no gi-address
<b>Parameters</b>	<p><i>ip-address</i> — Specifies the host IP address to be used for DHCP relay packets.</p> <p><i>src-ip-address</i> — Specifies that this GI address is to be the source IP address for DHCP relay packets.</p>

## gi-address

<b>Syntax</b>	<b>gi-address</b> <i>ipv4-address</i> <b>no gi-address</b>
<b>Context</b>	configure>subscr-mgmt>loc-user-db>ipoe>host
<b>Description</b>	This command allows selection of gi-addresses based on the host entry in LUDB.  The gi-address must be a valid address (associated with an interface) within the routing context that received the DHCP message on the access side.
<b>Default</b>	no gi-address
<b>Parameters</b>	<i>ipv4-address</i> — Specifies the IPv4 gi-address.

## relay-plain-bootp

<b>Syntax</b>	<b>[no] relay-plain-bootp</b>
<b>Context</b>	config>service>ies>if>dhcp
<b>Description</b>	This command enables the relaying of plain BOOTP packets.  The <b>no</b> form of the command disables the relaying of plain BOOTP packets.

## relay-unicast-msg

<b>Syntax</b>	<b>relay-unicast-msg [release-update-src-ip]</b> <b>no relay-unicast-msg</b>
<b>Context</b>	config>service>ies>if>dhcp config>service>ies>sub-if>dhcp config>service>ies>sub-if>grp-if>dhcp config>service>vprn>if>dhcp config>service>vprn>sub-if>dhcp config>service>vprn>sub-if>grp-if>dhcp
<b>Description</b>	Relay unicast client DHCPv4 request (renew) messages. In the upstream direction: update the source-ip address and add the gateway IP address (gi-address) field before sending the message to the intended DHCP server (the message is not broadcasted to all configured DHCP servers). In the downstream direction: remove the gi-address and update the destination IP address to the value of the yiaddr (your IP address) field.  By default, unicast DHCPv4 release messages are forwarded transparently. The optional “release-update-src-ip” flag, updates the source IP address with the value used for relayed DHCPv4 messages.  Additionally when the optional flag “relay-unicast-msg” is enabled, then the gi address and source IP address of relayed DHCPv4 messages can be configured to any local configured IP address in the same routing instance.
<b>Default</b>	no relay-unicast-msg

**Parameters** **release-update-src-ip** — Updates the source IP address with the value used for relayed DHCPv4 messages.

## server

**Syntax** **server** *ipv6z-address* [*ipv6z-address*...(up to 8 max)]

**Context** config>service>ies>if>ipv6>dhcp6

**Description** This command specifies a list of servers where DHCP6 requests will be forwarded. The list of servers can entered as either IP addresses or fully qualified domain names. There must be at least one server specified for DHCP6 relay to work. If there are multiple servers then the request is forwarded to all of the servers in the list.

There can be a maximum of 8 DHCP6 servers configured.

**Default** no server

**Parameters** *ipv6z-address* — Specifies a non-global IPv4 address including a zone index as defined by the InetAddressIPv4z textual convention. Up to 8 addresses can be specified.

**Values** ipv6-address: 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

## virtual-subnet

**Syntax** [**no**] **virtual-subnet**

**Context** config>service>ies>sub-if>dhcp  
config>service>vprn>sub-if>dhcp

**Description** This command enables a virtual-subnet for DHCPv4 hosts under the subscriber-interface. With this command configured, the system will snoop and record the default router address in the DHCP ACK message for the DHCPv4 ESM host. The system could answer or traceroute request even if the default router address is not configured on the subscriber-interface.

**Default** none

## server

**Syntax** **server** *server1* [*server2*...(up to 8 max)]

**Context** config>service>ies>if>dhcp  
config>service>vprn>if>dhcp  
config>service>ies>sub-if>grp-if>dhcp

**Description** This command specifies a list of servers where requests will be forwarded. The list of servers can entered as either IP addresses or fully qualified domain names. There must be at least one server

specified for DHCP relay to work. If there are multiple servers then the request is forwarded to all of the servers in the list.

There can be a maximum of 8 DHCP servers configured.

<b>Default</b>	no server
<b>Parameters</b>	<i>server</i> — Specify the DHCP server IP address.

## relay-plain-bootp

<b>Syntax</b>	<b>[no] relay-plain-bootp</b>
<b>Context</b>	config>service>vprn>if>dhcp
<b>Description</b>	This command enables the relaying of plain BOOTP packets. The <b>no</b> form of the command disables the relaying of plain BOOTP packets.

## use-arp

<b>Syntax</b>	<b>[no] use-arp</b>
<b>Context</b>	config>service>vprn>if>dhcp
<b>Description</b>	This command enables the use of ARP to determine the destination hardware address. The <b>no</b> form of the command disables the use of ARP to determine the destination hardware address

## trusted

<b>Syntax</b>	<b>[no] trusted</b>
<b>Context</b>	config>service>ies>if>dhcp
<b>Description</b>	This command enables relaying of untrusted packets. The <b>no</b> form of this command disables the relay.
<b>Default</b>	not enabled

## host-connectivity-verify

<b>Syntax</b>	<b>host-connectivity-verify [interval <i>interval</i>] [action {remove alarm}] [family <i>family</i>]</b>
<b>Context</b>	config>service>vprn>if>sap config>service>vprn>sub-if>grp-if config>service>vprn>sub-if>grp-if>dhcp
<b>Description</b>	This command enables subscriber host connectivity verification on a given SAP within a service.

This tool will periodically scan all known hosts (from dhcp-state) and perform a UC ARP request. The subscriber host connectivity verification will maintain state (connected vs. not-connected) for all hosts.

<b>Default</b>	no host-connectivity-verify
<b>Parameters</b>	<p><b>interval</b> <i>interval</i> — The interval, expressed in minutes, which specifies the time interval which all known sources should be verified. The actual rate is then dependent on number of known hosts and interval.</p> <p><b>Values</b> 1— 6000 ) Note that a zero value can be used by the SNMP agent to disable host-connectivity-verify.)</p> <p><b>action</b> {<b>remove</b>   <b>alarm</b>} — Defines the action taken on a subscriber host connectivity verification failure for a given host. The <b>remove</b> keyword raises an alarm and removes dhcp-state and releases all allocated resources (queues, table entries, etc.). DHCP-RELEASE will be signaled to corresponding DHCP server. Static hosts will never be removed. The <b>alarm</b> keyword raises an alarm indicating that the host is disconnected.</p> <p><b>family</b> <i>family</i> — The family configuration allows the host connectivity checks to be performed for IPv4 endpoint, IPv6 endpoint or both. With family IPv6 configured, host connectivity checks will be performed on the global unicast address (assigned via SLAAC or DHCPv6 IA_NA) and link-local address of a Layer 3 RG or bridged hosts. In case of SLAAC assignment, host connectivity can only be performed if the /128 is known (via downstream ND). DHCPv6 PD assigned prefixes will be removed if link-local address is determined to be unreachable via “host connectivity check”. Reachability checks for GUA and link-local address will be done simultaneously.</p>

## dhcp

<b>Syntax</b>	<b>dhcp</b>
<b>Context</b>	config>service>vprn>interface config>service>vprn> config>service>vprn>sub-if>grp-if
<b>Description</b>	This command enables the context to configure DHCP parameters.

## action

<b>Syntax</b>	<b>action</b> { <b>replace</b>   <b>drop</b>   <b>keep</b> } <b>no action</b>
<b>Context</b>	config>service>vprn>if>dhcp>option config>service>vprn>sub-if>grp-if>dhcp>option
<b>Description</b>	<p>This command configures the processing required when the SR-Series receives a DHCP request that already has a Relay Agent Information Option (Option 82) field in the packet.</p> <p>The <b>no</b> form of this command returns the system to the default value.</p>
<b>Default</b>	Per RFC 3046, <i>DHCP Relay Agent Information Option</i> , section 2.1.1, <i>Reforwarded DHCP requests</i> , the default is to keep the existing information intact. The exception to this is if the giaddr of the



received packet is the same as the ingress address on the router. In that case the packet is dropped and an error is logged.

<b>Parameters</b>	<p><b>replace</b> — In the upstream direction (from the user), the existing Option 82 field is replaced with the Option 82 field from the router. In the downstream direction (towards the user) the Option 82 field is stripped (in accordance with RFC 3046).</p> <p><b>drop</b> — The packet is dropped, and an error is logged.</p> <p><b>keep</b> — The existing information is kept in the packet and the router does not add any additional information. In the downstream direction the Option 82 field is not stripped and is sent on towards the client.</p> <p>The behavior is slightly different in case of Vendor Specific Options (VSOs). When the keep parameter is specified, the router will insert his own VSO into the Option 82 field. This will only be done when the incoming message has already an Option 82 field.</p> <p>If no Option 82 field is present, the router will not create the Option 82 field. In this in that case, no VSO will be added to the message.</p>
-------------------	---

## match-circuit-id

<b>Syntax</b>	<b>[no] match-circuit-id</b>
<b>Context</b>	config>service>vprn>sub-if>grp-if>dhcp
<b>Description</b>	<p>This command enables Option 82 circuit ID on relayed DHCP packet matching. For routed CO, the group interface DHCP relay process is stateful. When packets are relayed to the server the virtual router ID, transaction ID, SAP ID, and client hardware MAC address of the relayed packet are tracked.</p> <p>When a response is received from the server the virtual router ID, transaction ID, and client hardware MAC address must be matched to determine the SAP on which to send the packet out. In some cases, the virtual router ID, transaction ID, and client hardware MAC address are not guaranteed to be unique.</p> <p>When the <b>match-circuit-id</b> command is enabled this part of the key is used to guarantee correctness in our lookup. This is really only needed when dealing with an IP aware DSLAM that proxies the client hardware MAC address.</p>
<b>Default</b>	no match-circuit-id

## option

<b>Syntax</b>	<b>[no] option</b>
<b>Context</b>	config>service>vprn>if>dhcp config>service>vprn>sub-if>dhcp config>service>vprn>sub-if>grp-if>dhcp
<b>Description</b>	<p>This command enables DHCP Option 82 (Relay Agent Information Option) parameters processing and enters the context for configuring Option 82 sub-options.</p> <p>The <b>no</b> form of this command returns the system to the default.</p>

**Default** no option

### vendor-specific-option

**Syntax** [no] vendor-specific-option

**Context** config>service>vprn>if>dhcp>option  
config>service>vprn>sub-if>grp-if>dhcp>option

**Description** This command configures the Alcatel-Lucent vendor specific suboption of the DHCP relay packet.

### client-mac-address

**Syntax** [no] client-mac-address

**Context** config>service>vprn>if>dhcp>option  
config>service>vprn>if>dhcp>option>vendor  
config>service>vprn>sub-if>grp-if>dhcp>option>vendor

**Description** This command enables the sending of the MAC address in the Alcatel-Lucent vendor specific suboption of the DHCP relay packet.

The **no** form of the command disables the sending of the MAC address in the Alcatel-Lucent vendor specific suboption of the DHCP relay packet.

### sap-id

**Syntax** [no] sap-id

**Context** config>service>vprn>if>dhcp>option>vendor  
config>service>vprn>sub-if>grp-if>dhcp>option>vendor

**Description** This command enables the sending of the SAP ID in the Alcatel-Lucent vendor specific suboption of the DHCP relay packet.

The **no** form of the command disables the sending of the SAP ID in the Alcatel-Lucent vendor specific suboption of the DHCP relay packet.

### service-id

**Syntax** [no] service-id

**Context** config>service>vprn>if>dhcp>option>vendor  
config>service>vprn>sub-if>grp-if>dhcp>option>vendor

**Description** This command enables the sending of the service ID in the Alcatel-Lucent vendor specific suboption of the DHCP relay packet.

The **no** form of the command disables the sending of the service ID in the Alcatel-Lucent vendor specific suboption of the DHCP relay packet.

## string

<b>Syntax</b>	<b>[no] string</b> <i>text</i>
<b>Context</b>	config>service>vprn>if>dhcp>option>vendor config>service>vprn>sub-if>grp-if>dhcp>option>vendor
<b>Description</b>	This command specifies the vendor specific suboption string of the DHCP relay packet. The <b>no</b> form of the command returns the default value.
<b>Parameters</b>	<i>text</i> — The string can be any combination of ASCII characters up to 32 characters in length. If spaces are used in the string, enclose the entire string in quotation marks (“ ”).

## system-id

<b>Syntax</b>	<b>[no] system-id</b>
<b>Context</b>	config>service>vprn>if>dhcp>option>vendor config>service>vprn>sub-if>grp-if>dhcp>option>vendor
<b>Description</b>	This command specifies whether the system-id is encoded in the Alcatel-Lucent vendor specific sub-option of Option 82.
<b>Default</b>	None

## proxy-server

	<b>proxy-server</b>
<b>Context</b>	config>service>vpls>sap>dhcp config>subscr-mgmt>msap-policy>vpls-only>dhcp config>service>vprn>if>dhcp config>service>ies>if>dhcp config>service>vprn>sub-if>grp-if>dhcp config>service>ies>sub-if>grp-if>dhcp config>service>vprn>sub-if>dhcp config>service>vprn>sub-if>grp-if>ipv6>dhcp6 config>service>ies>sub-if>grp-if>ipv6>dhcp6 config>service>vprn>sub-if>ipv6>dhcp6 config>service>ies>sub-if>ipv6>dhcp6
<b>Description</b>	This command configures the DHCP proxy server.

## emulated-server

<b>Syntax</b>	<b>emulated-server</b> <i>ip-address</i> <b>no emulated-server</b>
<b>Context</b>	config>service>vprn>if>dhcp>proxy

```
config>service>vprn>sub-if>grp-if>dhcp>proxy-server
```

<b>Description</b>	<p>This command configures the IP address to be used as the DHCP server address in the context of this service. Typically, the configured address should be in the context of the subnet.</p> <p>The <b>no</b> form of this command reverts to the default setting. The local proxy server will not become operational without a specified emulated server address.</p>
<b>Parameters</b>	<p><i>ip-address</i> — Specifies the emulated server address.</p> <p><b>Default</b>      Note that for a retail interface, the default is the local interface.</p>

## lease-time

<b>Syntax</b>	<p><b>lease-time</b> [<b>days</b> <i>days</i>] [<b>hrs</b> <i>hours</i>] [<b>min</b> <i>minutes</i>] [<b>sec</b> <i>seconds</i>] [<b>override</b>]</p> <p><b>no lease-time</b></p>
<b>Context</b>	<pre>config&gt;service&gt;vprn&gt;if&gt;dhcp&gt;proxy config&gt;service&gt;vprn&gt;sub-if&gt;grp-if&gt;dhcp&gt;proxy-server</pre>
<b>Description</b>	<p>This command defines the length of lease-time that will be provided to DHCP clients. By default the local-proxy-server will always make use of the lease-time information provide by either a RADIUS or DHCP server.</p> <p>The no form of this command disables the use of the lease-time command. The local-proxy-server will use the lease-time offered by either a RADIUS or DHCP server.</p>
<b>Default</b>	7 days 0 hours 0 seconds
<b>Parameters</b>	<p><b>override</b> — Specifies that the local-proxy-server will use the configured lease-time information to provide DHCP clients.</p> <p><i>days</i> — Specifies the number of days that the given IP address is valid.</p> <p><b>Values</b>      0 — 3650</p> <p><i>hours</i> — Specifies the number of hours that the given IP address is valid.</p> <p><b>Values</b>      0 — 23</p> <p><i>minutes</i> — Specifies the number of minutes that the given IP address is valid.</p> <p><b>Values</b>      0 — 59</p> <p><i>seconds</i> — Specifies the number of seconds that the given IP address is valid.</p> <p><b>Values</b>      0 — 59</p>

## server

<b>Syntax</b>	<b>server</b> <i>server1</i> [ <i>server2</i> ...(up to 8 max)]
<b>Context</b>	<pre>config&gt;service&gt;vprn&gt;if&gt;dhcp config&gt;service&gt;vprn&gt;sub-if&gt;grp-if&gt;dhcp</pre>
<b>Description</b>	<p>This command specifies a list of servers where requests will be forwarded. The list of servers can entered as either IP addresses or fully qualified domain names. There must be at least one server</p>

specified for DHCP relay to work. If there are multiple servers then the request is forwarded to all of the servers in the list. There can be a maximum of 8 DHCP servers configured.

The flood command is applicable only in the VPLS case. There is a scenario with VPLS where the VPLS node only wants to add Option 82 information to the DHCP request to provider per-subscriber information, but it does not do full DHCP relay. In this case, the server is set to "flood". This means the DHCP request is still a broadcast and is sent through the VPLS domain. A node running at L3 further upstream then can perform the full L3 DHCP relay function.

<b>Default</b>	no server
<b>Parameters</b>	<i>server</i> — Specify the DHCP server IP address.

## host-connectivity-verify

<b>Syntax</b>	<b>host-connectivity-verify</b> [ <i>interval interval</i> ] [ <i>action {remove alarm}</i> ] [ <i>family family</i> ]
<b>Context</b>	config>service>vprn>if>sap config>service>vprn>sub-if>grp-if config>service>vprn>sub-if>grp-if>dhcp
<b>Description</b>	This command enables enables subscriber host connectivity verification on a given SAP within a service.  This tool will periodically scan all known hosts (from dhcp-state) and perform a UC ARP request. The subscriber host connectivity verification will maintain state (connected vs. not-connected) for all hosts.
<b>Default</b>	no host-connectivity-verify
<b>Parameters</b>	<b>interval <i>interval</i></b> — The interval, expressed in minutes, which specifies the time interval which all known sources should be verified. The actual rate is then dependent on number of known hosts and interval.  <b>Values</b> 1— 6000 ) Note that a zero value can be used by the SNMP agent to disable host-connectivity-verify.)  <b>action {remove   alarm}</b> — Defines the action taken on a subscriber host connectivity verification failure for a given host. The <b>remove</b> keyword raises an alarm and removes dhcp-state and releases all allocated resources (queues, table entries, etc.). DHCP-RELEASE will be signaled to corresponding DHCP server. Static hosts will never be removed. The <b>alarm</b> keyword raises an alarm indicating that the host is disconnected.  <b>family <i>family</i></b> — The family configuration allows the host connectivity checks to be performed for IPv4 endpoint, IPv6 endpoint or both. With family IPv6 configured, host connectivity checks will be performed on the global unicast address (assigned via SLAAC or DHCPv6 IA_NA) and link-local address of a Layer 3 RG or bridged hosts. In case of SLAAC assignment, host connectivity can only be performed if the /128 is known (via downstream ND). DHCPv6 PD assigned prefixes will be removed if link-local address is determined to be unreachable via “host connectivity check”. Reachability checks for GUA and link-local address will be done simultaneously.

## source-address

<b>Syntax</b>	<b>source-address</b> <i>ipv6-address</i> <b>no source-address</b>
<b>Context</b>	config>service>ies>if>ipv6>dhcp6 config>service>ies>sub-if>grp-if>ipv6>dhcp6 config>service>ies>sub-if>ipv6>dhcp6>relay config>service>ies>sub-if>grp-if>ipv6>dhcp6>relay config>service>vprn>sub-if>grp-if>ipv6>dhcp6>proxy-server config>service>vprn>sub-if>grp-if>ipv6>dhcp6>proxy-server config>service>vprn>if>ipv6>dhcp6>dhcp6-relay config>service>vprn>sub-if>grp-if>ipv6>dhcp6>relay config>service>vprn>sub-if>ipv6>dhcp6>relay
<b>Description</b>	This command configures the source IPv6 address of the DHCPv6 relay messages.
<b>Parameters</b>	<i>ipv6-address</i> — Specifies the source IPv6 address of the DHCPv6 relay messages.
<b>Values</b>	ipv6-address:    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

## dhcp6-server

<b>Syntax</b>	<b>[no] dhcp6-server</b>
<b>Context</b>	config>service>ies>if>ipv6
<b>Description</b>	This command enables the context to configure DHCPv6 server parameters for the IES interface. The <b>no</b> form of the command disables the DHCP6 server.

## max-nbr-of-leases

<b>Syntax</b>	<b>max-nbr-of-leases</b> <i>max-nbr-of-leases</i> <b>no max-nbr-of-leases</b>
<b>Context</b>	config>service>ies>if>ipv6>dhcp6-server
<b>Description</b>	This command configures the maximum number of lease states installed by the DHCP6 server function allowed on this interface. The <b>no</b> form of the command returns the value to the default.
<b>Default</b>	8000
<b>Parameters</b>	<i>max-nbr-of-leases</i> — Specifies the maximum number of lease states installed by the DHCP6 server function allowed on this interface.
<b>Values</b>	0 — 8000

## prefix-delegation

<b>Syntax</b>	<b>[no] prefix-delegation</b>
<b>Context</b>	config>service>ies>if>ipv6>dhcp6-server
<b>Description</b>	This command configures prefix delegation options for delegating a long-lived prefix from a delegating router to a requesting router, where the delegating router does not require knowledge about the topology of the links in the network to which the prefixes will be assigned.  The <b>no</b> form of the command disables prefix-delegation.

## prefix

<b>Syntax</b>	<b>[no] prefix</b> <i>ipv6-address/prefix-length</i>
<b>Context</b>	config>service>ies>if>ipv6>dhcp6-server>pfx-delegate
<b>Description</b>	This command specifies the IPv6 prefix that will be delegated by this system.
<b>Parameters</b>	<i>ipv6-address/prefix-length</i> — Specify the IPv6 address on the interface.
<b>Values</b>	<div> <div>ipv6-address/prefix: ipv6-address</div> <div> x:x:x:x:x:x:x (eight 16-bit pieces)  x:x:x:x:x:d.d.d  x [0 — FFFF]H  d [0 — 255]D </div> </div> <div> <div>prefix-length</div> <div>1 — 128</div> </div>

---

## Interface Commands

### local-proxy-arp

<b>Syntax</b>	<b>[no] local-proxy-arp</b>
<b>Context</b>	config>service>vprn>interface config>service>vprn>sub-if>grp-if
<b>Description</b>	This command enables local proxy ARP. When local proxy ARP is enabled on an IP interface, the system responds to all ARP requests for IP addresses belonging to the subnet with its own MAC address, and thus will become the forwarding point for all traffic between hosts in that subnet. When local-proxy-arp is enabled, ICMP redirects on the ports associated with the service are automatically blocked.
<b>Default</b>	no local-proxy-arp

### mac

<b>Syntax</b>	<b>[no] mac <i>ieee-mac-address</i></b>
<b>Context</b>	config>service>vprn>interface config>service>vprn>if>vrrp config>service>vprn>sub-if>grp-if
<b>Description</b>	This command assigns a specific MAC address to a VPRN IP interface. The <b>no</b> form of this command returns the MAC address of the IP interface to the default value.
<b>Default</b>	The physical MAC address associated with the Ethernet interface that the SAP is configured on.
<b>Parameters</b>	<i>ieee-mac-address</i> — Specifies the 48-bit MAC address for the static ARP in the form <i>aa:bb:cc:dd:ee:ff</i> or <i>aa-bb-cc-dd-ee-ff</i> where <i>aa</i> , <i>bb</i> , <i>cc</i> , <i>dd</i> , <i>ee</i> and <i>ff</i> are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

### proxy-arp-policy

<b>Syntax</b>	<b>[no] proxy-arp <i>policy-name</i> [<i>policy-name</i>...(up to 5 max)]</b>
<b>Context</b>	config>service>vprn>interface config>service>vprn>sub-if>grp-if
<b>Description</b>	This command enables a proxy ARP policy for the interface. The no form of this command disables the proxy ARP capability.
<b>Default</b>	no proxy-arp
<b>Parameters</b>	<i>policy-name</i> — The export route policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.



The specified name(s) must already be defined.

## redundant-interfacer

<b>Syntax</b>	<b>redundant-interface</b> <i>red-ip-int-name</i> <b>no redundant-interface</b>
<b>Context</b>	config>service>vprn config>service>vprn>sub-if>grp-if
<b>Description</b>	This command configures a redundant interface used for dual homing.
<b>Parameters</b>	<i>red-ip-int-name</i> — Specifies the redundant IP interface name.

## remote-proxy-arp

<b>Syntax</b>	<b>[no] remote-proxy-arp</b>
<b>Context</b>	config>service>vprn>interface config>service>vprn>sub-if>grp-if
<b>Description</b>	<p>This command enables remote proxy ARP on the interface.</p> <p>Remote proxy ARP is similar to proxy ARP. It allows the router to answer an ARP request on an interface for a subnet that is not provisioned on that interface. This allows the router to forward to the other subnet on behalf of the requester. To distinguish remote proxy ARP from local proxy ARP, local proxy ARP performs a similar function but only when the requested IP is on the receiving interface.</p>
<b>Default</b>	no remote-proxy-arp

---

## Subscriber Interface Commands

### subscriber-interface

<b>Syntax</b>	<b>[no] subscriber-interface</b> <i>ip-int-name</i>
<b>Context</b>	config>service>ies config>service>vprn
<b>Description</b>	<p>This command allows the operator to create special subscriber-based interfaces. It is used to contain multiple group interfaces. Multiple subnets associated with the subscriber interface can be applied to any of the contained group interfaces in any combination. The subscriber interface allows subnet sharing between group interfaces.</p> <p>Use the <b>no</b> form of the command to remove the subscriber interface.</p>
<b>Parameters</b>	<i>ip-int-name</i> — Specifies the name of the IP interface. Interface names can be from 1 to 32 alphanumeric characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

### group-interface

<b>Syntax</b>	<b>[no] group-interface</b> <i>ip-int-name</i>
<b>Context</b>	config>service>ies>sub-if
<b>Description</b>	<p>This command enables the context to configure a group interface. A group interface is an interface that may contain one or more SAPs. This interface is used in triple-play services where multiple SAPs are part of the same subnet.</p>
<b>Default</b>	none
<b>Parameters</b>	<i>ip-int-name</i> — Configures the interface group name. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

### authentication-policy

<b>Syntax</b>	<b>authentication-policy</b> <i>name</i> <b>no authentication-policy</b>
<b>Context</b>	config>service>ies>sub-if>grp-if
<b>Description</b>	<p>This command assigns a RADIUS authentication policy to the interface.</p> <p>The <b>no</b> form of this command removes the policy name from the group interface configuration.</p>
<b>Default</b>	no authentication-policy
<b>Parameters</b>	<i>name</i> — Specifies the authentication policy name. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

---

## Local User Database Commands

### local-user-db

<b>Syntax</b>	<b>local-user-db</b> <i>local-user-db-name</i> [ <b>create</b> ] <b>no local-user-db</b> <i>local-user-db-name</i>
<b>Context</b>	config>subscr-mgmt
<b>Description</b>	This command enables the context to configure a local user database.
<b>Default</b>	not enabled
<b>Parameters</b>	<i>local-user-db-name</i> — Specifies the name of a local user database.

### user-db

<b>Syntax</b>	<b>user-db</b> <i>local-user-db-name</i> <b>no user-db</b>
<b>Context</b>	config>service>ies>sub-if>grp-if>dhcp config>service>ies>sub-if>grp-if>dhcp6 config>service>ies>sub-if>grp-if>ipv6>dhcp6 config>service>ies>sub-if>grp-if>ipv6>dhcp6
<b>Description</b>	This command assigns a local user database.
<b>Default</b>	not enabled
<b>Parameters</b>	<i>local-user-db-name</i> — Specifies the name of a local user database.

### ipoe

<b>Syntax</b>	<b>ipoe</b>
<b>Context</b>	config>subscr-mgmt>loc-user-db
<b>Description</b>	This command configures IPoE host parameters.

### ppp

<b>Syntax</b>	<b>ppp</b>
<b>Context</b>	config>subscr-mgmt>loc-user-db
<b>Description</b>	This command configures PPP host parameters.

## mask

<b>Syntax</b>	<b>mask type</b> <i>dhcp-match-type</i> {[ <b>prefix-string</b> <i>prefix-string</i>   <b>prefix-length</b> <i>prefix-length</i> ] [ <b>suffix-string</b> <i>suffix-string</i>   <b>suffix-length</b> <i>suffix-length</i> ]} <b>no mask type</b> <i>dhcp-match-type</i>
<b>Context</b>	config>subscr-mgmt>loc-user-db>dhcp config>subscr-mgmt>loc-user-db>ppp config>subscr-mgmt>loc-user-db>ipoe
<b>Description</b>	This command configures the mask.
<b>Parameters</b>	<p><i>dhcp-match-type</i> — Specifies up to four matching types to identify a host.</p> <p><b>Values</b> DHCP: circuit-id, option60, remote-id, sap-id, string, system-id PPP: circuit-id, remote-id, service-name, username</p> <p><b>Values</b> <i>prefix-string prefix-string</i></p> <p>Specifies a substring that is stripped of the start of the incoming circuit ID before it is matched against the value configured in the DHCP or PPPOE circuit ID.</p> <p>This string can only contain printable ASCII characters. The “*” character is a wildcard that matches any substring. If a “\” character is masked, use the escape key so it becomes “\\”.</p> <p><b>Values</b> 127 characters maximum , * is wildcard.</p> <p><b>prefix-length prefix-length</b> — Specifies the number of characters to remove from the start of the incoming circuitId before it is matched against the value configured in the DHCP circuit ID.</p> <p><b>Values</b> 1— 127</p> <p><b>suffix-string suffix-string</b> — Specifies a substring that is stripped of the end of the incoming circuit ID before it is matched against the value configured in DHCP circuit ID.</p> <p>This string can only contain printable ASCII characters. The “*” character is a wildcard that matches any substring. If a “\” character is masked, use the escape key so it becomes “\\”.</p> <p><b>Values</b> 127 characters maximum</p> <p><b>suffix-length suffix-length</b> — Specifies the number of characters to remove from the end of the incoming circuit ID before it is matched against the value configured in the DHCP circuit ID.</p> <p><b>Values</b> 1— 127</p>

## host

<b>Syntax</b>	<b>host</b> <i>host-name</i> [ <b>create</b> ] <b>no host</b> <i>host-name</i>
<b>Context</b>	config>subscr-mgmt>loc-user-db>dhcp config>subscr-mgmt>loc-user-db>ppp
<b>Description</b>	This command defines a DHCP or PPP subscriber.
<b>Parameters</b>	<p><i>host-name</i> —</p> <p><b>create</b> — Keyword used to create the host name. The <b>create</b> keyword requirement can be enabled/disabled in the <b>environment&gt;create</b> context.</p>

## access-loop-encapsulation

<b>Syntax</b>	<b>[no] access-loop-encapsulation</b>
<b>Context</b>	config>subscr-mgmt>loc-user-db>ppp>host
<b>Description</b>	This command enables access loop information.

## encap-offset

<b>Syntax</b>	<b>encap-offset [type type]</b> <b>no encap-offset</b>
<b>Context</b>	config>subscr-mgmt>loc-user-db>ppp>host>ale
<b>Description</b>	This command configures the egress encapsulation offset.
<b>Parameters</b>	<b>type type</b> — Selects the encap type.  <b>Values</b> pppoa-llc, pppoa-null, pppoeoa-llc, pppoeoa-llc-fcs, pppoeoa-llc-tagged, pppoeoa-llc-tagged-fcs, pppoeoa-null, pppoeoa-null-fcs, pppoeoa-null-tagged, pppoeoa-null-tagged-fcs, ipoa-llc, ipoa-null, ipoeoa-llc, ipoeoa-llc-fcs, ipoeoa-llc-tagged, ipoeoa-llc-tagged-fcs, ipoeoa-null, ipoeoa-null-fcs, ipoeoa-null-tagged, ipoeoa-null-tagged-fcs, pppoe, pppoe-tagged, ipoe, ipoe-tagged

## rate-down

<b>Syntax</b>	<b>rate-down rate</b> <b>no rate-down</b>
<b>Context</b>	config>subscr-mgmt>loc-user-db>ppp>host>ale
<b>Description</b>	This command configures the last mile link downstream rate in the access loop.
<b>Parameters</b>	<b>rate</b> — Specifies the the last mile link downstream rate needed for proper (shaping) rate calculations and interleaving delay in the access loop.  <b>Values</b> 1 — 100000 kbps

## access-loop-information

<b>Syntax</b>	<b>access-loop-information</b>
<b>Context</b>	config>subscr-mgmt>loc-user-db>ppp>host>ali
<b>Description</b>	This command enables the context to configure access loop information in the local user database

### circuit-id

<b>Syntax</b>	<b>circuit-id sap-id</b> <b>circuit-id string</b> <i>ASCII string</i> <b>no circuit-id</b>
<b>Context</b>	config>subscr-mgmt>loc-user-db>ppp>host
<b>Description</b>	This command specifies a circuit-id for PPPoE hosts. A circuit-id received in PPPoE tags has precedence over the ludb specified circuit-id.
<b>Default</b>	no circuit-id
<b>Parameters</b>	<b>sap-id</b> — Specifies to use the SAP ID of the PPPoE session as the circuit ID.  <i>string ASCII string</i>  Specifies the circuit-id as a string, up to 63 characters. in length.

### remote-id

<b>Syntax</b>	<b>remote-id string mac</b> <b>remote-id string</b> <i>ASCII string</i> <b>no remote-id</b>
<b>Context</b>	config>subscr-mgmt>loc-user-db>ppp>host>ali
<b>Description</b>	This command specifies a remote-id for PPPoE hosts. A remote-id received in PPPoE tags has precedence over the ludb specified remote-id.
<b>Default</b>	no remote-id
<b>Parameters</b>	<i>string ASCII string</i> — specifies the circuit-id as a string, up to 63 characters. in length.  <b>mac</b> — specifies MAC address of the PPPoE session as the remote ID.

### acct-policy

<b>Syntax</b>	<b>acct-policy acct-policy-name</b> [ <b>duplicate acct-policy-name</b> ] <b>no acct-policy</b>
<b>Context</b>	config>subscr-mgmt>loc-user-db>ppp>hostconfig>subscr-mgmt>loc-user-db>dhcp>host config>subscr-mgmt>loc-user-db>dhcp>host
<b>Description</b>	This command specifies the accounting policy used for sending an Accounting Stop message to report RADIUS authentication failures of PPPoE sessions. A duplicate policy can be specified if a copy of the Accounting Stop message must be sent to another destination.  Reporting RADIUS authentication failures with an Accounting Stop message must be enabled in the RADIUS authentication policy (“send-acct-stop-on-fail”)  A duplicate RADIUS accounting policy can be specified if the accounting stop resulting from a RADIUS authentication failure must also be sent to a second RADIUS destination.
<b>Default</b>	no acct-policy

**Parameters** *acct-policy-name* — Specifies the name of a RADIUS accounting policy up to 32 characters in length.

## address

**Syntax** **address gi-address [scope scope]**  
**address ip-address[/prefix-length]**  
**address pool pool-name [secondary-pool sec-pool-name] [delimiter delimiter]**  
**address use-pool-from-client [delimiter delimiter]**  
**no address**

**Context** config>subscr-mgmt>loc-user-db>dhcp>host  
config>subscr-mgmt>loc-user-db>ppp>host

**Description** This command configures how the IP address is defined for this host.

When the user-db is used from a local-dhcp-server, then this command defines how to define the IP address the server will “offer” to the DHCP-client.

When the user-db is used for PPPoE authentication, the **gi-address** parameter cannot be used. A fixed IP address will then cause PPPoE to use this IP address. If no IP address is specified, the PPPoE will look for IP address by other means (DHCP). If a pool name is given, this pool will be sent in the DHCP request so that it can be used in by the DHCP server to determine which address to give to the host.

The **no** form of the command causes no IP address to be assigned to this host. In a user-db referred to from a local-dhcp-server, creating a host without address information will cause the matching client never to get an IP address.

**Default** no address

**Parameters** **gi-address** — When specified, the gi-address of the DHCP message is taken to look for a subnet in the local DHCP server. The first available free address of the subnet is taken and “offered” to the host. When **local-user-db** is used for PPPoE authentication, this has the same result as **no address**.

*ip-address* — Specifies the fixed IP address to use for this host.

*pool-name/sec-pool-name* — Specifies the primary (and secondary) pool (in the local DHCP server) to use to look for an available address. The first available IP address from any subnet in the pool will be used. When local-user-db is used for PPPoE authentication, this causes the specified pool name to be sent to the DHCP server in a vendor-specific suboption under Option 82

**use-pool-from-client** — Use the pool-name in the Option 82 vendor-specific sub-option.

**delimiter delimiter** — A single ascii character specifies the delimiter of separating primary and secondary pool names in option82 VSO.

## auth-domain-name

**Syntax** **auth-domain-name domain-name**  
**no auth-domain-name**

**Context** config>subscr-mgmt>loc-user-db>ipoe>host

config>subscr-mgmt>loc-user-db>ppp>host

**Description** This command configures the authentication policy of this host.

### auth-policy

**Syntax** **auth-policy** *policy-name*  
**no auth-policy**

**Context** config>subscr-mgmt>loc-user-db>dhcp>host

**Description** This command configures the authentication policy of this host and PPPoE hosts. This authentication policy is only used if no authentication policy is defined at the interface level. For DHCP hosts, the host entry should not contain any other information needed for setup of the host (IP address, ESM strings, etc.). For PPPoE hosts, the authentication policy configured here must have its pppoe-authentication-method set to **pap-chap**, otherwise the request will be dropped.

**Parameters** *policy-name* — Specifies the authentication policy name.

### force-ipv6cp

**Syntax** [**no**] **force-ipv6cp**

**Context** config>subscr-mgmt>loc-user-db>ppp>host

**Description** This command specifies if the IPv6 control protocol should be negotiated after PPP reaches the Network-Layer Protocol phase.

### diameter-application-policy

**Syntax** **diameter-application-policy** *policy-name*  
**no diameter-application-policy**

**Context** config>subscr-mgmt>loc-user-db>dhcp>host  
config>subscr-mgmt>loc-user-db>ipoe>host  
config>subscr-mgmt>loc-user-db>ppp>host  
config>service>ies>subscr-if>group-if  
config>service>vprn>subscr-if>group-if

**Description** This command configures the Diameter application policy.

**Parameters** *policy-name* — Specifies the Diameter application policy name.

### diameter-auth-policy

**Syntax** **diameter-auth-policy** *name*  
**no diameter-auth-policy**

**Context** config>subscr-mgmt>loc-user-db>dhcp>host



```
config>subscr-mgmt>loc-user-db>ipoe>host
config>subscr-mgmt>loc-user-db>ppp>host
config>service>ies>subscr-if>group-if
config>service>vprn>subscr-if>group-if
```

<b>Description</b>	This command is used to configure the Diameter NASREQ application policy to use for authentication.
<b>Parameters</b>	<i>name</i> — Specifies the name of the Diameter NASREQ application policy to use for authentication.

## auth-domain-name

<b>Syntax</b>	<b>auth-domain-name</b> <i>domain-name</i> <b>no auth-domain-name</b>
<b>Context</b>	config>subscr-mgmt>loc-user-db>dhcp>host
<b>Description</b>	This command sets the domain name which can be appended to user-name in RADIUS-authentication-request message for the given host.
<b>Parameters</b>	<i>domain-name</i> — Specifies the domain name to be appended to user-name in RADIUS-authentication-request message for the given host.

## host-identification

<b>Syntax</b>	<b>host-identification</b>
<b>Context</b>	config>subscr-mgmt>loc-user-db>dhcp>host config>subscr-mgmt>loc-user-db>ppp>host
<b>Description</b>	This command enables the context to configure host identification parameters.

## server

<b>Syntax</b>	<b>server</b> <i>ip-address</i> <b>no server</b>
<b>Context</b>	config>subscr-mgmt>loc-user-db>dhcp>host config>subscr-mgmt>loc-user-db>ipoe>host
<b>Description</b>	This command configures the IP address of the DHCP server in which to relay. The <b>no</b> form of the command removes the value from the configuration.
<b>Default</b>	no server
<b>Parameters</b>	<i>ip-address</i> — Specifies the IP address of the DHCP host server. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).

## circuit-id

<b>Syntax</b>	<b>circuit-id string</b> <i>ascii-string</i> <b>circuit-id hex</b> <i>hex-string</i> <b>no circuit-id</b>
<b>Context</b>	config>subscr-mgmt>loc-user-db>dhcp>host>host-ident config>subscr-mgmt>loc-user-db>ppp>host>host-ident
<b>Description</b>	This command specifies the circuit-id to match.
<b>Parameters</b>	<i>ascii-string</i> — specifies the circuit ID from the Option 82. <i>hex-string</i> — Specifies the circuit ID in hexadecimal format from the Option 82.
<b>Values</b>	0x0..0xFFFFFFFF (maximum 254 hex nibbles)

## derived-id

<b>Syntax</b>	<b>derived-id</b> <i>derived-id-string</i> <b>no derived-id</b>
<b>Context</b>	config>subscr-mgmt>loc-user-db>dhcp>host>host-ident
<b>Description</b>	This command configures an ASCII string that uniquely identifies a host, and is derived by a Python script from packet content available during a DHCP transaction.
<b>Parameters</b>	<i>derived-id-string</i> — Specifies the host ID to be derived by a python script from DHCP packets during a DHCP transaction up to 255 characters in length.

## encap-tag-range

<b>Syntax</b>	<b>encap-tag-range start-tag end-tag</b> <b>no encap-tag-range</b>
<b>Context</b>	config>subscr-mgmt>loc-user-db>dhcp>host>host-ident config>subscr-mgmt>loc-user-db>ppp>host>host-ident
<b>Description</b>	This command specifies a range of encapsulation tag as the host identifications. The encapsulation tag is dot1q or qinq on Ethernet port; VPI/VCI on ATM port. For dot1q, the start/end-tag is single number, range from 0-4094; for QinQ, the start/end-tag format is x.y, x or y could be “*”, which means ignore inner or outer tag; For ATM the start/end-tag format is vpi/vci, vpi or vci could be “*”, which means ignore VPI or VCI. Note: This command will only be used when “encap-tag-range” is configured as one of the match-list The <b>no</b> form of the command removes the values from the configuration.
<b>Default</b>	none
<b>Parameters</b>	<b>start-tag start-tag</b> — Specifies the value of the start label in the range of SAP’s allowed on this host.
<b>Values</b>	<i>start-tag</i> 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]

```

**end-tag** *end-tag* — Specifies the value of the end label in the range of SAP's allowed on this host.

<b>Values</b>	<i>end-tag</i>	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]

## mac

<b>Syntax</b>	<b>mac</b> <i>ieee-address</i> <b>no mac</b>
<b>Context</b>	config>subscr-mgmt>loc-user-db>dhcp>host>host-ident config>subscr-mgmt>loc-user-db>ppp>host>host-ident
<b>Description</b>	This command specifies the MAC address to match.
<b>Parameters</b>	<i>ieee-address</i> — Specifies the 48-bit MAC address in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee, and ff are hexadecimal numbers.

## options6

<b>Syntax</b>	<b>options6</b>
<b>Context</b>	config>subscr-mgmt>loc-user-db>ppp>host config>subscr-mgmt>loc-user-db>dhcp>host config>subscr-mgmt>loc-user-db>ipoe>host>options
<b>Description</b>	This command enables the context to configure IPv6 DNS server information in the local user database

## option60

<b>Syntax</b>	<b>option60</b> <i>hex-string</i> <b>no option60</b>
<b>Context</b>	config>subscr-mgmt>loc-user-db>dhcp>host>host-ident
<b>Description</b>	This command specifies the Vendor-Identifying Vendor Option to match. Option 60 is encoded as Type-Length-Value (TLV). The <i>hex-string</i> portion of Option 60 in the received DHCP request is used

## Local User Database Commands

for matching. Only the first 32 bytes can be defined here. If Option 60 from the message is longer, those bytes are ignored.

<b>Default</b>	no option60
<b>Parameters</b>	<i>hex-string</i> — Specifies the hex value of this option.
<b>Values</b>	0x0..0xFFFFFFFF...(maximum 254 hex nibbles)

## remote-id

<b>Syntax</b>	<b>remote-id</b> <i>hex-string</i> <b>remote-id</b> <i>ascii-string</i> <b>no remote-id</b>
<b>Context</b>	config>subscr-mgmt>loc-user-db>dhcp>host>host-ident config>subscr-mgmt>loc-user-db>ppp>host>host-ident
<b>Description</b>	This command specifies the remote id of this host. The <b>no</b> form of this command returns the system to the default.
<b>Default</b>	no remote-id
<b>Parameters</b>	<i>remote-id</i> — Specifies the remote-id.

## service-name

<b>Syntax</b>	<b>service-name</b> <i>service-name</i> <b>no service-name</b>
<b>Context</b>	config>subscr-mgmt>loc-user-db>ppp>host>host-ident
<b>Description</b>	This command specifies the service-name tag in PADI and/or PADR packets to match for PPPoE hosts.
<b>Parameters</b>	<i>service-name</i> — Specifies a PPPoE service name, up to 255 characters maximum.

## sap-id

<b>Syntax</b>	<b>sap-id</b> <i>sap-id</i> <b>no sap-id</b>
<b>Context</b>	config>subscr-mgmt>loc-user-db>dhcp>host>host-ident
<b>Description</b>	This command specifies the SAP ID from the Alcatel Vendor Specific Sub-option in Option 82 to match.
<b>Parameters</b>	<i>sap-id</i> — Specifies a SAP ID, up to 255 characters maximum.

## service-id

<b>Syntax</b>	<b>service-id</b> <i>service-id</i> <b>no service-id</b>
<b>Context</b>	config>subscr-mgmt>loc-user-db>dhcp>host>host-ident
<b>Description</b>	This command specifies an existing service ID from the Alcatel Vendor Specific Sub-Option in Option 82 to match.
<b>Parameters</b>	<i>service-id</i> — Specifies an existing service ID. <b>Values</b> 1 — 2147483647

## string

<b>Syntax</b>	<b>string</b> <i>string</i> <b>no string</b>
<b>Context</b>	config>subscr-mgmt>loc-user-db>dhcp>host>host-ident
<b>Description</b>	This command specifies the string from the Alcatel Vendor Specific Sub-Option in Option 82 to match.
<b>Parameters</b>	<i>string</i> — Specifies the string, up to 255 characters maximum.

## system-id

<b>Syntax</b>	<b>system-id</b> <i>system-id</i> <b>no system-id</b>
<b>Context</b>	config>subscr-mgmt>loc-user-db>dhcp>host>host-ident
<b>Description</b>	This command specifies the system ID from the Alcatel Vendor Specific Sub-Option in Option 82 to match.
<b>Parameters</b>	<i>system-id</i> — Specifies the system ID, up to 255 characters maximum.

## username

<b>Syntax</b>	<b>username</b> <i>user-name</i> <b>username</b> <i>user-name</i> [ <b>no-domain</b> ] <b>username</b> <i>user-name</i> <b>domain-only</b> <b>no username</b>
<b>Context</b>	config>subscr-mgmt>loc-user-db>ppp>host>host-ident
<b>Description</b>	This command specifies how the username is specified.
<b>Parameters</b>	<i>username</i> — Specifies the user name of this host. <b>no-domain</b> — No username is specified.

**domain-only** — Only the domain part of the username is specified, for example, alcatel-lucent.com.

## identification-strings

<b>Syntax</b>	<b>identification-strings</b> <i>option-number</i> [ <b>create</b> ] <b>no identification-strings</b>
<b>Context</b>	config>subscr-mgmt>loc-user-db>dhcp>host config>subscr-mgmt>loc-user-db>ppp>host
<b>Description</b>	This command specifies identification strings for the subscriber. This is useful when the server is centralized with Enhanced Subscriber Management (ESM) in a lower level in the network. These strings will be parsed by a downstream Python script or they can be used literally if the “strings-from-option” option in the <b>config&gt;subscriber-mgmt&gt;sub-ident-policy</b> context is set to this option number. In this case, the option number may be set to any allowed number (between 224 and 254 is suggested, as these are not dedicated to specific purposes). If the option number is not given, a default value of 254 is used. Note, for PPPoE only, if the local user database is attached to the PPPoE node under the group interface and not to a local DHCP server, the strings will be used internally so the option number is not used.
<b>Default</b>	254
<b>Parameters</b>	<i>option-number</i> — Specifies identification strings for the subscriber <b>Values</b> 1 — 254

## ancp-string

<b>Syntax</b>	<b>ancp-string</b> <i>ancp-string</i> <b>no ancp-string</b>
<b>Context</b>	config>subscr-mgmt>loc-user-db>dhcp>host>ident-strings config>subscr-mgmt>loc-user-db>ppp>host>ident-strings
<b>Description</b>	This command specifies the ANCP string which is encoded in the identification strings.
<b>Parameters</b>	<i>ancp-string</i> — Specifies the the ANCP string, up to 63 characters, maximum.

## app-profile-string

<b>Syntax</b>	<b>app-profile-string</b> <i>app-profile-string</i> <b>no app-profile-string</b>
<b>Context</b>	config>subscr-mgmt>loc-user-db>dhcp>host>ident-strings config>subscr-mgmt>loc-user-db>ppp>host>ident-strings
<b>Description</b>	This command specifies the application profile string which is encoded in the identification strings.
<b>Parameters</b>	<i>app-profile-string</i> — Specifies the the application profile string, up to 16 characters, maximum.

## category-map

<b>Syntax</b>	<b>category-map</b> <i>category-map-name</i> <b>no category-map</b> <i>category-map-name</i>
<b>Context</b>	config>subscr-mgmt>loc-user-db>dhcp>host>ident-strings config>subscr-mgmt>loc-user-db>ppp>host>ident-strings
<b>Description</b>	This command specifies the category map name.
<b>Default</b>	none
<b>Parameters</b>	<i>category-map-name</i> — Specifies an existing category map name up to 32 characters in length.

## inter-dest-id

<b>Syntax</b>	<b>inter-dest-id</b> <i>intermediate-destination-id</i> <b>no inter-dest-id</b>
<b>Context</b>	config>subscr-mgmt>loc-user-db>dhcp>host>ident-strings config>subscr-mgmt>loc-user-db>ppp>host>ident-strings
<b>Description</b>	This command specifies the intermediate destination identifier which is encoded in the identification strings.
<b>Parameters</b>	<i>intermediate-destination-id</i> — Specifies the intermediate destination identifier, up to 32 characters, maximum.

## sla-profile-string

<b>Syntax</b>	<b>sla-profile-string</b> <i>sla-profile-string</i> <b>no sla-profile-string</b>
<b>Context</b>	config>subscr-mgmt>loc-user-db>dhcp>host>ident-strings config>subscr-mgmt>loc-user-db>ppp>host>ident-strings
<b>Description</b>	This command specifies the SLA profile string which is encoded in the identification strings.
<b>Parameters</b>	<i>sla-profile-string</i> — Specifies the SLA profile string, up to 16 characters, maximum.

## sub-profile-string

<b>Syntax</b>	<b>sub-profile-string</b> <i>sub-profile-string</i> <b>no sub-profile-string</b>
<b>Context</b>	config>subscr-mgmt>loc-user-db>dhcp>host>ident-strings config>subscr-mgmt>loc-user-db>ppp>host>ident-strings
<b>Description</b>	This command specifies the subscriber profile string which is encoded in the identification strings.
<b>Parameters</b>	<i>sub-profile-string</i> — Specifies the subscriber profile string, up to 16 characters, maximum.

## subscriber-id

<b>Syntax</b>	<b>subscriber-id</b> <i>sub-ident-string</i> <b>no subscriber-id</b>
<b>Context</b>	config>subscr-mgmt>loc-user-db>dhcp>host>ident-strings config>subscr-mgmt>loc-user-db>ppp>host>ident-strings
<b>Description</b>	This command specifies the subscriber ID which is encoded in the identification strings.
<b>Parameters</b>	<i>sub-ident-string</i> — Specifies the subscriber ID string, up to 32 characters, maximum.

## interface

<b>Syntax</b>	<b>interface</b> <i>ip-int-name</i> <b>service-id</b> <i>service-id</i> <b>no interface</b>
<b>Context</b>	config>subscr-mgmt>loc-user-db>ppp>host
<b>Description</b>	This command configures the interface where PPP sessions are terminated. The no form of the command reverts to the default.
<b>Default</b>	none
<b>Parameters</b>	<i>ip-int-name</i> — Specifies the name of the group interface where the PPP sessions are established <b>service-id</b> <i>service-id</i> — Specifies the service ID of the service where the PPP sessions are established.

## ipv6-address

<b>Syntax</b>	<b>ipv6-address</b> <i>ipv6-address</i> <b>no ipv6-address</b>
<b>Context</b>	config>subscr-mgmt>loc-user-db>dhcp>host config>subscr-mgmt>loc-user-db>ppp>host
<b>Description</b>	This command configures static DHCPv6 IA-NA address for the host. This address is delegated to the client as /128 via DHCPv6 proxy function within the 7x50. This IP address must not be part of any DHCP pool within internal DHCP server. The <b>no</b> form of the command removes the IPv6 address from the host configuration.
<b>Parameters</b>	<i>ipv6-address</i> — Specifies the IPv6 address.
<b>Values</b>	ipv6-address:    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-delegated-prefix

Syntax	ipv6-delegated-prefix <i>ipv6-prefix/prefix-length</i> no ipv6-delegated-prefix		
Context	config>subscr-mgmt>loc-user-db>dhcp>host config>subscr-mgmt>loc-user-db>ipoe>host config>subscr-mgmt>loc-user-db>ppp>host		
Description	This command configures static DHCPv6 IA-PD prefix for the host. This prefix can be further delegated by the host itself to its clients. The prefix length is restricted to 48 to 64 bits. This prefix must not be part of any DHCP pool within internal DHCP server.		
Default	no ipv6-delegated-prefix		
Parameters	<i>ipv6-address</i> — Specifies the IPv6 address.		
	Values	ipv6-address:	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
		prefix-length	[48..64]

## ipv6-delegated-prefix-pool

<b>Syntax</b>	<b>ipv6-delegated-prefix-pool</b> <i>pool-name</i> <b>no ipv6-delegated-prefix-pool</b>
<b>Context</b>	config>subscr-mgmt>loc-user-db>dhcp>host config>subscr-mgmt>loc-user-db>ipoe>host config>subscr-mgmt>loc-user-db>ppp>host
<b>Description</b>	This command configures the pool name that will be used in DHCPv6 server for DHCPv6 IA-PD prefix selection.  The <b>no</b> form of the command removes the pool name from the configuration.
<b>Parameters</b>	<i>pool-name</i> — Specifies the pool name ot be assigned to the delegated prefix pool.

## ipv6-slaac-prefix

<b>Syntax</b>	<b>ipv6-slaac-prefix</b> <i>ipv6-prefix/prefix-length</i> <b>no ipv6-slaac-prefix</b>
<b>Context</b>	config>subscr-mgmt>loc-user-db>dhcp>host config>subscr-mgmt>loc-user-db>ppp>host config>subscr-mgmt>loc-user-db>ipoe>host
<b>Description</b>	This command configures static IPv6 SLAAC prefix (PIO) for the host. The host will assign an IPv6 address to itself based on this prefix. The prefix length is 64 bits.  The <b>no</b> form of the command removes the static IPv6 SLAAC prefix (PIO) for the host from the configuration.

<b>Default</b>	no ipv6-slaac-prefix
<b>Parameters</b>	<i>ipv6-prefix/prefix-length</i> — Specifies the IPv6 address and prefix length.
<b>Values</b>	<div> <div>&lt;ipv6-prefix/prefix*&gt; : ipv6-prefix</div> <div> <div>x:x:x:x:x:x:x (eight 16-bit pieces)</div> <div>x:x:x:x:x:d.d.d.d</div> <div>x [0..FFFF]H</div> <div>d [0..255]D</div> </div> </div> <div> <div>prefix-length</div> <div>64</div> </div>

### ipv6-slaac-prefix-pool

<b>Syntax</b>	<b>ipv6-slaac-prefix-pool</b> <i>pool</i> <b>no ipv6-slaac-prefix-pool</b>
<b>Context</b>	config>subscr-mgmt>loc-user-db>ipoe>host config>subscr-mgmt>loc-user-db>ppp>host
<b>Description</b>	This command configures the IPv6 slaac prefix pool of this host.

### ipv6-delegated-prefix-length

<b>Syntax</b>	<b>ipv6-delegated-prefix-length</b> <i>bits</i> <b>no ipv6-delegated-prefix-length</b>
<b>Context</b>	configure>subscr-mgmt>local-user-db>dhcp>host configure>subscr-mgmt>local-user-db>ppp>host
<b>Description</b>	This command allows configuration of delegated prefix length via local user database.
<b>Default</b>	no ipv6-delegated-prefix-length
<b>Parameters</b>	<i>bits</i> — Specifies the delegated prefix length in bits.
<b>Values</b>	48..64

### ipv6-prefix

<b>Syntax</b>	<b>ipv6-prefix</b> <i>ipv6-prefix/prefix-length</i> <b>no ipv6-prefix</b>
<b>Context</b>	config>subscr-mgmt>loc-user-db>dhcp>host
<b>Description</b>	This command configures the IPv6 prefix and length of this host. The <b>no</b> form of the command removes the IPv6 prefix and length of this host from the configuration.
<b>Parameters</b>	<i>ipv6-prefix/prefix-length</i> — Specifies the IPv6 prefix of this host.
<b>Values</b>	<div> <div>ipv6-prefix/prefix: ipv6-prefix</div> <div> <div>x:x:x:x:x:x:x (eight 16-bit pieces)</div> <div>x:x:x:x:x:d.d.d.d</div> </div> </div>

	x [0..FFFF]H
	d [0..255]D
prefix-length	48..64

## ipv6-wan-address-pool

<b>Syntax</b>	<b>ipv6-wan-address-pool</b> <i>pool-name</i> <b>no ipv6-wan-address-pool</b>
<b>Context</b>	config>subscr-mgmt>loc-user-db>dhcp>host
<b>Description</b>	This command configures the pool name that will be used in DHCPv6 server for DHCPv6 IA-PA address selection.  The <b>no</b> form of the command removes the pool name from the configuration.
<b>Default</b>	no ipv6-wan-address-pool
<b>Parameters</b>	<i>pool-name</i> — Specifies the WAN address pool up to 32 characters in length.

## link-address

<b>Syntax</b>	<b>link-address</b> <i>ipv6-address</i> <b>no link-address</b>								
<b>Context</b>	config>subscr-mgmt>loc-user-db>ipoe>host config>service>ies>sub-if>grp-if>ipv6>dhcp6>proxy-server config>service>vprn>sub-if>grp-if>ipv6>dhcp6>proxy-server								
<b>Description</b>	This command allows link-address selection based on the host entry in LUDB.  The link-address is a field in DHCP6 Relay-Forward message that is used in DHCP6 server to select the IPv6 address (IA-NA) or IPv6 prefix (IA-PD) from a pool with configured prefix range covering the link-address. The selection scope is the pool or a prefix range within the pool.								
<b>Parameters</b>	<i>ipv6-address</i> — Specifies the link-address.  <table> <tr> <td><b>Values</b></td> <td>&lt;ipv6-address&gt;    ipv6-address    - x:x:x:x:x:x:x:x (eight 16-bit pieces)</td> </tr> <tr> <td></td> <td>x:x:x:x:x:x.d.d.d.d</td> </tr> <tr> <td></td> <td>x - [0..FFFF]H</td> </tr> <tr> <td></td> <td>d - [0..255]D</td> </tr> </table>	<b>Values</b>	<ipv6-address>    ipv6-address    - x:x:x:x:x:x:x:x (eight 16-bit pieces)		x:x:x:x:x:x.d.d.d.d		x - [0..FFFF]H		d - [0..255]D
<b>Values</b>	<ipv6-address>    ipv6-address    - x:x:x:x:x:x:x:x (eight 16-bit pieces)								
	x:x:x:x:x:x.d.d.d.d								
	x - [0..FFFF]H								
	d - [0..255]D								

## server

<b>Syntax</b>	<b>server</b> <i>ipv6-address</i> [ <i>ipv6-address</i> ...(upto 8 max)] <b>no server</b>
<b>Context</b>	config>service>ies>subscriber-interface>ipv6>dhcp6 config>service>vprn>subscriber-interface>ipv6>dhcp6

## match-radius-proxy-cache

<b>Syntax</b>	<b>match-radius-proxy-cache</b>
<b>Context</b>	config>subscr-mgmt>loc-user-db>ipoe>host
<b>Description</b>	This command enables the context to configure RADIUS proxy cache match parameters.

## delete-hold-time

<b>Syntax</b>	<b>delete-hold-time</b> <i>seconds</i> <b>no delete-hold-time</b>
<b>Context</b>	config>subscr-mgmt>loc-user-db>dhcp>host>match-rdprox-cache
<b>Description</b>	This command specifies the time for which the UE state (including ESM host) is maintained after an accounting-stop has been received for the UE from RADIUS client on the AP. This allows UE state to exist when a UE moves to a new AP and re-authenticates within the hold-time, thereby providing seamless mobility. The hold-time is canceled if re-authentication or accounting-start is received for the UE before expiry. If the hold-time expires, the UE state is deleted.  The no form of the command reverts to the default.
<b>Default</b>	no delete-hold-time
<b>Parameters</b>	<i>seconds</i> — Specifies the time for which the UE state will be held after an accounting-stop has been received for the UE.  <b>Values</b> 1 — 600

## fail-action

<b>Syntax</b>	<b>fail-action {continue drop}</b> <b>no fail-action</b>
<b>Context</b>	config>subscr-mgmt>loc-user-db>ipoe>host>match-radprox-cache
<b>Description</b>	This command specifies the action to take when no match is found in the cache.  The no form of the command reverts to the default.
<b>Default</b>	drop

## mac-format

<b>Syntax</b>	<b>mac-format</b> <i>mac-format</i> <b>no mac-format</b>
<b>Context</b>	config>subscr-mgmt>loc-user-db>ipoe>host>match-radprox-cache
<b>Description</b>	This command specifies how a MAC address is represented.

## match

<b>Syntax</b>	<b>match {circuit-id mac remote-id}</b> <b>match option</b> [1..254] [ <b>option6</b> [1..65535]]> <b>match option6</b> [1..65535] <b>no match</b>
<b>Context</b>	config>subscr-mgmt>loc-user-db>ipoe>host>match-radprox-cache
<b>Description</b>	This command specifies in what DHCPv6 option to retrieve the value to be used as lookup key in the RADIUS proxy cache.
<b>Default</b>	none

## server

<b>Syntax</b>	<b>server [service service-id] name server-name]</b> <b>no server</b>
<b>Context</b>	config>subscr-mgmt>loc-user-db>ipoe>host>match-radprox-cache
<b>Description</b>	This command specifies the RADIUS proxy server.

## ipv6-lease-times

<b>Syntax</b>	<b>[no] ipv6-lease-times</b>
<b>Context</b>	config>subscr-mgmt>loc-user-db>ipoe>host config>subscr-mgmt>loc-user-db>ppp>host
<b>Description</b>	This command configures the lease times for DHCPv6.

## preferred-lifetime

<b>Syntax</b>	<b>preferred-lifetime [days days] [hrs hrs] [min min] [sec sec]</b> <b>preferred-lifetime infinite</b> <b>no preferred-lifetime</b>
<b>Context</b>	config>service>vprn>sub-if>ipv6>dhcp6>proxy config>service>ies>sub-if>ipv6>dhcp6>proxy config>service>vprn>sub-if>grp-if>ipv6>dhcp6>proxy config>service>ies>sub-if>grp-if>ipv6>dhcp6>proxy
<b>Description</b>	This command configures the preferred-lifetime for DHCPv6 leases (address/prefix) in a proxy-scenario (For example address/prefix obtained from Radius)  Preferred lifetime is the length of time that a valid address/prefix is preferred (for example, the time until deprecation).
<b>Default</b>	hrs 1

<b>Parameters</b>	<i>infinite</i> — Specifies that the valid lifetime is infinite.
<b>Values</b>	0xffffffff
	[ <b>days</b> <i>days</i> ][ <b>hrs</b> <i>hours</i> ] [ <b>min</b> <i>minutes</i> ] [ <b>sec</b> <i>seconds</i> ] — Specifies the preferred lifetime.
<b>Values</b>	days: [0..3650] hours: [0..23] minutes: [0..59] seconds: [0..59]

### rebind-timer

<b>Syntax</b>	<b>rebind-timer</b> [ <b>days</b> <i>days</i> ] [ <b>hrs</b> <i>hrs</i> ] [ <b>min</b> <i>min</i> ] [ <b>sec</b> <i>sec</i> ] <b>no rebind-timer</b>
<b>Context</b>	config>subscr-mgmt>loc-user-db>ipoe>host>ipv6-lease-times config>subscr-mgmt>loc-user-db>ppp>host>ipv6-lease-times config>service>ies>sub-if>grp-if>ipv6>dhcp6>proxy-server config>service>vprn>sub-if>grp-if>ipv6>dhcp6>proxy-server
<b>Description</b>	<p>This command configures the lease rebind timer (T2) via LUBD.</p> <p>The T2 time is the time at which the client contacts any available addressing authority to extend the lifetimes of DHCPv6 leases. T2 is a time duration relative to the current time expressed in units of seconds.</p> <p>The IP addressing authority controls the time at which the client contacts the addressing authority to extend the lifetimes on assigned addresses/prefixes through the T1 and T2 parameters assigned to an IA. At time T1 for an IA, the client initiates a Renew/Reply message exchange to extend the lifetimes on any addresses in the IA. The client includes an IA option with all addresses/prefixes currently assigned to the IA in its Renew message. Recommended values for T1 and T2 are .5 and .8 times the shortest preferred lifetime of the addresses/prefixes in the IA that the addressing authority is willing to extend, respectively.</p> <p>The configured rebind timer should always be longer than or equal to the renew timer.</p> <p>The T1 and T2 are carried in the IPv6 address option that is within the IA.</p>
<b>Default</b>	rebind-timer min 48
<b>Parameters</b>	<b>days</b> <i>days</i> ][ <b>hrs</b> <i>hours</i> ] [ <b>min</b> <i>minutes</i> ] [ <b>sec</b> <i>seconds</i> ] — Specifies the preferred lifetime.
<b>Values</b>	days: [0..7] hours: [0..23] minutes: [0..59] seconds: [0..59]

### renew-timer

<b>Syntax</b>	<b>renew-timer</b> [ <b>days</b> <i>days</i> ] [ <b>hrs</b> <i>hrs</i> ] [ <b>min</b> <i>min</i> ] [ <b>sec</b> <i>sec</i> ] <b>no renew-timer</b>
<b>Context</b>	config>subscr-mgmt>loc-user-db>ipoe>host>ipv6-lease-times config>subscr-mgmt>loc-user-db>ppp>host>ipv6-lease-times

```
config>service>ies>sub-if>grp-if>ipv6>dhcp6>proxy-server
config>service>vprn>sub-if>grp-if>ipv6>dhcp6>proxy-server
```

Description	<p>This command configures the lease renew time (T1) via LUDB.</p> <p>The T1 is the time at which the client contacts the addressing authority to extend the lifetimes of the DHCPv6 leases (addresses/prefixes). T1 is a time duration relative to the current time expressed in units of seconds.</p> <p>The IP addressing authority controls the time at which the client contacts the addressing authority to extend the lifetimes on assigned addresses through the T1 and T2 parameters assigned to an IA. At time T1 for an IA, the client initiates a Renew/Reply message exchange to extend the lifetimes on any addresses in the IA. The client includes an IA option with all addresses currently assigned to the IA in its Renew message. Recommended values for T1 and T2 are .5 and .8 times the shortest preferred lifetime of the addresses in the IA that the addressing authority is willing to extend, respectively.</p> <p>The configured renew timer should always be smaller than or equal to the rebind timer.</p> <p>The T1 and T2 are carried in the IPv6 address option that is within the IA.</p>												
Default	renew-timer min 30												
Parameters	<p><b>days</b> <i>days</i> [<b>hrs</b> <i>hours</i>] [<b>min</b> <i>minutes</i>] [<b>sec</b> <i>seconds</i>] — Specifies the preferred lifetime.</p> <table><tr><td><b>Values</b></td><td>days:</td><td>[0..7]</td></tr><tr><td></td><td>hours:</td><td>[0..23]</td></tr><tr><td></td><td>minutes:</td><td>[0..59]</td></tr><tr><td></td><td>seconds:</td><td>[0..59]</td></tr></table>	<b>Values</b>	days:	[0..7]		hours:	[0..23]		minutes:	[0..59]		seconds:	[0..59]
<b>Values</b>	days:	[0..7]											
	hours:	[0..23]											
	minutes:	[0..59]											
	seconds:	[0..59]											

## server-id

<b>Syntax</b>	<pre>server-id duid-en hex <i>hex-string</i> server-id duid-en string <i>ascii-string</i> server-id duid-ll no server-id</pre>		
<b>Context</b>	<pre>config&gt;service&gt;ies&gt;sub-if&gt;grp-if&gt;ipv6&gt;dhcp6 config&gt;service&gt;ies&gt;sub-if&gt;grp-if&gt;ipv6&gt;dhcp6&gt;proxy-server config&gt;service&gt;vprn&gt;sub-if&gt;grp-if&gt;ipv6&gt;dhcp6&gt;proxy-server</pre>		
<b>Description</b>	<p>This command allows operator to customize the “server-id” attribute of a DHCPv6 message from the DHCPv6 proxy server (such as DHCPv6 advertise and reply). By default, the server-id uses DUID-ll derive from the chassis link-layer address. Operators have the option to use a unique identifier by using DUID-en (vendor based on enterprise number). There is a maximum length associated with the customizable hex-string and ascii-string.</p>		
<b>Default</b>	server-id duid-ll		
<b>Parameters</b>	<p><b>duid-en hex</b> <i>hex-string</i> — Specifies a DUID system ID in a hex format.</p> <table> <tr> <td><b>Values</b></td><td>0x0..0xFFFFFFFF...(max 116 hex nibbles)</td></tr> </table> <p><b>duid-en string</b> <i>ascii-string</i> — Specifies a DUID system ID in an ASCII format up to 58 characters.</p> <p><b>duid-ll</b> — Specifies that the DUID system ID is derived from the system link layer address.</p>	<b>Values</b>	0x0..0xFFFFFFFF...(max 116 hex nibbles)
<b>Values</b>	0x0..0xFFFFFFFF...(max 116 hex nibbles)		

## valid-lifetime

<b>Syntax</b>	<b>valid-lifetime</b> [ <b>days</b> <i>days</i> ] [ <b>hrs</b> <i>hrs</i> ] [ <b>min</b> <i>min</i> ] [ <b>sec</b> <i>sec</i> ] <b>valid-lifetime</b> <i>infinite</i> <b>no valid-lifetime</b>
<b>Context</b>	config>subscr-mgmt>loc-user-db>ipoe>host>ipv6-lease-times config>subscr-mgmt>loc-user-db>ppp>host>ipv6-lease-times config>service>ies>sub-if>grp-if>ipv6>dhcp6>proxy-server config>service>vprn>sub-if>grp-if>ipv6>dhcp6>proxy-server
<b>Description</b>	<p>This command configured valid-lifetime for DHCPv6 lease (address/prefix).</p> <p>Valid lifetime is the the length of time an address/prefix remains in the valid state (i.e., the time until invalidation). The valid lifetime must be greater than or equal to the preferred lifetime. When the valid lifetime expires, the address/prefix becomes invalid and must not be used in communications. RFC 2461, sec 6.2.1 recommends default value of 30 days.</p> <p>Each address/prefix assigned to the client has associated preferred and valid lifetimes specified by the address assignment authority (DHCP Server, Radius, ESM). To request an extension of the lifetimes assigned to an address, the client sends a Renew message to the addressing authority. The addressing authority sends a Reply message to the client with the new lifetimes, allowing the client to continue to use the address/prefix without interruption.</p> <p>The lifetimes are transmitted from the addressing authority to the client in the IA option on the top level (not the address or prefix level).</p>
<b>Default</b>	valid-lifetime days 1
<b>Parameters</b>	<i>infinite</i> — Specifies that the valid lifetime is infinite.
	<b>Values</b> 0xffffffff
	<b>days</b> <i>days</i> ] [ <b>hrs</b> <i>hours</i> ] [ <b>min</b> <i>minutes</i> ] [ <b>sec</b> <i>seconds</i> ] — Specifies the preffered lifetime.
	<b>Values</b> days: [0..3650] hours: [0..23] minutes: [0..59] seconds: [0..59]

## l2tp

<b>Syntax</b>	<b>l2tp</b>
<b>Context</b>	config>subscr-mgmt>loc-user-db>ppp>host
<b>Description</b>	This command configures L2TP for the host.

## group

<b>Syntax</b>	<b>group</b> <i>tunnel-group-name</i> [ <b>service-id</b> <i>service-id</i> ] <b>no group</b>
<b>Context</b>	config>subscr-mgmt>loc-user-db>ppp>host>l2tp



<b>Description</b>	This command configures the L2TP tunnel group. The tunnel-group-name is configured in the <b>config&gt;router&gt;l2tp</b> context. Refer to the 7750 SR OS Router Configuration Guide.
<b>Parameters</b>	<i>tunnel-group-name</i> — Specifies an existing tunnel L2TP group up to 63 characters in length. <b>service-id</b> <i>service-id</i> — [Specifies an existing service ID or service name.
<b>Values</b>	service-id: 1 — 214748364 svc-name: A string up to 64 characters in length.

## authentication-policy

<b>Syntax</b>	<b>authentication-policy</b> <i>policy-name</i> <b>no authentication-policy</b>
<b>Context</b>	config>subscr-mgmt>loc-user-db>ppp>host
<b>Description</b>	This command configures the authentication policy for the host. A host name with name “default” will be matched when all other hosts do not match.

## pado-delay

<b>Syntax</b>	<b>pado-delay</b> <i>deci-seconds</i> <b>no pado-delay</b>
<b>Context</b>	config>subscr-mgmt>loc-user-db>ppp>host
<b>Description</b>	This command configures the delay timeout before sending a PPPoE Active Discovery Offer. (PADO)
<b>Parameters</b>	<i>deci-seconds</i> — Specifies the delay timeout before sending a PADO.
<b>Values</b>	1 — 30

## mask

<b>Syntax</b>	<b>mask type</b> <i>ppp-match-type</i> {[ <b>prefix-string</b> <i>prefix-string</i>   <b>prefix-length</b> <i>prefix-length</i> ] [ <b>suffix-string</b> <i>suffix-string</i>   <b>suffix-length</b> <i>suffix-length</i> ]} <b>no mask type</b> <i>ppp-match-type</i>
<b>Context</b>	config>subscr-mgmt>loc-user-db>ppp
<b>Syntax</b>	<b>mask type</b> <i>dhcp-match-type</i> {[ <b>prefix-string</b> <i>prefix-string</i>   <b>prefix-length</b> <i>prefix-length</i> ] [ <b>suffix-string</b> <i>suffix-string</i>   <b>suffix-length</b> <i>suffix-length</i> ]} <b>no mask type</b> <i>dhcp-match-type</i>
<b>Context</b>	config>subscr-mgmt>loc-user-db>dhcp
<b>Description</b>	This command configures the mask.
<b>Parameters</b>	<i>ppp-match-type</i> — Specifies the sub-option inserted by the PPPoE intermediate agent.
<b>Values</b>	circuit-id, remote-id, service-name, username

*dhcp-match-type* — The data type represents the type of matching done to identify a DHCP host.

**Values** circuit-id , option60 , remote-id , sap-id , string , system-id

**prefix-string** *prefix-string* — Specifies a substring that is stripped of the start of the incoming circuit ID before it is matched against the value configured in the DHCP or PPPOE circuit ID.

This string can only contain printable ASCII characters. The “\*” character is a wildcard that matches any substring. If a “\” character is masked, use the escape key so it becomes “\\”.

**Values** 127 characters maximum , \* is wildcard.

**prefix-length** *prefix-length* — Specifies the number of characters to remove from the start of the incoming circuitId before it is matched against the value configured in the circuit ID.

**Values** 1— 127

**suffix-string** *suffix-string* — Specifies a substring that is stripped of the end of the incoming circuit ID before it is matched against the value configured in circuit ID.

This string can only contain printable ASCII characters. The “\*” character is a wildcard that matches any substring. If a “\” character is masked, use the escape key so it becomes “\\”.

**Values** 127 characters maximum

**suffix-length** *suffix-length* — Specifies the number of characters to remove from the end of the incoming circuit ID before it is matched against the value configured in the circuit ID.

**Values** 1— 127

## match-list

<b>Syntax</b>	<b>match-list</b> <i>match-type-1</i> [ <i>match-type-2</i> ...(up to 4 max)] <b>no match-list</b>
<b>Context</b>	config>subscr-mgmt>loc-user-db>dhcp config>subscr-mgmt>loc-user-db>ppp config>subscr-mgmt>loc-user-db>ipoe
<b>Description</b>	This command specifies the type of matching done to identify a host. There are different match-types for PPPoE or IPoE hosts of which a maximum of 4 can be specified.
<b>Default</b>	no match-list
<b>Parameters</b>	<p><i>match-type-x</i> — Specifies up to four matching types to identify a host.</p> <p><b>Values</b> circuit-id, derived-id, dual-stack-remote-id, encaps-tag-range, mac, option60, remote-id, sap-id, service-id, string, system-id</p> <p><b>circuit-id</b> — Specifies the DHCP4 option (82,1) or DHCP6 option 18.</p> <p><b>mac</b> — Specifies the MAC address of the client. Chaddr in DHCP4 and DUID in IPv6.</p> <p><b>option60</b> — Specifies the DHCP4 option 60.</p> <p><b>remote-id</b> — Specifies the DHCP4 option (82,2) or DHCP6 option 37 (Note that the format of remote-id in IPv6 is different that the format of remote-id in IPv4; IPv6 remote-id contains enterprise-id field that is also honored in matching.)</p> <p><b>dual-stack-remote-id</b> — Specifies the enterprise-id in v6 Remote-id will be stripped off before LUDB matching is performed. Processing of IPv4 Remote-id remains unchanged. This will allow a single host entry in LUDB for dual-stack host where host identification is performed</p>

based on the Remote-id field.

**sap-id** — Specifies the SAP ID on which DHCPv4 packet are received. The sap-id is inserted as ALU VSO (82,9,4) by the DHCPv4 relay in 7x50. This is enabled via configuration under the vendor-specific-option CLI hierarchy of the DHCPv4 relay. Since the dhcp-relay configuration is enabled under the group-interface CLI hierarchy, the group-interface and the service-id must be known before the sap-id can be used for LUDB match.

**encap-tag-range** — Specifies the VLAN tags.

**service-id** — Specifies the service-id of the ingress SAP for DHCPv4 packets. The service-id is inserted as ALU VSO (82,9,3) by the DHCPv4 relay in 7x50. This is enabled via configuration under the vendor-specific-option CLI hierarchy of the DHCPv4 relay.

**string** — Specifies the custom string configured under the vendor-specific-option CLI hierarchy of the DHCPv4 relay. The string is inserted as ALU VSO (82,9,5) by the DHCPv4 relay in 7x50. Since the dhcp-relay configuration is enabled under the group-interface CLI hierarchy, the group-interface and the service-id must be known before the string can be used for LUDB match.

**system-id** — Specifies the system-id of the node name configured under the system>name CLI hierarchy. The system-id is inserted as ALU VSO (82,9,1) by the DHCPv4 relay in 7x50. This is enabled via configuration under the vendor-specific-option CLI hierarchy of the DHCPv4 relay. Since the dhcp-relay configuration is enabled under the group-interface CLI hierarchy, the group-interface and the service-id must be known before the system-id can be used for LUDB match.

**derived-id** — Specifies the value extracted by Python script during processing of DHCP Discover/Solicit/Request/Renew/Rebind Messages (client to server bound messages). The value is stored in the DHCP Transaction Cache (DTC) in a variable named alc.dtc.derivedId. This value has a lifespan of a DHCP transaction (a single pair of messages exchanged between the client and the server; for example DHCP Discover and DHCP Offer).

## password

<b>Syntax</b>	<b>password {ignore   chap <i>string</i>   pap <i>string</i>} [hash hash2] no password</b>
<b>Context</b>	config>subscr-mgmt>loc-user-db>ppp>host
<b>Description</b>	<p>This command specifies a password type or configures password string for <b>pap</b> or <b>chap</b>. The pap and chap passwords are stored in a hashed format in the config files. The <b>hash hash2</b> optional keywords are used for config execution.</p> <p>This command will only be interpreted if the local user database is connected directly to the PPPoE node under the VPRN/IES group interface. It is not used if the local user database is accessed by a local DHCP server.</p>
<b>Parameters</b>	<p><b>ignore</b> — Specifies that the password will be ignored, in which case authentication will always succeed, independent of the password used by the PPPoE client. The client must still perform authentication.</p> <p><b>chap <i>string</i></b> — Specifies that the password for Challenge-Handshake Authentication Protocol (CHAP) is used. Only a password received with the CHAP protocol will be accepted.</p> <p><b>pap <i>string</i></b> — Specifies that the Password Authentication Protocol (PAP) is used. Only a password received with the PAP protocol will be accepted, even though the CHAP protocol will be proposed to the client first because it is unknown at the time of the offer which password type will be allowed to the client.</p> <p><b>hash hash2</b> — Specifies hashing scheme.</p>

## pre-auth-policy

<b>Syntax</b>	<b>pre-auth-policy</b> <i>policy-name</i> <b>no pre-auth-policy</b>
<b>Context</b>	config>subscr-mgmt>loc-user-db>ppp>host
<b>Description</b>	This command configures the pre-authentication policy of this host.

## retail-service-id

<b>Syntax</b>	<b>retail-service-id</b> <i>service-id</i> <b>no retail-service-id</b>
<b>Context</b>	config>subscr-mgmt>loc-user-db>ipoe>host config>subscr-mgmt>loc-user-db>ppp>host
<b>Description</b>	This command indicates the service ID of the retailer VPRN service to which this session belongs. If the value of this object is non-zero, the session belongs to a retailer VPRN. The <b>no</b> form of the command removes the service ID from the configuration.
<b>Default</b>	no retail-service-id
<b>Parameters</b>	<i>service-id</i> — Specifies the the retailer service ID.
<b>Values</b>	service-id: 1 — 2147483647 service-name: Service name up to 64 characters in length.

## server6

<b>Syntax</b>	<b>server6</b> <i>ipv6-address</i> <b>no server6</b>
<b>Context</b>	config>subscr-mgmt>loc-user-db>ipoe>host
<b>Description</b>	This command allows DHCP6 server selection based on the host entry in LUDB. The configured DHCP6 message must reference one if the v6 addressees configured under the <b>configure&gt;service&gt;vprn&gt;sub-if&gt;grp-if&gt;ipv6&gt;dhcpv6&gt;relay</b> or <b>configure&gt;service&gt;ies&gt;sub-if&gt;grp-if&gt;ipv6&gt;dhcpv6&gt;relay</b> context.
<b>Default</b>	no server6
<b>Parameters</b>	<i>ipv6-address</i> — Specifies the the retailer service ID.
<b>Values</b>	service-id: 1 — 2147483647 service-name: Service name up to 64 characters in length.

## MLPPP on LNS Commands

### accept-mrru

<b>Syntax</b>	<b>[no] accept-mrru</b>
<b>Context</b>	configure>subscr-mgt>ppp-policy>mlppp
<b>Description</b>	<p>This command is applicable only to LAC. MRRU option is an indication that the session is of MLPPPoX type. The 7750 LAC will never initiate MRRU option in LCP negotiation process. However, it will respond to MRRU negotiation request by the client.</p> <p>This command provides an option to specifically enable or disable negotiation of MLPPPoX on a capture SAP level or on a group-interface level.</p>
<b>Default</b>	no accept-mrru — The MRRU option in LCP will not be negotiated by LAC.

### admin-state

<b>Syntax</b>	<b>admin-state {up   down}</b> <b>no admin-state</b>
<b>Context</b>	configure>router>l2tp>group>tunnel>mlppp configure>service>vpn>l2tp>group>tunnel>mlppp
<b>Description</b>	<p>This command is applicable only to LNS.</p> <p>The tunnel can be explicitly activated (assuming that the parent group is in a no shutdown state) or deactivated by the <b>up</b> and <b>down</b> keywords.</p> <p>If case that there is no admin-state configured, the tunnel will inherit its administrative state from its parent (group).</p>
<b>Default</b>	<p>no admin-state — Tunnel administrative state is inherited from the group.</p> <p><b>up</b> — Tunnel is in administratively up.</p> <p><b>down</b> — Tunnel is administratively down.</p>

### encap-offset

<b>Syntax</b>	<b>encap-offset [type <i>encap-type</i>]</b> <b>no encap-offset</b>
<b>Context</b>	configure>subscriber-mgmt>local-user-db>ppp>host>access-loop
<b>Description</b>	<p>This command is applicable within the LAC/LNS context. It provides the last mile link encapsulation information that is needed for proper (shaping) rate calculations and interleaving delay in the last mile.</p> <p>The encapsulation value will be taken from the following sources in the order of priority:</p>

- Statically provisioned value in local user database (LUDB).
- RADIUS
- PPPoE tags on LAC or ICRQ message (RFC 5515) on LNS

In case that the encapsulation information is not provided by any of the existing means (LUDB, RADIUS, AVP signaling, PPPoE Tags), then by default pppoea-null encapsulation will be in effect.

The following values are supported encapsulation values on LNS in the 7750.

encap-type:

```
pppoa-llc LLC (NLPID) PPPoA encapsulation.
pppoa-null VC-MUX PPPoA encapsulation.
pppoeoa-llc LLC/SNAP based bridged Ethernet PPPoEoA encapsulation without FCS.
pppoeoa-llc-fcs LLC/SNAP based bridged Ethernet PPPoEoA encapsulation with FCS.
pppoeoa-null VC-MUX PPPoEoA encapsulation without FCS.
pppoeoa-null-fcs VC-MUX PPPoEoA encapsulation with FCS.
pppoe tagged PPPoE encapsulation.
pppoe-tagged Tagged PPPoE Encapsulation.
```

The values are not supported encapsulation values on LNS in the 7750.

```
pppoeoa-llc-tagged
pppoeoa-llc-tagged-fcs
pppoeoa-null-tagged
pppoeoa-null-tagged-fcs
ipoa-llc
ipoa-null
ipoeoa-llc
ipoeoa-llc-fcs
ipoeoa-llc-tagged
ipoeoa-llc-tagged-fcs
ipoeoa-null
ipoeoa-null-fcs
ipoeoa-null-tagged
ipoeoa-null-tagged-fcs
ipoe
ipoe-tagged
```

**Default** no encap-offset No offset is configured.

## endpoint

**Syntax** **endpoint ip** *ip-address*  
**endpoint mac** *ieee-address*  
**endpoint system-ip**  
**endpoint system-mac**  
**no endpoint**

**Context** configure>router>l2tp>group>mlppp  
configure>router>l2tp>group>tunnel>mlppp  
configure>service>vpn>l2tp>group>mlppp  
configure>service>vpn>l2tp>group>tunnel>mlppp  
configure>subscr-mgt>ppp-policy>mlppp

<b>Description</b>	<p>When configured under the l2tp hierarchy, this command is applicable to LNS.</p> <p>Within the ppp-policy, this command is applicable only to LAC.</p> <p>The endpoint, according to RFC 1990, represents the system transmitting the packet. It is used during MLPPPoX negotiation phase to distinguish this peer from all others.</p> <p>In the case that the client rejects the endpoint option during LCP negotiation, the LAC and the LNS must be able to negotiate the LCP session without the endpoint option.</p> <p>The <b>no</b> form of this command disables sending endpoint option in LCP negotiation.</p>
<b>Default</b>	no endpoint
<b>Parameters</b>	<p><b>ip</b> <i>ip-address</i> — Specifies the IPv4 address (class 2)</p> <p><b>system-ip</b> — Specifies to use the system IPv4 address (class 2)</p> <p><b>mac</b> <i>ieee-address</i> — Specifies the MAC address of the interface (class 3).</p> <p><b>system-mac</b> — Specifies to use the MAC address of the system (class 3)</p>

## interleave

<b>Syntax</b>	<b>[no] interleave</b>
<b>Context</b>	<pre>configure&gt;router&gt;l2tp&gt;group&gt;mlppp configure&gt;service&gt;vpn&gt;l2tp&gt;group&gt;mlppp</pre>
<b>Description</b>	<p>This command is applicable only to LNS. Interleaving is supported only on MLPPPoX bundles that contain a single member link. If more than one link is present in the MLPPPoX bundle, interleaving will be automatically disabled and a TRAP/log (tmnxMlpppBundleIndicatorsChange) will be generated.</p> <p>The minimum supported rate of the link on which interleaving is performed is 1kbps.</p> <p>If configured at this level, interleaving will be enabled on all tunnels within the group, unless it is explicitly disable per tunnel.</p>
<b>Default</b>	no interleave — Interleaving per group is disabled.

## interleave

<b>Syntax</b>	<b>interleave {always   never}</b> <b>no interleave</b>
<b>Context</b>	<pre>configure&gt;router&gt;l2tp&gt;group&gt;tunnel&gt;mlppp configure&gt;service&gt;vpn&gt;l2tp&gt;group&gt;tunnel&gt;mlppp</pre>
<b>Description</b>	<p>This command is applicable only to LNS. Interleaving is supported only on MLPPPoX bundles that contain a single member link. If more than one link is present in the MLPPPoX bundle, interleaving will be automatically disabled and a TRAP/log (tmnxMlpppBundleIndicatorsChange ) will be generated.</p> <p>The minimum supported rate of the link on which interleaving is performed is 1kbps.</p>

Interleaving configured on this level will overwrite the configuration option under the group hierarchy. If the no form of the command is configured for interleaving at this level, the interleaving configuration will inherit the configuration option configured under the l2tp group.

<b>Default</b>	no interleave — Interleaving configuration is inherited from the group.
<b>Parameters</b>	<p><b>always</b> — Always perform interleaving on single linked MLPPPoX sessions within this tunnel, regardless of the configuration option for interleaving under the group level.</p> <p><b>never</b> — Never perform interleaving on single linked MLPPPoX sessions within this tunnel, regardless of the configuration option for interleaving under the group level.</p>

## load-balance-method

<b>Syntax</b>	<b>load-balance-method {session   tunnel}</b> <b>no load-balance-method</b>
<b>Context</b>	configure>router>l2tp>group configure>router>l2tp>group>tunnel configure>service>vpnn>l2tp>group configure>service>vpnn>l2tp>group>tunnel
<b>Description</b>	<p>This command is applicable only to LNS. By default traffic load balancing between the BB-ISAs is based on sessions. Each session is individually assigned to an BB-ISA during session establishment phase.</p> <p>By introducing MLPPPoX, all sessions of a bundle must be terminated on the same LNS BB-ISA. This is necessary for two reasons:</p> <ul style="list-style-type: none"> <li>• QoS in the carrier IOM has a uniform view of the subscriber</li> <li>• a single BB-ISA is responsible for MLPPPoX encapsulation/fragmentation for a given bundle.</li> </ul> <p>Therefore, if fragmentation is enabled, load-balancing per tunnel must be configured. In the per tunnel load-balancing mode, all sessions within the same tunnel are terminated on the same LNS BB-ISA.</p> <p>In the case that we have MLPPPoX sessions with a single member link, both load-balancing methods are valid.</p> <p>The <b>no</b> form of this command set the per session load balancing.</p>
<b>Default</b>	session — Per session load balancing is enabled by default.
<b>Parameters</b>	<p><b>session</b> — Traffic load balancing between the LNS BB-ISAs is based on individual PPPoE sessions.</p> <p><b>tunnel</b> — Traffic load balancing between the LNS BB-ISAs is based on tunnels.</p>

## max-fragment-delay

<b>Syntax</b>	<b>max-fragment-delay <i>mili-seconds</i></b> <b>no max-fragment-delay</b>
---------------	---



<b>Context</b>	configure>router>l2tp>group>mlppp configure>router>l2tp>group>tunnel>mlppp configure>service>vpn>l2tp>group>mlppp configure>service>vpn>l2tp>group>tunnel>mlppp
<b>Description</b>	This command is applicable only to LNS. It determines the maximum fragment delay caused by the transmission that will be imposed on a link.  Fragmentation can be used to interleave high priority packet in-between low priority fragments on a MLPPPoX session with a single link or on a MLPPPoX session with multiple links to better load balance traffic over multiple member links.
<b>Default</b>	no max-fragment-delay — Fragmentation is disabled.
<b>Parameters</b>	<i>mili-seconds</i> — Specifies the interval in mili-seconds.  <b>Values</b> 5-1000ms

## max-link

<b>Syntaxs</b>	<b>max-links</b> <i>max-links</i> <b>no max-links</b>
<b>Context</b>	configure>router>l2tp>group>mlppp configure>router>l2tp>group>tunnel>mlppp configure>service>vpn>l2tp>group>mlppp configure>service>vpn>l2tp>group>tunnel>mlppp
<b>Description</b>	This command is applicable only to LNS. It determines the maximum number of links that can be put in a bundle.  Any attempt of a session to join a bundle that is above the max-link limit will be rejected.  If interleaving is configured, it is recommended that max-links be set to 1 or a ?o?version of the command is used (no max-links). Both have the same effect.  The configuration under the tunnel hierarchy will override the configuration under the group hierarchy.  The <b>no</b> form of this command limits the number of links in the bundle to 1.
<b>Default</b>	no max-links — A single link per bundle is allowed.
<b>Parameters</b>	<i>max-links</i> — Specifies the maximum number of links in a bundle.  <b>Values</b> 1 — 8

## reassembly-timeout

<b>Syntax</b>	<b>reassembly-timeout</b> {{100   1000} milliseconds} <b>no reassembly-timeout</b>
<b>Context</b>	configure>router>l2tp>group>mlppp configure>router>l2tp>group>tunnel>mlppp configure>service>vpn>l2tp>group>mlppp

```
configure>service>vpn>l2tp>group>tunnel>mlppp
```

<b>Description</b>	<p>This command is applicable only to LNS. It determines the time during which the LNS keeps fragments of the same packet in the buffer before it discards them. The assumption is that if the fragments do not arrive within certain time, the chance is that they were lost somewhere in the network. In this case the partial packet cannot be reassembled and all fragments that has arrived up to this point and are stored in the buffer will be discarded in order to free up the buffer. Otherwise, a condition will arise in which partial packets will be held in the buffer until the buffer is exhausted.</p> <p>The configuration under the tunnel hierarchy will override the configuration under the group hierarchy.</p> <p>The <b>no</b> form of this command also sets the time-out to 1000ms.</p>
<b>Default</b>	1000
<b>Parameters</b>	{{100   1000} milliseconds} — Specifies the reassembly timeout value.

## rate-down

<b>Syntax</b>	<b>rate-down rate</b> <b>no rate-down</b>
<b>Context</b>	configure>subscriber-mgmt>local-user-db>ppp>host>access-loop
<b>Description</b>	<p>This command is applicable to LAC and LNS. It provides the last mile link rate in the downstream direction that is needed for proper shaping and calculating the interleaving delay.</p> <p>The rate information in the last mile will be taken from the following sources in the order of priority:</p> <ul style="list-style-type: none"> <li>• Statically provisioned value in local user database (LUDB).</li> <li>• RADIUS.</li> <li>• PPPoE tags on LAC or ICRQ message (RFC 5515) /ICCN message (TX Connect Seed) on LNS.</li> </ul>
<b>Default</b>	no rate-down
<b>Parameters</b>	<i>rate</i> — Specifies last mile link downstream rate in the access loop <b>Values</b> 1 — 100000 kbps

## short-sequence-numbers

<b>Syntax</b>	<b>[no] short-sequence-numbers</b>
<b>Context</b>	configure>subscr-mgt>ppp-policy>mlppp
<b>Description</b>	<p>This command enables a peer request to send short sequence numbers. This command is applicable to LAC and LNS. By default, MLPPPoX will negotiate 24bit long sequence numbers. This command allows this to be changed to shorter, 12-bit sequence numbers.</p>
<b>Default</b>	short-sequence-numbers

---

## Show Commands

### id

<b>Syntax</b>	<b>id</b> <i>service-id</i>
<b>Context</b>	show>service
<b>Description</b>	This command displays information for a particular service-id.
<b>Parameters</b>	<i>service-id</i> — The unique service identification number that identifies the service in the service domain.  <b>all</b> — Display detailed information about the service. <b>base</b> — Display basic service information. <b>fdb</b> — Display FDB entries. <b>labels</b> — Display labels being used by this service. <b>sap</b> — Display SAPs associated to the service. <b>sdp</b> — Display SDPs associated with the service. <b>split-horizon-group</b> — Display split horizon group information. <b>stp</b> — Display STP information.

### dhcp

<b>Syntax</b>	<b>dhcp</b>
<b>Context</b>	show>router show>service>id
<b>Description</b>	This command enables the context to show DHCP statistics.

### dhcp6

<b>Syntax</b>	<b>dhcp6</b>
<b>Context</b>	show>router show>service>id
<b>Description</b>	This command enables the context to show DHCP6 statistics.

dhcp6

<b>Syntax</b>	<b>dhcp6</b>
<b>Context</b>	show>system
<b>Description</b>	This command displays system-wide DHCPv6 configuration information.

**Sample Output**

```
A:PE-1# show system dhcp6
=====
DHCP6 system
=====
Global NoAddrsAvail status : esm-relay server
=====
```

lease-state

<b>Syntax</b>	<b>lease-state [detail]</b> <b>lease-state [detail] interface <i>interface-name</i></b> <b>lease-state [detail] ipv6-address <i>ipv6-prefix[/prefix-length]</i></b> <b>lease-state [detail] mac <i>ieee-address</i></b>
<b>Context</b>	show>service>id>dhcp6
<b>Description</b>	This command displays DHCP6 lease state related information.

**Sample Output**

```
*A:Dut-C# show service id 202 dhcp6 lease-state
=====
DHCP lease state table, service 202
=====
IP Address      Mac Address      Sap/Sdp Id      Remaining      Lease      MC
                  LifeTime      Origin      Stdby
-----
1::/120
                  1/1/6          30d33h12m      DHCP
-----
Number of lease states : 1
=====
*A:Dut-C#

*A:Dut-C# show service id 202 dhcp6 lease-state detail
=====
DHCP lease states for service 202
=====
Service ID      : 202
IP Address      : 1::/120
Mac Address     :
Interface       : ip-11.3.202.3
SAP             : 1/1/6
Remaining Lifetime : 30d33h12m
Persistence Key  : N/A
```

```

Sub-Ident          : ""
Sub-Profile-String : ""
SLA-Profile-String : ""
Lease ANCP-String  : ""
Dhcp6 ClientId (DUID): 0101
Dhcp6 IAID         : 1
Dhcp6 IAID Type    : prefix
Dhcp6 Client Ip    : FE80::200:FF:FE00:202

ServerLeaseStart   : 09/01/2002 04:27:00
ServerLastRenew    : 09/01/2002 04:27:00
ServerLeaseEnd     : 10/01/2002 04:27:00
-----
Number of lease states : 1
=====
*A:Dut-C#

```

## statistics

<b>Syntax</b>	<b>statistics</b> [ <b>sap</b> <i>sap-id</i> ]   [ <b>sdp</b> [ <i>sdp-id</i> [: <i>vc-id</i> ]]   <b>interface</b> <i>ip-int-name</i> ]
<b>Context</b>	show>service>id>dhcp show>router>dhcp
<b>Description</b>	This command displays statistics for DHCP relay and DHCP snooping. If no IP address or interface name is specified, then all configured interfaces are displayed. If an IP address or interface name is specified, then only data regarding the specified interface is displayed.
<b>Parameters</b>	<i>sap-id</i> — Specifies the physical port identifier portion of the SAP definition. See <a href="#">Common Service Commands on page 2062</a> for <i>sap-id</i> command syntax. <i>sdp-id</i> — The SDP ID to be shown. <b>Values</b> 1 — 17407 <i>vc-id</i> — The virtual circuit ID on the ID to be shown. <b>Values</b> 1 — 4294967295 <i>ip-int-name</i>   <i>ip-address</i> — Displays statistics for the specified IP interface.
<b>Output</b>	<b>Show DHCP Statistics Output</b> — The following table describes the output fields for DHCP statistics.

Label	Description
Received Packets	The number of packets received from the DHCP clients.
Transmitted Packets	The number of packets transmitted to the DHCP clients.
Received Malformed Packets	The number of malformed packets received from the DHCP clients.

Label	Description (Continued)
Received Untrusted Packets	The number of untrusted packets received from the DHCP clients.
Client Packets Discarded	The number of packets received from the DHCP clients that were discarded.
Client Packets Relayed	The number of packets received from the DHCP clients that were forwarded.
Client Packets Snooped	The number of packets received from the DHCP clients that were snooped.
Server Packets Discarded	The number of packets received from the DHCP server that were discarded.
Server Packets Relayed	The number of packets received from the DHCP server that were forwarded.
Server Packets Snooped	The number of packets received from the DHCP server that were snooped.

### Sample Output

```
A:ALA-A# show router 1000 dhcp statistics
=====
DHCP Global Statistics (Service: 1000)
=====
Rx Packets                : 16000
Tx Packets                : 15041
Rx Malformed Packets      : 0
Rx Untrusted Packets      : 0
Client Packets Discarded  : 423
Client Packets Relayed    : 0
Client Packets Snooped    : 0
Client Packets Proxied (RADIUS) : 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
=====
A:ALA-A#
```

## summary

<b>Syntax</b>	<b>summary</b>
<b>Context</b>	show>router>dhcp show>service>id>dhcp
<b>Description</b>	Display the status of the DHCP Relay and DHCP Snooping functions on each interface.

**Output** **Show DHCP Summary Output** — The following table describes the output fields for DHCP summary.

Label	Description
Interface Name	Name of the router interface.
ARP Populate	Indicates whether ARP populate is enabled.
Used/Provided	Indicates the number of used and provided DHCP leases.
Info Option	Indicates whether Option 82 processing is enabled on the interface.
Admin State	Indicates the administrative state.

### Sample Output

```
A:ALA-48>show>router>dhcp# summary
=====
Interface Name          Arp      Used/    Info    Admin
                        Populate Provided Option  State
-----
ccaiesif                No        0/0      Keep    Down
ccanet6                 No        0/0      Keep    Down
iesBundle               No        0/0      Keep    Up
spokeSDP-test          No        0/0      Keep    Down
test                    No        0/0      Keep    Up
test1                   No        0/0      Keep    Up
test2                   No        0/0      Keep    Up
testA                   No        0/0      Keep    Up
testB                   No        0/0      Keep    Up
testIES                 No        0/0      Keep    Up
to-web                  No        0/0      Keep    Up
-----
Interfaces: 11
=====
A:ALA-48>show>router>dhcp#
```

## virtual-subnet

- Syntax** **virtual-subnet subscriber** *sub-ident*  
**virtual-subnet** [**sap** *sap-id*]
- Context** show>service>id
- Description** This command displays currently recorded default gateway and subnets for all virtual subnets enabled for DHCPv4 hosts in the specified service.
- Parameters** **subscriber** *sub-ident* — Displays information relating to the specified subscriber ID.  
**sap** *sap-id* — Displays information relating to the specified SAP ID.

### Sample Output

```
show service id 500 virtual-subnet
```

## Show Commands

```
=====
Virtual subnets in service 500
=====
Subscriber                : 00:20:fc:1e:cd:52|1/1/9:200
-----
Default router            : 192.168.100.254
Subnet                    : 192.168.100.0/24
SAP                       : 1/1/9:200
-----
No. of subnets: 1
=====
```

## statistics

<b>Syntax</b>	<b>statistics [interface <i>ip-int-name</i>]</b>
<b>Context</b>	show>router>dhcp6 show>service>id>dhcp6
<b>Description</b>	This command displays statistics for DHCP relay and DHCP snooping.

### Sample Output

```
A:ALA-A# show router 1000 dhcp statistics
=====
DHCP Global Statistics (Service: 1000)
=====
Rx Packets                : 16000
Tx Packets                : 15041
Rx Malformed Packets     : 0
Rx Untrusted Packets     : 0
Client Packets Discarded  : 423
Client Packets Relayed    : 0
Client Packets Snooped    : 0
Client Packets Proxied (RADIUS) : 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
=====
A:ALA-A#
```

## summary

<b>Syntax</b>	<b>summary</b>
<b>Context</b>	show>router>dhcp6 show>service>id>dhcp6
<b>Description</b>	Display the status of the DHCP6 relay and DHCP snooping functions on each interface.



**Output**    **Show DHCP6 Summary Output** — The following table describes the output fields for DHCP6 summary.

Label	Description
Interface Name	Name of the router interface.
Nbr. Resol.	Indicates whether or not neighbor resolution is enabled.
Used/Provided	Indicates the number of used and provided DHCP leases.
Admin State	Indicates the administrative state.
Oper State	Indicates the operational state.

### Sample Output

```
*A:Dut-C# show router dhcp6 summary
=====
DHCP6 Summary (Router: Base)
=====
Interface Name      Nbr      Used/Max Relay   Admin  Oper Relay
  SapId      Resol.   Used/Max Server  Admin  Oper Server
-----
ip-1.1.1.10        No          0/0             Down   Down
  sap:1/1/5                0/8000             Down   Down
ip-11.3.202.3      No          0/0             Down   Down
  sap:1/1/6                1/8000             Up     Up
-----
Interfaces: 2
=====
*A:Dut-C#
```

## local-dhcp-server

**Syntax**    **local-dhcp-server** *server-name*

**Context**    show>router>dhcp

**Description**    This command displays local DHCP server information.

**Parameters**    *server-name* — Specifies information about the local DHCP server.

### Sample Output

```
*A:ALA-48>show>router>dhcp>local-dhcp-server# declined-addresses pool test
=====
Declined addresses for server test Base
=====
Pool                Subnet                IP Address
PPPoE User Name/    Time                 MAC Address           Type
Option 82 Circuit ID
-----
No Matching Entries
=====
```

```
*A:ALA-48>show>router>dhcp>local-dhcp-server#
```

associations

<b>Syntax</b>	<b>associations</b>
<b>Context</b>	show>router>dhcp>local-dhcp-server show>router>dhcp
<b>Description</b>	This command displays the interfaces associated with this DHCP or DHCP6 server.

Sample Output

```
*A:SUB-Dut-A# show router dhcp local-dhcp-server dhcpS1 associations
=====
DHCP server s1  router 3
=====
Associations                               Admin
-----
tosim5                                     Up
=====
*A:SUB-Dut-A#
```

declined-addresses

<b>Syntax</b>	<b>declined-addresses</b> <i>ip-address[/mask]</i> [ <b>detail</b> ] <b>declined-addresses pool</b> <i>pool-name</i>
<b>Context</b>	show>router>dhcp>local-dhcp-server
<b>Description</b>	This command display information about declined addresses.
<b>Parameters</b>	<b>pool</b> <i>pool-name</i> — Specifies a DHCP pool name on the router.  <i>ip-address</i> — Specifies the IP address of the DNS server. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).  <b>detail</b> — Displays detailed information.

Sample Output

```
*A:ALA-48>show>router>dhcp>local-dhcp-server# declined-addresses pool test
=====
Declined addresses for server test Base
=====
Pool                               Subnet          IP Address
PPPoE User Name/                   Time            MAC Address    Type
Option 82 Circuit ID
-----
No Matching Entries
=====
*A:ALA-48>show>router>dhcp>local-dhcp-server#
```

## free-addresses

<b>Syntax</b>	<b>free-addresses</b> <i>ip-address[/mask]</i> <b>free-addresses summary</b> [ <b>subnet</b> <i>ip-address[/mask]</i> ] <b>free-addresses pool</b> <i>pool-name</i>
<b>Context</b>	show>router>dhcp>local-dhcp-server
<b>Description</b>	This command displays the free addresses in a subnet.
<b>Parameters</b>	<b>pool</b> <i>pool-name</i> — Specifies a DHCP pool name on the router. <b>subnet</b> <i>subnet</i> — Specifies a subnet of IP addresses that are served from the pool. <b>summary</b> — Displays summary output of the free addresses.

## Sample Output

```
*A:ALA-48>show>router>dhcp>local-dhcp-server# free-addresses pool test subnet
1.0.0.0/24
=====
Free addresses in subnet 1.0.0.0/24
=====
IP Address
-----
No. of free addresses: 0
=====
*A:ALA-48>show>router>dhcp>local-dhcp-server#
```

## interface-id-mapping

<b>Syntax</b>	<b>interface-id-mapping</b>
<b>Context</b>	show>router>dhcp6>local-dhcp-server
<b>Description</b>	This command displays the DHCP6 interface-id mappings.

## Sample Output

```
show router 600 dhcp6 local-dhcp-server "d6" interface-id-mapping
=====
Interface-ID Mappings for DHCPv6 server d6
=====
Mapped Prefix      : 2001:AAAA::/64
Relay Interface ID : 1/1/10
LDRA Interface ID  : (Not Specified)
Active Leases      : 2001:AAAA::1 (stable)
=====
1 prefix found
=====
```

## leases

<b>Syntax</b>	<b>leases</b> <b>leases</b> <i>ip-address[/mask]</i> <b>address-from-user-db</b> [ <b>detail</b> ] <b>leases</b> <i>ip-address[/mask]</i> <b>dhcp-host</b> <i>dhcp-host-name</i> [ <b>detail</b> ] <b>leases</b> <i>ip-address[/mask]</i> <b>ppp-host</b> <i>ppp-host-name</i> [ <b>detail</b> ] <b>leases</b> <i>ip-address[/mask]</i> [ <b>detail</b> ]
<b>Context</b>	show>router>dhcp>local-dhcp-server
<b>Description</b>	This command displays the DHCP leases.
<b>Parameters</b>	<i>ip-address</i> — Specifies the base IP address of the subnet. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets). <i>mask</i> — The subnet mask in dotted decimal notation. <b>Values</b> 0 — 32 <b>address-from-user-db</b> [ <b>detail</b> ] — Displays only leases that have ip-addresses from the local-user-db. <b>dhcp-host</b> <i>dhcp-host-name</i> [ <b>detail</b> ] — Shows all leases that match a certain DHCP host from the local-user-db. <b>ppp-host</b> <i>ppp-host-name</i> [ <b>detail</b> ] — Displays all leases that match a certain PPPoE host from the local-user-db. <b>detail</b> — Displays detailed information of all leases that fall into the indicated subnet. The command with no parameters will show all leases from the local-user-db.

## Sample Output

```
*A:ALA-48>show>router>dhcp>local-dhcp-server# leases ip-address 1.0.0.4
=====
Leases for DHCP server test router Base
=====
IP Address      Lease State      Mac Address      Remaining Clnt
  PPPoE user name/Opt82 Circuit Id      LifeTime  Type
-----
No leases found
*A:ALA-48>show>router>dhcp>local-dhcp-server#
```

## leases

<b>Syntax</b>	<b>leases</b> [ <i>ipv6-address/prefix-length</i> ] [ <i>type</i> ] [ <i>state</i> ] [ <b>detail</b> ]
<b>Context</b>	show>router>dhcp6>local-dhcp-server
<b>Description</b>	This command displays the DHCP6 leases.
<b>Parameters</b>	<i>ipv6-address</i> — Specifies the base IP address of the subnet. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).

*mask* — The subnet mask in dotted decimal notation.

**Values** 0 — 32

*type* — Displays the lease type.

**Values** pd, wan-host

*state* — Displays the state of the lease.

**Values** advertised, remove-pending, held

**detail** — Displays detailed information of all leases that fall into the indicated subnet.

The command with no parameters will show all leases from the local-user-db.

### Sample Output

```
show router 600 dhcp6 local-dhcp-server "d6" leases
=====
Leases for DHCPv6 server d6
=====
IP Address/Prefix          Lease State      Remaining      Fail
Link-local Address          LifeTime         Ctrl
-----
2001:AAAA::1/128
FE80::220:FCFF:FE1E:CD52    stable           23h58m52s     local
-----
1 leases found
=====
```

## pool-ext-stats

**Syntax** **pool-ext-stats** [*pool-name*]

**Context** show>router>dhcp>server

**Description** This command displays extended statistics per DHCPv4 pool in local DHCPv4 server.

The following statistics are included in output:

- The number of stable leases in the pool
- The number of provisioned address in the pool
- The number of used address in the pool
- The number of free address in the pool
- The percentage of used address
- The percentage of free address

For each statistic (except for Provisioned Addresses), there is current value and peak value, peak value is the highest value since pool creation or last reset via the **clear router *rt-id* dhcp local-dhcp-server *svr-name* pool-ext-stats** command.

**Parameters** *pool-name* — Specify the name of DHCPv4 local server pool.

### Sample Output

```
show router 500 dhcp local-dhcp-server "d4" pool-ext-stats "pool-1"
=====
Extended pool statistics for server "d4"
=====
-----
Current          Peak          TimeStamp
-----
Pool              pool-1
Local:
  Stable Leases    0              0              01/07/2013 19:07:11
  Provisioned Addresses 101            101            01/07/2013 19:07:11
  Used Addresses    0              0              01/07/2013 19:07:11
  Free Addresses    101            101            01/07/2013 19:07:11
  Used Pct          0              0              01/07/2013 19:07:11
  Free Pct          100            100            01/07/2013 19:07:11
Last Reset Time    01/07/2013 19:07:11
-----
Number of entries   1
=====
```

## pool-ext-stats

**Syntax** `pool-ext-stats [pool-name]`

**Context** `show>router>dhcp6>server`

**Description** This command displays extended statistics per DHCPv6 pool in local DHCPv6 server.

The following statistics are included in output:

- The number of stable leases in the pool
- The number of provisioned /64 address block in the pool
- The number of used /64 address block in the pool
- The number of free /64 address block in the pool
- The percentage of used address (with /64 address block)
- The percentage of free address (with /64 address block)

For each statistic (except for Provisioned Addresses), there is current value and peak value, peak value is the highest value since pool creation or last reset via command “clear router <rt-id> dhcp6 local-dhcp-server <svr-name> pool-ext-stats”.

**Parameters** *pool-name* — Specify the name of DHCPv6 local server pool.

### Sample Output

```
show router 500 dhcp6 local-dhcp-server "d6" pool-ext-stats "pool-v6"
=====
Extended pool statistics for server "d6"
=====
-----
Current          Peak          TimeStamp
-----
Pool              pool-v6
Local:
  Stable Leases    0              0              01/07/2013 19:54:52
  Provisioned Blks 4              4              01/07/2013 19:54:52
  Used Blks        0              0              01/07/2013 19:54:52
```

```

Free Blks          4          4          01/07/2013 19:54:52
Used Pct           0          0          01/07/2013 19:54:52
Free Pct           100        100        01/07/2013 19:54:52
Last Reset Time    01/07/2013 19:54:52
-----
Number of entries   1
=====

```

## prefix-ext-stats

- Syntax** **prefix-ext-stats** *ipv6-address/prefix-length*  
**prefix-ext-stats pool** *pool-name*
- Context** show>router>dhcp6>server
- Description** This command displays extended statistics per DHCPv6 prefix in local DHCPv6 server.  
The following statistics are included in output:
- The number of stable leases in the prefix
  - The number of provisioned /64 address block in the prefix
  - The number of used /64 address block in the prefix
  - The number of free /64 address block in the prefix
  - The percentage of used address (with /64 address block)
  - The percentage of free address (with /64 address block)
- For each statistic (except for “Provisioned Addresses”), there is current value and peak value, peak value is the highest value since prefix creation or last reset via command “clear router <rt-id> dhcp6 local-dhcp-server <svr-name> prefix-ext-stats”.
- When parameter “pool” is used, the statistics of each prefix in the pool will be displayed.
- Parameters** *ipv6-address/prefix-length* — Specifies the IPv6 prefix  
*pool-name* — The name of DHCPv6 local server pool

### Sample Output

```

show router 500 dhcp6 local-dhcp-server "d6" prefix-ext-stats 2001:ABCD::/62
=====
Extended statistics for prefix 2001:ABCD::/62
=====

```

	Current	Peak	TimeStamp
-----			
Local:			
Failover Oper State	Active		
Stable Leases	0	0	01/07/2013 19:54:52
Provisioned Blks	4		
Used Blks	0	0	01/07/2013 19:54:52
Free Blks	4	4	01/07/2013 19:54:52
Used Pct	0	0	01/07/2013 19:54:52
Free Pct	100	100	01/07/2013 19:54:52
Last Reset Time			01/07/2013 19:54:52
-----			
Number of entries	1		

=====

pool-threshold-stats

- Syntax

pool-threshold-stats [pool-name] detail [format {exact|scientific}]  
pool-threshold-stats [pool-name]
- Context

show>router>dhcp6>server
- Description

This commands displays pool level threshold stats of local DHCPv6 server. A minimum-free threshold needs to be configured before system collects threshold stats for the prefix.

The stats for each threshold are calculated based on the configured minimum-free prefix length.

For example, a /59 prefix is provision in the local DHCPv6 server, and the server allocated two PD leases, one /62 and one /63. And there is a /63 minimum threshold configured. So the threshold stats are calculated based on /63 as the base unit(block). So the value of “current used block” would be 3 because there is one /62 lease and one /63 lease, so it equals to total three /63.
- Parameters

*pool-name* — Specifies the name of the pool in local DHCPv6 server.

**detail** — Displays detailed output.

**format** — Specifies the format in the display to be either **exact** or **scientific**.

Sample Output

```
show router 500 dhcp6 local-dhcp-server "d6" pool-threshold-stats "1"
=====
Server "d6"
=====
  Operational state      : inService
-----
Pool                    : 1
-----
  Stable leases         : 2
  Advertised leases     : 0
-----
  Threshold   Used   Peak   Too low   Depleted   Peak timestamp
-----
  /62         25%    25%    N        N          01/21/2015 21:52:12
  /63         19%    19%    N        N          01/21/2015 21:52:12
```

- The command shown above displays an overview of pool level thresholds in the specified pool:
- The **Peak** field indicates the peak value of used
  - The **Too low** field indicate if the configured minimum-free threshold is exceed
  - The **Depleted** field indicate if there is no available prefix with the length in the provisioned prefix
  - The **Peak timestamp** field indicates the time of peak used value

```
show router 500 dhcp6 local-dhcp-server "d6" pool-threshold-stats "1" detail
=====
Server "d6"
```



```

=====
Operational state      : inService
-----
Pool                  : 1
-----
Stable leases         : 2
Advertised leases     : 0
-----
Threshold             : /62
-----
Current Provisioned Blks : 8.000000x10^0
Current Used Blks       : 2.000000x10^0
Current Free Blks       : 6.000000x10^0
Current Used Percent    : 25%
Current Used Peak Blks  : 2.000000x10^0
Current Used Peak Percent : 25%
Current Used Peak Time  : 01/21/2015 21:52:12
Current Free Percent    : 75%
Current Free Too Low    : N
Current Free Depleted   : N
Local Provisioned Blks  : 8.000000x10^0
Local Used Blks         : 2.000000x10^0
Local Free Blks         : 6.000000x10^0
Local Used Peak Blks    : 2.000000x10^0
Local Used Peak Percent : 25%
Local Used Peak Time    : 01/21/2015 21:52:12
Remote Provisioned Blks : 0.000000x10^0
Remote Used Blks        : 0.000000x10^0
Remote Free Blks        : 0.000000x10^0
Remote Used Peak Blks   : 0.000000x10^0
Remote Used Peak Percent : 0%
Remote Used Peak Time   : 01/21/2015 21:47:39
Peak Reset Time         : 01/21/2015 21:47:39
Valid Data              : Y
-----
Threshold             : /63
-----
Current Provisioned Blks : 1.600000x10^1
Current Used Blks       : 3.000000x10^0
Current Free Blks       : 1.300000x10^1
Current Used Percent    : 19%
Current Used Peak Blks  : 3.000000x10^0
Current Used Peak Percent : 19%
Current Used Peak Time  : 01/21/2015 21:52:12
Current Free Percent    : 81%
Current Free Too Low    : N
Current Free Depleted   : N
Local Provisioned Blks  : 1.600000x10^1
Local Used Blks         : 3.000000x10^0
Local Free Blks         : 1.300000x10^1
Local Used Peak Blks    : 3.000000x10^0
Local Used Peak Percent : 19%
Local Used Peak Time    : 01/21/2015 21:52:12
Remote Provisioned Blks : 0.000000x10^0
Remote Used Blks        : 0.000000x10^0
Remote Free Blks        : 0.000000x10^0
Remote Used Peak Blks   : 0.000000x10^0
Remote Used Peak Percent : 0%
Remote Used Peak Time   : 01/21/2015 21:47:39
Peak Reset Time         : 01/21/2015 21:47:39
Valid Data              : Y

```

The above command displays detailed statistics of all pool level thresholds in the specified pool:

- **Blks** in the output means the minimum free prefix length.
- **Valid Data** output indicates whether the data you see is valid or not. The data is invalid when a background stats update is scheduled or busy.

## prefix-threshold-stats

**Syntax** `prefix-threshold-stats pool pool-name detail [format {exact|scientific}]`  
`prefix-threshold-stats pool pool-name`  
`prefix-threshold-stats ipv6-address/prefix-length detail [format {exact|scientific}]`  
`prefix-threshold-stats ipv6-address/prefix-length`

**Context** show>router>dhcp6>server

<b>Description</b>	This command displays prefix level threshold stats of local DHCPv6 server prefix. A minimum-free threshold needs to be configured before system collects threshold stats for the prefix.
--------------------	--

The stats for each threshold are calculated based on the configured minimum-free prefix length.

For example, a /59 prefix is provision in the local DHCPv6 server, and the server allocated two PD leases, one /62 and one /63. And there is a /63 minimum threshold configured. So the threshold stats are calculated based on /63 as the base unit(block). So the value of “current used block” would be 3 because there is one /62 lease and one /63 lease, so it equals to total three /63.

**Parameters**    **pool** *pool-name* — Specifies the name of the pool in local DHCPv6 server up to 32 characters in length.

**detail** — Displays detailed output statistics.

**format** — Specifies that the number format in the display will be either **exact** or **scientific**

*ipv6-address/prefix-length* — Specifies the IPv6 prefix with prefix length

<b>Values</b>	ipv6-address 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 prefix-length [1..128]
---------------	--

### Sample Output

```

show router 500 dhcp6 local-dhcp-server "d6" leases
=====
Leases for DHCPv6 server d6
=====
IP Address/Prefix                               Lease State      Remaining      Fail
  Link-local Address                             LifeTime         Ctrl
-----
8888:0:0:ffe0::/62
  fe80::3:ffff:fe00:111                         stable          18h19m2s      local
8888:0:0:ffe4::/63
  fe80::3:ffff:fe00:211                         stable          19h49m37s     local
-----
2 leases found
=====
show router 500 dhcp6 local-dhcp-server "d6" prefix-threshold-stats pool "1"

```

```

=====
Server "d6"
=====
Operational state      : inService
-----
Pool                  : 1
-----
Stable leases         : 2
Advertised leases     : 0
-----
Prefix                : 8888:0:0:ffe0::/59
-----
Stable leases         : 2
Advertised leases     : 0
Draining              : N
-----
Threshold   Used   Peak   Too low   Depleted   Peak timestamp
-----
/62         25%    25%    Y         N         01/20/2015 23:51:36
/63         19%    19%    N         N         01/21/2015 05:00:53
=====

```

The command shown above displays an overview of prefix level thresholds in the specified pool:

- The **Peak** field indicates the peak value of used.
- The **Too low** field indicate if the configured minimum-free threshold is exceed.
- The **Depleted** field indicate if there is no available prefix with the length in the provisioned prefix.
- The **Peak** timestamp field indicates the time of peak used value.

```

show router 500 dhcp6 local-dhcp-server "d6" prefix-threshold-stats pool "1" detail
=====
Server "d6"
=====
Operational state      : inService
-----
Pool                  : 1
-----
Stable leases         : 2
Advertised leases     : 0
-----
Prefix                : 8888:0:0:ffe0::/59
-----
Stable leases         : 2
Advertised leases     : 0
Draining              : N
-----
Threshold             : /62
-----
Current Provisioned Blks : 8.000000x10^0
Current Used Blks       : 2.000000x10^0
Current Free Blks       : 6.000000x10^0
Current Used Percent    : 25%
Current Used Peak Blks  : 2.000000x10^0
Current Used Peak Percent : 25%
Current Used Peak Time  : 01/21/2015 21:59:02
Current Free Percent    : 75%
Current Free Too Low    : N
Current Free Depleted   : N
Local Provisioned Blks  : 8.000000x10^0
Local Used Blks         : 2.000000x10^0

```

## Show Commands

```
Local Free Blks      : 6.000000x10^0
Local Used Peak Blks : 2.000000x10^0
Local Used Peak Percent : 25%
Local Used Peak Time  : 01/21/2015 21:59:02
Remote Provisioned Blks : 0.000000x10^0
Remote Used Blks      : 0.000000x10^0
Remote Free Blks      : 0.000000x10^0
Remote Used Peak Blks : 0.000000x10^0
Remote Used Peak Percent : 0%
Remote Used Peak Time  : 01/21/2015 21:59:02
Peak Reset Time       : 01/21/2015 21:59:02
Valid Data            : Y
```

```
-----
Threshold            : /63
-----
```

```
Current Provisioned Blks : 1.600000x10^1
Current Used Blks        : 3.000000x10^0
Current Free Blks        : 1.300000x10^1
Current Used Percent      : 19%
Current Used Peak Blks   : 3.000000x10^0
Current Used Peak Percent : 19%
Current Used Peak Time    : 01/21/2015 21:59:13
Current Free Percent      : 81%
Current Free Too Low     : N
Current Free Depleted    : N
Local Provisioned Blks   : 1.600000x10^1
Local Used Blks          : 3.000000x10^0
Local Free Blks          : 1.300000x10^1
Local Used Peak Blks     : 3.000000x10^0
Local Used Peak Percent  : 19%
Local Used Peak Time     : 01/21/2015 21:59:13
Remote Provisioned Blks  : 0.000000x10^0
Remote Used Blks         : 0.000000x10^0
Remote Free Blks         : 0.000000x10^0
Remote Used Peak Blks    : 0.000000x10^0
Remote Used Peak Percent : 0%
Remote Used Peak Time    : 01/21/2015 21:59:13
Peak Reset Time         : 01/21/2015 21:59:13
Valid Data              : Y
```

The command shown above displays detailed statistics of all prefix level thresholds in the specified pool:

- **Blks** means the minimum free prefix length.
- **Valid Data** output indicates whether the data is or is not valid. The data is invalid when a background stats update is scheduled or busy.

```
show router 500 dhcp6 local-dhcp-server "d6" prefix-threshold-stats 8888:0:0:ffe0::/
59
=====
Server "d6"
=====
Operational state      : inService
-----
Pool                   : 1
-----
Stable leases          : 2
Advertised leases      : 0
-----
Prefix                 : 8888:0:0:ffe0::/59
-----
```

```

Stable leases           : 2
Advertised leases       : 0
Draining                : N
-----
Threshold   Used   Peak   Too low   Depleted   Peak timestamp
-----
/62         25%    25%    N         N         01/21/2015 21:59:02
/63         19%    19%    N         N         01/21/2015 21:59:13

```

The command shown above displays an overview of prefix level thresholds in the specified provision prefix.

```
show router 500 dhcp6 local-dhcp-server "d6" prefix-threshold-stats 8888:0:0:ffe0::/59 detail
```

```

=====
Server "d6"
=====
Operational state      : inService
-----
Pool                  : 1
-----
Stable leases          : 2
Advertised leases      : 0
-----
Prefix                : 8888:0:0:ffe0::/59
-----
Stable leases          : 2
Advertised leases      : 0
Draining               : N
-----
Threshold              : /62
-----
Current Provisioned Blks : 8.000000x10^0
Current Used Blks       : 2.000000x10^0
Current Free Blks       : 6.000000x10^0
Current Used Percent    : 25%
Current Used Peak Blks  : 2.000000x10^0
Current Used Peak Percent : 25%
Current Used Peak Time  : 01/21/2015 21:59:02
Current Free Percent    : 75%
Current Free Too Low    : N
Current Free Depleted   : N
Local Provisioned Blks  : 8.000000x10^0
Local Used Blks         : 2.000000x10^0
Local Free Blks         : 6.000000x10^0
Local Used Peak Blks    : 2.000000x10^0
Local Used Peak Percent : 25%
Local Used Peak Time    : 01/21/2015 21:59:02
Remote Provisioned Blks : 0.000000x10^0
Remote Used Blks        : 0.000000x10^0
Remote Free Blks        : 0.000000x10^0
Remote Used Peak Blks   : 0.000000x10^0
Remote Used Peak Percent : 0%
Remote Used Peak Time   : 01/21/2015 21:59:02
Peak Reset Time         : 01/21/2015 21:59:02
Valid Data              : Y
-----
Threshold              : /63
-----
Current Provisioned Blks : 1.600000x10^1
Current Used Blks        : 3.000000x10^0
Current Free Blks        : 1.300000x10^1

```

```

Current Used Percent      : 19%
Current Used Peak Blks   : 3.000000x10^0
Current Used Peak Percent : 19%
Current Used Peak Time    : 01/21/2015 21:59:13
Current Free Percent     : 81%
Current Free Too Low     : N
Current Free Depleted    : N
Local Provisioned Blks   : 1.600000x10^1
Local Used Blks          : 3.000000x10^0
Local Free Blks          : 1.300000x10^1
Local Used Peak Blks     : 3.000000x10^0
Local Used Peak Percent  : 19%
Local Used Peak Time     : 01/21/2015 21:59:13
Remote Provisioned Blks  : 0.000000x10^0
Remote Used Blks         : 0.000000x10^0
Remote Free Blks        : 0.000000x10^0
Remote Used Peak Blks    : 0.000000x10^0
Remote Used Peak Percent : 0%
Remote Used Peak Time    : 01/21/2015 21:59:13
Peak Reset Time         : 01/21/2015 21:59:13
Valid Data              : Y

```

The command displayed above displays detailed statistics of prefix level thresholds in the specified provision prefix.

## subnet-ext-stats

<b>Syntax</b>	<b>subnet-ext-stats</b> <i>ip-address[/mask]</i> <b>subnet-ext-stats pool</b> <i>pool-name</i>
<b>Context</b>	show>router>dhcp>server
<b>Description</b>	<p>This command displays extended statistics per DHCPv4 subnet in local DHCPv4 server.</p> <p>The following statistics are included in output:</p> <ul style="list-style-type: none"> <li>• The number of stable leases in the subnet</li> <li>• The number of provisioned address in the subnet</li> <li>• The number of used address in the subnet</li> <li>• The number of free address in the subnet</li> <li>• The percentage of used address</li> <li>• The percentage of free address</li> </ul> <p>For each statistic (except for Provisioned Addresses), there is current value and peak value, peak value is the highest value since subnet creation or last reset via the <b>clear router <i>rt-id</i> dhcp local-dhcp-server <i>svr-name</i> subnet-ext-stats</b> command.</p> <p>When parameter pool is used, the statistics of each subnet in the pool will be displayed.</p>
<b>Parameters</b>	<i>ip-address[/mask]</i> — Specifies the subnet <i>pool-name</i> — The name of local DHCPv4 server pool
<b>Sample Output</b>	

```

show router 500 dhcp local-dhcp-server "d4" subnet-ext-stats 220.10.10.0/24
=====
Extended statistics for subnet 220.10.10.0/24
=====

```

	Current	Peak	TimeStamp
-----			
Local:			
Stable Leases	1	1	01/07/2013 19:38:36
Provisioned Addresses	101		
Used Addresses	1	1	01/07/2013 19:38:36
Free Addresses	100	100	01/07/2013 19:38:36
Used Pct	1	1	01/07/2013 19:38:36
Free Pct	99	99	01/07/2013 19:38:36
Last Reset Time			01/07/2013 19:07:11
-----			
Number of entries	1		
=====			

## server-stats

<b>Syntax</b>	<b>server-stats</b>
<b>Context</b>	show>router>dhcp>server
<b>Description</b>	This command displays server statistics.

### Sample Output

```

*A:SUB-Dut-A# show router dhcp local-dhcp-server dhcpS1 server-stats
=====
Statistics for DHCP Server dhcpS1 router Base
=====
Rx Discover Packets           : 0
Rx Request Packets           : 0
Rx Release Packets           : 0
Rx Decline Packets           : 0
Rx Inform Packets            : 0

Tx Offer Packets              : 0
Tx Ack Packets                : 0
Tx Nak Packets                : 0
Tx Forcerenew Packets        : 0

Client Ignored Offers         : 0
Leases Timed Out              : 0

Dropped Bad Packet            : 0
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       : 0
Dropped Not Serving Pool      : 0
Dropped Invalid User          : 0
Dropped Overload              : 0
Dropped Persistence Overload  : 0
Dropped Generic Error         : 0
Dropped Destined To Other     : 0

```

Show Commands

```
Dropped Address Unavailable : 0
Dropped Max Leases Reached : 0
Dropped Server Shutdown : 0
Dropped No Subnet For Fixed IP: 0

=====
*A:SUB-Dut-A#
```

subnet-stats

- Syntax** subnet-stats ip-address[/mask]  
subnet-stats pool pool-name
- Context** show>router>dhcp>server
- Description** This command displays subnet statistics.

Sample Output

```
*A:SUB-Dut-A# show router dhcp local-dhcp-server dhcpS2 subnet-stats pool POOL2
=====
Statistics for pool POOL2
=====
Subnet                Free      Offered    Stable
                      FRPending  RemPending Declined
-----
2.0.0.0/8             16384      0          0
                      0          0          0
-----
No. of entries: 1
=====
*A:SUB-Dut-A#
```

summary

- Syntax** summary
- Context** show>router>dhcp>server
- Description** This command displays DHCP summary information.

Sample Output

```
*A:SUB-Dut-A# show router dhcp local-dhcp-server dhcpS2 summary
=====
DHCP server dhcpS2  router Base
=====
dhcpS2-POOL2
Admin State          : inService
Persistency State    : ok
User Data Base       : N/A
Use gateway IP address : disabled
Send force-renewals   : disabled
```



```
-----
Pool name : POOL2
-----
```

Subnet	Free	Stable	Declined	Offered	Remove-pending
2.0.0.0/8	16384	0	0	0	0
Totals for pool	16384	0	0	0	0
Totals for server	16384	0	0	0	0

```
-----
Associations                               Admin
-----
```

```
No associations found
=====
```

```
*A:SUB-Dut-A#
```

## servers

**Syntax** **servers**

**Context** show>router>dhcp

**Description** This command lists the local DHCP servers.

### Sample Output

```
*A:ALA-49>show>router>dhcp# servers
=====
Overview of DHCP Servers
=====
Active Leases:      0
Maximum Leases:     159744

Router              Server                               Admin State
-----
Router: Base        base_router_dhcp_server          outOfService
Service: 3          s1                                   inService
=====
*A:ALA-49>show>router>dhcp#
```

## servers

**Syntax** **servers**

**Context** show>router>dhcp>local-dhcp-server>statistics

**Description** This command displays server statistics.

### Sample Output

```
*A:ALA-48>show>router>dhcp>local-dhcp-server>statistics# servers
=====
Statistics for DHCP Server test router Base
=====
```

## Show Commands

```
Rx Discover Packets      : 0
Rx Request Packets      : 0
Rx Release Packets      : 0
Rx Decline Packets      : 0
Rx Inform Packets       : 0

Tx Offer Packets        : 0
Tx Ack Packets          : 0
Tx Nack Packets         : 0
Tx Forcerenew Packets   : 0

Client ignored offers   : 0

Dropped Bad Packet      : 0
Dropped Invalid Type    : 0
Dropped Unknown Host    : 0
Dropped User Not Allowed: 0
Dropped Lease Not Ready : 0
Dropped Lease Not Found : 0
Dropped Not Serving Pool: 0
Dropped Invalid User    : 0
Dropped Generic Error   : 0
=====
*A:ALA-48>show>router>dhcp>local-dhcp-server>statistics#
```

## subnet

<b>Syntax</b>	<b>subnet pool</b> <i>pool-name</i> [ <b>subnet</b> <i>subnet</i> ]
<b>Context</b>	show>router>dhcp>local-dhcp-server>statistics
<b>Description</b>	This command displays subnet statistics.
<b>Parameters</b>	<b>pool</b> <i>pool-name</i> — Specifies the pool name on the router. <b>subnet</b> <i>subnet</i> — Specifies a subnet of IP addresses that are served from the pool.

### Sample Output

```
*A:ALA-48>show>router>dhcp>local-dhcp-server>statistics# subnet pool test
=====
Statistics for pool test
=====
Subnet                Free      Offered      Stable
                      FRPending  RemPending  Declined
-----
1.0.0.0/24            0          0            0
                      0          0            0
-----
No. of entries: 1
=====
*A:ALA-48>show>router>dhcp>local-dhcp-server>statistics#
```

## lease-state

<b>Syntax</b>	<b>lease-state</b> <b>[[sap sap-id]   [sdp dp-id:vc-id ]   [interface interface-name]   [ip-address ip-address[/mask]&gt;]   [mac ieee-address]] [detail]</b>
<b>Context</b>	show>service>id>dhcp
<b>Description</b>	This command displays DHCP lease state related information.
<b>Parameters</b>	<p><b>sap-id</b> — Specifies the physical port identifier portion of the SAP definition. See <a href="#">Common Service Commands on page 2062</a> for <i>sap-id</i> command syntax.</p> <p><b>sdp-id</b> — The SDP ID to be shown.</p> <p><b>Values</b> 1 — 17407</p> <p><b>vc-id</b> — The virtual circuit ID on the SDP ID to be shown.</p> <p><b>Values</b> 1 — 4294967295</p>

## servers

<b>Syntax</b>	<b>servers</b>
<b>Context</b>	show>router>dhcp
<b>Description</b>	This command lists the local DHCP servers.

## Sample Output

```
*A:SUB-Dut-A# show router dhcp servers
=====
Overview of DHCP Servers
=====
Active Leases:      0
Maximum Leases:    159744

Router              Server              Admin State
-----
Router: Base        dhcpS1              inService
Router: Base        dhcpS10            inService
Router: Base        dhcpS100           inService
Router: Base        dhcpS101           inService
Router: Base        dhcpS102           inService
Router: Base        dhcpS103           inService
Router: Base        dhcpS104           inService
Router: Base        dhcpS105           inService
Router: Base        dhcpS106           inService
Router: Base        dhcpS107           inService
Router: Base        dhcpS108           inService
Router: Base        dhcpS109           inService
Router: Base        dhcpS11            inService
Router: Base        dhcpS110           inService
Router: Base        dhcpS111           inService
Router: Base        dhcpS112           inService
Router: Base        dhcpS113           inService
Router: Base        dhcpS114           inService
Router: Base        dhcpS115           inService
```

## Show Commands

```
Router: Base      dhcpS116      inService
Router: Base      dhcpS117      inService
Router: Base      dhcpS118      inService
Router: Base      dhcpS119      inService
...
Service: 1022     dhcpS1022     inService
Service: 1023     dhcpS1023     inService
Service: 1024     dhcpS1024     inService
=====
*A:SUB-Dut-A#

*A:SUB-Dut-A#
=====
Overview of DHCP Servers
=====
Active Leases: 0
Maximum Leases: 159744

Router Server Admin State
-----
Router: Base base_router_dhcp_server outOfService
Service: 3 s1 inService
=====
```

## statistics

<b>Syntax</b>	<b>statistics</b> [ <b>interface</b> <i>ip-int-name</i> ]
<b>Context</b>	show>router>dhcp6 show>service>id>dhcp6
<b>Description</b>	This command displays statistics for DHCP relay and DHCP snooping.

### Sample Output

```
*A:Dut-C# show service id 202 dhcp6 statistics
=====
DHCP Statistics, service 202, all interfaces
=====
Packets received           : 1
Packets transmitted       : 1
Packets dropped            : 0
=====
*A:Dut-C#
```

## summary

<b>Syntax</b>	<b>summary</b>
<b>Context</b>	show>router>dhcp6 show>service>id>dhcp6

**Description** This command displays the status of the DHCP6 relay and DHCP snooping functions on each interface.

**Output** **Show DHC6P Summary Output** — The following table describes the output fields for DHCP6 summary.

Label	Description
Interface Name	Name of the router interface.
ARP Populate	Indicates whether ARP populate is enabled.
Used/Provided	Indicates the number of used and provided DHCP leases.
Info Option	Indicates whether Option 82 processing is enabled on the interface.
Admin State	Indicates the administrative state.

### Sample Output

```
*A:Dut-C# show router dhcp6 summary
=====
DHCP6 Summary (Router: Base)
=====
Interface Name      Nbr      Used/Max Relay   Admin  Oper  Relay
  SapId            Resol.   Used/Max Server  Admin  Oper  Server
-----
ip-1.1.1.10         No        0/0              Down   Down
  sap:1/1/5          0/8000
ip-11.3.202.3       No        0/0              Down   Down
  sap:1/1/6          1/8000
  Up               Up
-----
Interfaces: 2
=====
*A:Dut-C#
```

## remap-lease-state

**Syntax** **remap-lease-state old-mac** *ieee-address* **mac** *ieee-address*  
**remap-lease-state sap** *sap-id* [**mac** *ieee-address*]

**Context** tools>perform>subscr-mgmt

**Description** This command allows the remapping of all existing hosts if network card on CMTS/WAC side is changed is required.

When this command is executed, the following restrictions apply

- When **sap** is taken, all leases associated with the SAP are re-written.
  - For a SAP with a configured MAC in lease-populate command, this MAC will be taken.
  - For a SAP without a configured MAC the MAC from tools command will be taken.
  - For a SAP without a configured MAC and no MAC in tools command no action will be perform.
- When using the **old-mac** option, providing a new MAC *ieee-address* is mandatory.

This command is applicable only when dealing with DHCP lease states which were instantiated using l2header mode of DHCP operation.

### Parameters

**old-mac** *ieee-address*

**old-mac** *ieee-address* — specifies the old MAC address to remap.

**mac** *ieee-address* — Specifies that the provisioned MAC address will be used in the anti-spoofing entries for this SAP when l2-header is enabled. The parameter may be changed mid-session. Existing sessions will not be re-programmed unless a **tools perform** command is issued for the lease.

**sap** *sap-id* — Specifies the physical port identifier portion of the SAP definition. See [Common Service Commands on page 2062](#) for *sap-id* command syntax.

When configured, the SAP parameter will remap all MAC addresses of DHCP lease states on the specified SAP. When no optional MAC parameter is specified, the **sap** *sap-id* command remaps all MAC addresses of lease states towards the MAC address specified in the l2-header configuration.

## Clear Commands

### dhcp

<b>Syntax</b>	<b>dhcp</b>
<b>Context</b>	clear>router
<b>Description</b>	This command enables the context to clear and reset DHCP entities.

### dhcp6

<b>Syntax</b>	<b>dhcp6</b>
<b>Context</b>	clear>router
<b>Description</b>	This command enables the context to clear and reset DHCP6 entities.

### lease-state

<b>Syntax</b>	<b>lease-state [no-dhcp-release]</b> <b>lease-state [port <i>port-id</i>] [inter-dest-id <i>intermediate-destination-id</i>] [no-dhcp-release]</b> <b>lease-state [port <i>port-id</i>] no-inter-dest-id [no-dhcp-release]</b> <b>lease-state ip-address <i>ip-address</i> [no-dhcp-release]</b> <b>lease-state mac <i>ieee-address</i> no-dhcp-release</b> <b>lease-state sap <i>sap-id</i> [no-dhcp-release]</b> <b>lease-state sdp <i>sdp-id:vc-id</i> [no-dhcp-release]</b>
<b>Context</b>	clear>service>id>dhcp
<b>Description</b>	This command clears DHCP lease state information.
<b>Parameters</b>	<p><b>no-dhcp-release</b> — Clears the state without sending the DHCP release message.</p> <p><b>ip-address <i>ip-address</i></b> — Clears the DHCP IP address lease state information. The <i>ip-address</i> portion of the <b>address</b> command specifies the IP host address that will be used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).</p> <p><b>mac <i>ieee-address</i></b> — Clears DHCP MAC address lease state information. The 48-bit MAC address for the static ARP in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee, and ff are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.</p> <p><b>sap <i>sap-id</i></b> — clears DHCP SAP lease state information. See <a href="#">Common Service Commands on page 2062</a> for <i>sap-id</i> command syntax.</p>

## Clear Commands

*sdp-id* — Clears DHCP SDP lease state information.

**Values** 1 — 17407

*port-id* — Clears DHCP port lease state information. [Common Service Commands on page 2062](#)

*intermediate-destination-id* — Specifies the intermediate destination identifier which is encoded in the identification strings.

*vc-id* — Clears virtual circuit ID information on the specified SDP.

**Values** 1 — 4294967295

## local-dhcp-server

<b>Syntax</b>	<b>local-dhcp-server</b> <i>server-name</i>
<b>Context</b>	clear>router>dhcp
<b>Description</b>	This command clears DHCP server data.
<b>Parameters</b>	<i>server-name</i> — Clears data for the specified local DHCP server.

## declined-addresses

<b>Syntax</b>	<b>declined-addresses</b> <i>ip-address[/mask]</i> <b>declined-addresses pool</b> <i>pool-name</i>
<b>Context</b>	clear>router>dhcp>local-dhcp-server
<b>Description</b>	This command clears declined DHCP addresses.
<b>Parameters</b>	<i>pool-name</i> — Specifies the declined pool name. <i>ip-address[/mask]</i> — Specifies the declined IP address and mask.

## leases

<b>Syntax</b>	<b>leases</b> <i>ip-address[/mask]</i> [ <b>offered</b> ]
<b>Context</b>	clear>router>dhcp>local-dhcp-server
<b>Description</b>	This command clears DHCP leases.
<b>Parameters</b>	<i>ip-address[/mask]</i> — Clears the specified IP address and mask. <b>offered</b> — Clears leases in offered state only.

## pool-ext-stats

<b>Syntax</b>	<b>pool-ext-stats</b> [ <i>pool-name</i> ]
---------------	--



<b>Context</b>	clear>router>dhcp>local-dhcp-server
<b>Description</b>	This command clears extended pool statistics.
<b>Parameters</b>	<i>pool-name</i> — Specifies the pool name.

## server-stats

<b>Syntax</b>	<b>server-stats</b>
<b>Context</b>	clear>router>dhcp>local-dhcp-server
<b>Description</b>	This command clears all server statistics.

## subnet-ext-stats

<b>Syntax</b>	<b>subnet-ext-stats</b> <i>ip-address[/mask]</i> <b>subnet-ext-stats</b> <b>pool</b> <i>pool-name</i>
<b>Context</b>	clear>router>dhcp>local-dhcp-server
<b>Description</b>	This command clears extended subnet statistics.

## lease-state

<b>Syntax</b>	<b>lease-state</b> [ <b>ip-address</b> <i>ipv6-address/prefix-length</i> ] [ <b>mac</b> <i>ieee-address</i> ]
<b>Context</b>	clear>service>id>dhcp6
<b>Description</b>	This command clears DHCP6 lease state information.
<b>Parameters</b>	<p><b>ip-address</b> <i>ipv6-address</i> — Clears the DHCP6 IP address lease state information. The <i>ipv6-address</i> portion of the <b>address</b> command specifies the IP host address that will be used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IPv6 addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).</p> <p><b>mac</b> <i>ieee-address</i> — Clears DHCP6 MAC address lease state information. The 48-bit MAC address for the static ARP in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee, and ff are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.</p>

## statistics

<b>Syntax</b>	<b>statistics</b> [ <b>sap</b> <i>sap-id</i>   <b>sdp</b> [ <i>sdp-id</i> [: <i>vc-id</i> ]]   <b>interface</b> <i>ip-int-name</i>   <i>ip-address</i> ]
<b>Context</b>	clear>router>dhcp
<b>Description</b>	This command clears DHCP statistics.

## Clear Commands

<b>Parameters</b>	<b>sap</b> <i>sap-id</i> — clears DHCP statistics. See <a href="#">Common Service Commands on page 2062</a> for <i>sap-id</i> command syntax.  <i>sdp-id</i> — Clears DHCP SDP statistics.  <b>Values</b> 1 — 17407  <i>vc-id</i> — Clears DHCP the SDP VC ID statistics.  <b>Values</b> 1 — 4294967295  <i>ip-int-name</i> — Clears DHCP statistics for the specified interface name.  <i>ip-address</i> — Clears DHCP statistics for the specified IP address.
-------------------	--

## local-dhcp-server

<b>Syntax</b>	<b>local-dhcp-server</b> <i>server-name</i>
<b>Context</b>	clear>router>dhcp6
<b>Description</b>	This command enables the context to clear local DHCP server data.

## leases

<b>Syntax</b>	<b>leases</b> [ <i>ipv6-address/prefix-length</i> ] [ <i>type</i> ] [ <i>state</i> ] <b>leases all</b> [ <i>type</i> ] [ <i>state</i> ]
<b>Context</b>	clear>router>dhcp6>server
<b>Description</b>	This command removes the specified leases in the specified local DHCPv6 server.
<b>Parameters</b>	<i>ipv6-address/prefix-length</i> — The prefix of the leases to be removed. <i>type</i> — The type of the lease to be remove. <b>Values</b> pd, wan-host <i>state</i> — The state of the lease to be removed. <b>Values</b> advertised, remove-pending, held <b>all</b> — Remove all leases of specified type and(or) state.

## pool-ext-stats

<b>Syntax</b>	<b>pool-ext-stats</b> [ <i>pool-name</i> ]
<b>Context</b>	clear>router>dhcp6>server
<b>Description</b>	This command reset the begin time of peak values in output of the <b>show router <i>rt-id</i> dhcp6 local-dhcp-server <i>svr-name</i> pool-ext-stats</b> command.
<b>Parameters</b>	<i>pool-name</i> — The name of the local DHCPv6 server pool.  — <b>pool-threshold-stats</b> [ <i>pool-name</i> ]

## pool-threshold-stats

<b>Syntax</b>	<b>pool-threshold-stats</b> [ <i>pool-name</i> ]
<b>Context</b>	clear>router>dhcp6>server
<b>Description</b>	This commands resets the peak stats in the pool level threshold stats in the specified pool. If the pool name is not specified, then the peak stats in all pools in the server will be reset. <i>pool-name</i> — The name of the local DHCPv6 server pool

## prefix-ext-stats

<b>Syntax</b>	<b>prefix-ext-stats</b> <i>ipv6-address/prefix-length</i> <b>prefix-ext-stats</b> <b>pool</b> <i>pool-name</i>
<b>Context</b>	clear>router>dhcp6>server
<b>Description</b>	This command reset the begin time of peak values in output of the <b>show router rt-id dhcp6 local-dhcp-server svr-name prefix-ext-stats</b> command/
<b>Parameters</b>	<i>ipv6-address/prefix-length</i> — Specify the IPv6 prefix. <i>pool-name</i> — The name of the local DHCPv6 server pool

## prefix-threshold-stats

<b>Syntax</b>	<b>prefix-threshold-stats</b> <i>ipv6-address/prefix-length</i> <b>prefix-threshold-stats</b> <b>pool</b> <i>pool-name</i>
<b>Context</b>	clear>router>dhcp6>server
<b>Description</b>	This commands resets the peak stats in the prefix level threshold stats in the specified provision prefix or pool.
<b>Parameters</b>	<i>pool-name</i> — Specifies the name of the pool in local DHCPv6 server. <i>ipv6-address/prefix-length</i> — Specifies the name of theIPv6 prefix with prefix length.

## server-stats

<b>Syntax</b>	<b>server-stats</b>
<b>Context</b>	clear>router>dhcp6>server
<b>Description</b>	This command reset all stats of the specified local DHCPv6 server.

## statistics

<b>Syntax</b>	<b>statistics</b>
---------------	-------------------

## Clear Commands

**Context** clear>router>dhcp6

**Description** This command clears DHCP6 statistics.

---

## Debug Commands

### dhcp

<b>Syntax</b>	<b>[no] dhcp</b> <i>[ip-int-name]</i>
<b>Context</b>	debug>router>ip
<b>Description</b>	This command enables DHCP debugging. The <b>no</b> form of the command disables debugging.
<b>Parameters</b>	<i>ip-int-name</i> — Specifies the name of the IP interface. Interface names can be from 1 to 32 alphanumeric characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

### dhcp6

<b>Syntax</b>	<b>dhcp6</b> <i>[ip-int-name]</i> <b>no dhcp6</b>
<b>Context</b>	debug>router>ip
<b>Description</b>	This command enables DHCP debugging. The <b>no</b> form of the command disables debugging.
<b>Parameters</b>	<i>ip-int-name</i> — Specifies the name of the IP interface. Interface names can be from 1 to 32 alphanumeric characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

### detail-level

<b>Syntax</b>	<b>detail-level</b> { <b>low</b>   <b>medium</b>   <b>high</b> } <b>no detail-level</b>
<b>Context</b>	debug>router>ip>dhcp debug>router>local-dhcp-server debug>router>ip>dhcp6
<b>Description</b>	This command debugs the DHCP tracing detail level.

### local-dhcp-server

<b>Syntax</b>	<b>[no] local-dhcp-server</b> <i>server-name</i> [ <b>lease-address</b> <i>ip-address</i> ] <b>[no] local-dhcp-server</b> <i>server-name</i> [ <b>mac</b> <i>ieee-address</i> ]
<b>Context</b>	debug>router

## Debug Commands

<b>Description</b>	This command enables, disables or configures debugging for a local DHCP server.
<b>Parameters</b>	<i>server-name</i> — [32 chars max] <i>ip-address</i> — a.b.c.d <i>ieee-address</i> — xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx (cannot be all zeroes)

## mode

<b>Syntax</b>	<b>mode</b> {dropped-only   ingr-and-dropped   egr-ingr-and-dropped} <b>no mode</b>
<b>Context</b>	debug>router>ip>dhcp debug>router>local-dhcp-server debug>router>ip>dhcp6
<b>Description</b>	This command debugs the DHCP tracing detail level.

## wpp

<b>Syntax</b>	<b>[no] wpp</b>
<b>Context</b>	debug>router
<b>Description</b>	This command enables the context to configure debugging for the Web Portal Protocol.

## packet

<b>Syntax</b>	<b>[no] packet</b>
<b>Context</b>	debug>router>wpp
<b>Description</b>	This command configures WPP packet debugging.

## detail-level

<b>Syntax</b>	<b>detail-level</b> <i>detail-level</i>
<b>Default</b>	debug>router>wpp>packet
<b>Description</b>	This command specifies the detail level of the WPP packet debug output.
<b>Parameters</b>	<i>detail-level</i> — Specifies the detail level for WPP packet debugging. <b>Values</b> high, low

## portal

<b>Syntax</b>	<b>[no] portal</b> <i>wpp-portal-name</i>
<b>Context</b>	debug>router>wpp
<b>Description</b>	This command enables WPP debugging for the specified portal.
<b>Parameters</b>	<i>portal-name</i> — Specifies the name of this WPP portal.

## packet

<b>Syntax</b>	<b>[no] packet</b>
<b>Context</b>	debug>router>wpp>portal
<b>Description</b>	This command configures the WPP portal packet debugging.

## detail-level

<b>Syntax</b>	<b>detail-level</b> <i>detail-level</i>
<b>Context</b>	debug>router>wpp>portal>packet
<b>Description</b>	This command configures the detail level for WPP portal packet debugging.
<b>Parameters</b>	<i>detail-level</i> — Specifies the detail level for WPP portal packet debugging.
<b>Values</b>	high, low

## Tools Commands

---

### tools

<b>Syntax</b>	<b>tools</b>
<b>Context</b>	<root>
<b>Description</b>	This command enables the context to enable useful tools for debugging purposes.
<b>Default</b>	none
<b>Parameters</b>	<b>dump</b> — Enables dump tools for the various protocols. <b>perform</b> — Enables tools to perform specific tasks.

### perform

<b>Syntax</b>	<b>perform</b>
<b>Context</b>	tools
<b>Description</b>	This command enables the context to enable tools to perform specific tasks.
<b>Default</b>	none

### subscriber-mgmt

<b>Syntax</b>	<b>subscriber-mgmt</b>
<b>Context</b>	tools>perform
<b>Description</b>	This command enables tools to control subscriber management.

### edit-ppp-session

<b>Syntax</b>	<b>edit-ppp-session sap</b> <i>sap-id</i> <b>ip</b> <i>ip-address</i> [ <b>subscriber</b> <i>sub-ident-string</i> ] [ <b>sub-profile-string</b> <i>sub-profile-string</i> ] [ <b>sla-profile-string</b> <i>sla-profile-string</i> ] [ <b>inter-dest-id</b> <i>intermediate-destination-id</i> ] [ <b>ancp-string</b> <i>ancp-string</i> ] [ <b>app-profile-string</b> <i>app-profile-string</i> ]  <b>edit-ppp-session svc-id</b> <i>service-id</i> <b>ip</b> <i>ip-address</i> [ <b>subscriber</b> <i>sub-ident-string</i> ] [ <b>sub-profile-string</b> <i>sub-profile-string</i> ] [ <b>sla-profile-string</b> <i>sla-profile-string</i> ] [ <b>app-profile-string</b> <i>app-profile-string</i> ] [ <b>inter-dest-id</b> <i>intermediate-destination-id</i> ] [ <b>ancp-string</b> <i>ancp-string</i> ]
<b>Context</b>	tools>perform>subscriber-mgmt
<b>Description</b>	This command modifies PPP session information.



<b>Parameters</b>	<p><i>sap-id</i> — Specifies the physical port identifier portion of the SAP definition. See <a href="#">Common Service Commands on page 2062</a> for <i>sap-id</i> command syntax.</p> <p><i>ip-address</i> — Specifies the IP address.</p> <p><i>sub-ident-string</i> — Specifies a subscriber identification profile.</p> <p><i>sub-profile-string</i> — Specifies the subscriber profile string, up to 16 characters, maximum.</p> <p><i>service-id</i> — The service identification number that identifies the service in the domain.</p> <p><i>intermediate-destination-id</i> — Specifies the intermediate destination identifier which is encoded in the identification strings.</p> <p><b>ancp-string</b> <i>ancp-string</i> — Specifies the ASCII string of the DSLAM circuit ID name.</p> <p><i>app-profile-string</i> — Specifies an application profile string.</p>
-------------------	--

## eval-lease-state

<b>Syntax</b>	<b>eval-lease-state</b> [ <b>svc-id</b> <i>service-id</i> ] [ <b>sap</b> <i>sap-id</i> ] [ <b>subscriber</b> <i>sub-ident-string</i> ] [ <b>ip</b> <i>ip-address</i> ]
<b>Context</b>	tools>perform>subscriber-mgmt
<b>Description</b>	This command evaluates lease state.
<b>Parameters</b>	<p><i>sap-id</i> — Specifies the physical port identifier portion of the SAP definition. See <a href="#">Common Service Commands on page 2062</a> for <i>sap-id</i> command syntax.</p> <p><i>ip-address</i> — Specifies the a server's IP address. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).</p> <p><i>sub-ident-string</i> — Specifies the subscriber ID string, up to 32 characters, maximum.</p> <p><i>service-id</i> — Specifies an existing service ID.</p> <p><b>Values</b>      1 — 2147483647</p>

## local-user-db

<b>Syntax</b>	<b>local-user-db</b> <i>local-user-db-name</i>
<b>Context</b>	tools>perform>subscriber-mgmt
<b>Description</b>	This command enables tools for controlling the local user database.
<b>Parameters</b>	<i>local-user-db-name</i> — Specifies the name of a local user database.

## dhcp

<b>Syntax</b>	<b>dhcp</b>
<b>Context</b>	tools>perform>subscriber-mgmt>local-user-db
<b>Description</b>	This command contains the tools used for controlling DHCP entries in the local user database.

## host-lookup

<b>Syntax</b>	<b>host-lookup</b> [ <b>mac</b> <i>ieee-address</i> ] [ <b>remote-id</b> <i>remote-id</i> ] [ <b>sap-id</b> <i>sap-id</i> ] [ <b>service-id</b> <i>service-id</i> ] [ <b>string</b> <i>vso-string</i> ] [ <b>system-id</b> <i>system-id</i> ] [ <b>option60</b> <i>hex-string</i> ] [ <b>circuit-id</b> <i>circuit-id</i>   <b>circuit-id-hex</b> <i>circuit-id-hex</i> ]
<b>Context</b>	tools>perform>subscriber-mgmt>local-user-db>dhcp
<b>Description</b>	This command performs a lookup in the local user database. This command looks up the host with the match-list configured in the local user database.
<b>Parameters</b>	<p><b>mac</b> <i>ieee-address</i> — Specifies the 48-bit MAC address for the static ARP in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee, and ff are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.</p> <p><i>remote-id</i> — specifies what information goes into the remote-id sub-option in the DHCP relay packet.</p> <p><b>Values</b> Up to 255 characters maximum</p> <p><b>sap-id</b> — Specifies a SAP identifier to be used. See <a href="#">Common Service Commands on page 2062</a> for <i>sap-id</i> command syntax.</p> <p><i>service-id</i> — Specifies an existing subscriber service ID.</p> <p><b>Values</b> 1 — 2147483647</p> <p><i>vso-string</i> — Specifies a Vendor Specific Option (VSO) string.</p> <p><i>system-id</i> — Specifies the system ID.</p> <p><b>Values</b> up to 255 characters maximum.</p> <p><b>option60</b> <i>hex-string</i> — Specifies the content of option 60 for this lookup.</p> <p><b>Values</b> 0x0..0xFFFFFFFF (maximum 64 hex nibbles)</p> <p><b>circuit-id</b> <i>circuit-id</i> — specifies the circuit ID from the Option 82.</p> <p><b>circuit-id-hex</b> <i>circuit-id-hex</i> — Specifies the circuit ID in hexadecimal format from the Option 82.</p> <p><b>Values</b> 0x0..0xFFFFFFFF (maximum 254 hex nibbles)</p>

## ppp

<b>Syntax</b>	<b>ppp</b>
<b>Context</b>	tools>perform>subscriber-mgmt>local-user-db

**Description** This command contains the tools used to control PPP entries in the local user database.

## authentication

**Syntax** **authentication password** *password* [**mac** *ieee-address*] [**remote-id** *remote-id*] [**circuit-id** *circuit-id*] **user-name** *user-name* [**service-name** *service-name*]  
**authentication password** *password* [**mac** *ieee-address*] [**remote-id** *remote-id*] [**circuit-id-hex** *circuit-id-hex*] **user-name** *user-name* [**service-name** *service-name*]

**Context** tools>perform>subscriber-mgmt>local-user-db>ppp

**Description** This command authenticates PPP user name. As local user database PAP/CHAP authentication can only be used when the local user database is connected to the PPP node under the group interface, the user lookup will be performed with match-list username.

**Parameters** **password** *password* — specifies the password of this host up to 32 characters in length.  
**mac** *ieee-address* — Specifies information about the MAC address of the PPP session.  
*remote-id* — specifies what information goes into the remote-id sub-option in the DHCP relay packet.  
**Values** Up to 255 characters maximum  
**circuit-id** *circuit-id* — specifies the circuit ID from the Option 82.  
**circuit-id-hex** *circuit-id-hex* — Specifies the circuit ID in hexadecimal format from the Option 82.  
**Values** 0x0..0xFFFFFFFF (maximum 254 hex nibbles)  
**user-name** *user-name* — Specifies the PPP user name.  
**service-name** *service-name* —

## host-lookup

**Syntax** **host-lookup** [**mac** *ieee-address*] [**remote-id** *remote-id*] [**user-name** *user-name*] [**service-name** *service-name*] [**circuit-id** *circuit-id* | **circuit-id-hex** *circuit-id-hex*]

**Context** tools>perform>subscr-mgmt>loc-user-db>ppp

**Description** This command performs a lookup in the local user database.

**mac** *ieee-address* — Specifies the 48-bit MAC address for the static ARP in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee, and ff are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

**remote-id** *remote-id* — specifies what information goes into the remote-id sub-option in the DHCP relay packet.  
**Values** Up to 255 characters maximum  
**user-name** *user-name* — Specifies a user name up to 128 characters in length.  
**service-name** *service-name* — Specifies a PPP service name, up to 255 characters maximum.  
**circuit-id** *circuit-id* — specifies the circuit ID from the Option 82.

**circuit-id-hex** *circuit-id-hex* — Specifies the circuit ID in hexadecimal format from the Option 82.

**Values**      0x0..0xFFFFFFFF (maximum 254 hex nibbles)

# Triple Play Security

---

## In This Chapter

This chapter provides information about configuring specific security aspects for Triple Play services, including configuration process overview, and application notes.

Please note that the 7450 ESS supports many additional security features, which are described in the 7450 ESS System Management Guide.

Topics in this chapter include:

- [Triple Play Security Features on page 546](#)
  - [Anti-Spoofing Filters on page 546](#)
  - [Layer 2 Triple Play Security Features on page 548](#)
    - [MAC Pinning on page 548](#)
    - [MAC Protection on page 548](#)
    - [DoS Protection on page 549](#)
    - [VPLS Redirect Policy on page 551](#)
  - [ARP Handling on page 552](#)
  - [Web Portal Redirect on page 554](#)
- [Configuring Triple Play Security with CLI on page 557](#)
- [Common Configuration Tasks on page 558](#)

## Triple Play Security Features

---

### Anti-Spoofing Filters

This section describes the Alcatel-Lucent 7450 ESS acting as a Broadband Subscriber Aggregator (BSA).

Anti-spoofing is useful to prevent certain packets gaining unauthorized network access. Such packets originate from within a network and with an invalid source address.

Enabling anti-spoof filtering on a subscriber-facing SAP causes the anti-spoof table to be populated with all static and dynamic host information available on the SAP. Enabling anti-spoof filtering on the SAP will fail if any static hosts are defined without the proper addresses specified for the selected anti-spoof filter type.

When enabled, forwarding IP packets that ingress the SAP is dependent on a successful anti-spoof table match with an entry in the table. DHCP and non-IP packets (including ARP) are not subject to anti-spoof filtering. If an entry does not match the ingress packet, the packet will be silently discarded while incrementing the SAP discard counter.

Anti-spoof filtering is only allowed on VPLS SAPs and IES SAP-based IP interfaces. Anti-spoof filtering is not available on IES SDP bound IP interfaces. Anti-spoof filtering is not supported on Epipe services. Support for anti-spoofing is dependent on SAP based service interfaces.

Note that anti-spoofing filters, with type **ip-mac**, must be enabled to do Enhanced Subscriber Management (as described in section [Triple Play Enhanced Subscriber Management on page 729](#)).

Topics in this section are:

- [Anti-spoofing Filter Types on page 547](#)
- [Filtering Packets on page 547](#)

## Anti-spoofing Filter Types

A SAP or interface that supports anti-spoof filtering can be configured to use one of three types of anti-spoof tables. The type of table used by the SAP is dependent on the type of anti-spoof filtering desired, only one anti-spoofing table type is supported per SAP:

- When only the incoming source MAC address is to be verified, the source MAC table must be defined (anti-spoof type = **mac**).
- When only the incoming source IP address is to be verified, the source IP table must be defined (anti-spoof type = **ip**).
- When both the incoming source MAC and source IP addresses are to be verified, the combination source IP and source MAC table must be defined (anti-spoof type = **ip-mac**).

Note that setting the anti-spoof filter type for the SAP is dependent on pre-existing static host definitions, for example, attempting to set the SAP anti-spoof filtering to mac will fail if any static hosts exist that do not have a defined MAC address.

The anti-spoof table of a SAP or interface will be populated from the DHCP lease state table and from any statically defined hosts on the SAP or interface.

---

## Filtering Packets

Packets from a client that match an anti-spoof filter entry when anti-spoof filtering is enabled are allowed to be further processed by the system. The matching packet is still subject to other forwarding criteria including potentially ACL filtering.

All packets that are not exempt from anti-spoofing and do not match a entry in the anti-spoof table are discarded. Every discard event will increment the SAP discard packet counter. The discard event is not logged or alarmed, but a threshold alarm could be configured for the counter (see Configuring System Monitoring Thresholds in the 7450 ESS Basic System Configuration Guide).

Not all ingress packets are subject to the anti-spoof filtering when enabled. Non-IP packets are exempt for anti-spoof filter lookups and are allowed to be further processed by the system. This includes ARP requests and replies, as well as PPPoE packets. The only IP packets exempt from anti-spoof filtering are DHCP packets destined to the server UDP port 67. DHCP packets destined to the client UDP port number (port 68) are not exempt.

## Layer 2 Triple Play Security Features

Topics in this section include:

- [MAC Pinning on page 548](#)
  - [MAC Protection on page 548](#)
  - [DoS Protection on page 549](#)
  - [VPLS Redirect Policy on page 551](#)
- 

### MAC Pinning

This section describes the Alcatel-Lucent 7450 ESS acting as a Broadband Subscriber Aggregator (BSA).

DHCP snooping and IP/MAC auto-filters can be employed to prevent Theft of Service (by a malicious user spoofing another user's address). However, these auto-filters do not discard non-IP packets such as PPPoE packets, thus potentially allowing a MAC address to be relearned on another SAP. MAC pinning closes this loophole, by not allowing a MAC address to be relearned on another SAP.

When MAC pinning is enabled, a MAC address learned on one SAP or SDP can not be re-learned on another SAP or SDP in the same VPLS, until the FIB entry for the MAC address times out. (Note that in case MAC aging is disabled, MAC entries on a SAP/SDP with MAC pinning enabled will effectively become permanent.)

MAC pinning is implicitly enabled when DHCP auto-filters are enabled, and cannot be disabled. For MAC addressing learned during DHCP address assignment (when DHCP snooping function is active at least on one port of the VPLS), the MAC address is tied to a given SAP for the duration of the DHCP lease.

When a SAP or spoke SDP is part of a Residential Split Horizon Group (RSHG), MAC pinning is automatically enabled at creation of the SAP, but can be disabled if desired.

---

### MAC Protection

In a Layer 2 environment, a malicious subscriber could create a denial-of-service attack by sending Ethernet frames, with as source MAC address the address of a gateway (for example, the IP next hop upstream). As MAC learning is typically enabled, this would move the learned gateway MAC from the uplink SAP or SDP to the subscriber's SAP, causing all communication to the gateway to be disrupted. If a local content server is attached to the same VPLS, a similar attack could be launched against it.



Communication between subscribers can be disallowed using Split Horizon Groups, but this by itself will not be sufficient to prevent such an attack.

The solution is to create a mechanism to explicitly protect some MAC addresses against being relearned on other SAPs.

The **mac-protect** feature on the Alcatel-Lucent 7450 ESS allows a list of special MAC addresses to be configured in a VPLS. Two checks can then be made on incoming packets against these protected MAC addresses:

- **[no] auto-learn-mac-protect**: Used to enable the automatic protection of source MAC addresses learned on the associated object. MAC protection is used in conjunction with **restrict-protected-src**, **restrict-unprotected-dst** and **mac-protect**. When this command is applied or removed, the MAC addresses are cleared from the related object.
- **restrict-protected-src**: Used to prevent denial-of-service attacks. If the source MAC address of a packet from a subscriber matches a protected entry, probably this subscriber tried to impersonate the gateway or server. Such packets are discarded, a trap is generated, and the SAP on which it arrived is placed operationally down.
- **restrict-unprotected-dst**: Used to force traffic from subscribers to only go towards a few defined destinations (the gateways or servers). Any packet from a subscriber whose destination MAC address does not match a protected entry is discarded.

---

## DoS Protection

This section describes the mechanisms and limitations of DoS protection related to subscriber management snooping functions. This feature is only supported on 7750 SR-Series and 7450 ESS-Series redundant chassis models. In subscriber aggregation networks, these routers play an active role in several protocols. Subscribers either intentionally or unknowingly interfere with the operation of the node's processing capacity (for example, excessive ARP handling) or other user traffic.

Routing protocols such as OSPF and ISIS could also be a threat as packets can be injected by customers (erroneously or maliciously) which could cause high CPM overload. Service providers are concerned about DoS protection including DoS attacks when acting as a subscriber aggregation device and guarding against DoS attacks using unprovisioned protocols.

---

## Subscriber Aggregation Network

In a subscriber aggregation network, multiple devices such as the 7750 SR and 7450 ESS routers provide access to a DHCP or a RADIUS server. These servers usually do not scale high enough to provide the means to control access to snooping functions through a controlled queue. It is possible, under severe conditions, that the network could become unavailable if the node cannot handle requests from subscribers.

Because the IOMs cannot be scaled to provide a per-subscriber queue to control traffic, a monitoring function, handled by the CPM, is provided. With this monitoring system, the CPM tracks the number of control plane messages set per subscriber and limits the rate to a specified level and provides feedback using event generation to alert a centralized system of a possible DoS attack.

The CPM provides a prioritized access to the CPU. Since the number of control packets expected from a subscriber should have a low rate, and under normal conditions, the system will provide a rate limit on a per subscriber/MAC basis and will drop a subscriber control packet before it is queued or processed by the CPU. The system will be configured with expected arrival rate of per MAC/subscriber control packet rates and optionally total rate per interface/SAP.

The system maintains a per-second running rate monitor per SAP and per MAC. If an entry is using more than the configured rate, the system will not forward that packet to be queued. Every existing subscriber host will be monitored. A subscriber host will be flagged and the system observed with an excessive rate of control packets. In the case of PPPoE, the CPM will monitor subscriber hosts before the IP address is provided by the SAP/MAC/session-id combination.

The control protocols affected by this mechanism include:

- ARP (in arp-reply-agent)
- DHCP (for discover and renew)
- ICMP
- PPPoE
- IGMP

---

## Network Control Filtering

Alcatel-Lucent's 7750 SR and 7450 ESS can block network control traffic for unconfigured protocols. For example, if OSPF is not configured on an IP interface, all OSPF-related traffic should be dropped before the traffic reaches the CPU.

Protocols are blocked based on whether that particular protocol is configured to run on the given IP interface. It is not required to re-configure the permitted protocols.

Protocol traffic control by this mechanism includes:

- OSPFv2
- OSPFv3
- IS-IS
- RSVP-TE
- LDP
- RIP
- PIM

- MLD
  - IGMP
  - BGP
  - BFD
  - L2PT
  - PPP
  - DHCP
- 

## VPLS Redirect Policy

This section describes the Alcatel-Lucent 7450 ESS acting as a Broadband Subscriber Aggregator (BSA) with Layer 2 aggregation towards a Broadband Subscriber Router (BSR).

In a Triple Play network it may be desired to route some traffic from/to subscribers through a Deep Packet Inspection (DPI) device, for example, to limit peer-to-peer traffic. However such a DPI device typically has limited bandwidth available, so only those packets that need inspection should be sent to it.

In a Layer 3 network, such policy-based redirection can be achieved using “next-hop redirect” ACL entries. In a layer 2 (VPLS) aggregation network, the same result can be achieved using “redirect to SAP” or “redirect to SDP” policy.

Refer to the ACL Next-Hop for VPLS section in the 7450 ESS Services Overview Guide.

## ARP Handling

---

### ARP Reply Agent

This section describes the Alcatel-Lucent 7450 ESS acting as a Broadband Subscriber Aggregator (BSA).

In Triple Play networks, typically downstream broadcast is not allowed on subscriber SAPs. As a result, subscribers can not receive ARP requests from the network. Instead, the 7450 ESS will respond to ARP requests from the network, with information from the DHCP lease state table.

In the upstream direction (towards the network), the ARP reply agent intercepts ARP Requests on subscriber SAPs, and checks them against the DHCP lease state table. The purpose is to prevent a malicious subscriber spoofing ARP Request or ARP reply messages and thus populating the upstream router's ARP table with incorrect entries.

When the keyword **sub-ident** is added in the ARP reply agent configuration, also the subscriber identity is checked. If an upstream ARP request is targeted to the same subscriber, it is dropped. Otherwise, it is flooded to all VPLS interfaces outside the received Split Horizon Group (SHG).

Static hosts can be defined on the SAP using the **host** command. Dynamic hosts are enabled on the system by enabling the **lease-populate** command in the SAP's **dhcp** context. In the event that both a static host and a dynamic host share the same IP and MAC address, the VPLS ARP reply agent will retain the host information until both the static and dynamic information are removed. In the event that both a static and dynamic host share the same IP address, but different MAC addresses, the VPLS ARP reply agent is populated with the static host information.

In brief, the ARP Replay Agent operation is as follows:

- For ARP request received from a customer SAP:
  - first check in DHCP lease state table - if no match: discard
  - if (**sub-ident** enabled) and (destination = same subscriber): discard
  - otherwise: flood to all SAPs/SDPs outside this SHG
- For ARP request received from the network:
  - lookup IP address in DHCP lease state table - if no match: discard
  - otherwise: respond with MAC address from the DHCP lease state table

## Dynamic ARP Table Population

This section describes the Alcatel-Lucent 7450 ESS acting as a Broadband Subscriber Aggregator (BSA) with Layer 3 forwarding towards the network.

In an IES service, the system's ARP table can be populated dynamically using entries in the DHCP lease state table (in turn populated from snooping DHCP ACK messages (see [DHCP Snooping on page 322](#))), and from static hosts defined on the SAP. In the router ARP table these are indicated with state managed.

In the event that both a static host is created with the same IP and MAC address as an existing managed entry, creation will fail and a trap is generated.

In the event that a DHCP Lease needs to be populated with the same IP and MAC address as an existing static host entry, creation will fail and a trap is generated.

No **static-arp** creation is possible when combined with **arp-populate**.

---

## Local Proxy ARP

This section describes the Alcatel-Lucent 7450 ESS acting as a Broadband Subscriber Aggregator (BSA) with Layer 3 forwarding towards the network.

Local proxy ARP allows the Alcatel-Lucent 7450 ESS to respond to ARP requests received on an interface, for an IP address which is part of a subnet assigned to the interface. When the local proxy ARP feature is enabled, the switch responds to all ARP requests for IP addresses belonging to the subnet with the MAC address of the interface, and forwards all traffic between hosts in the subnet.

This feature is intended to be used in situations (such as DSL aggregation networks) when hosts belonging to the same subnet are prevented from directly communicating with each other over the subnet by the configuration of the switch (or DSLAM) to which they are connected.

Note: When local-proxy-arp is enabled under a IES service, all ICMP redirects on the ports associated with the service are automatically blocked. This prevents users from learning each other's MAC address (from ICMP redirects).

The implementation of proxy ARP with support for local proxy ARP allows the 7450 ESS to respond to ARP requests in the subnet assigned to an IES or VPRN interface, thus allowing multiple customers to share the same IP subnet.

## Web Portal Redirect

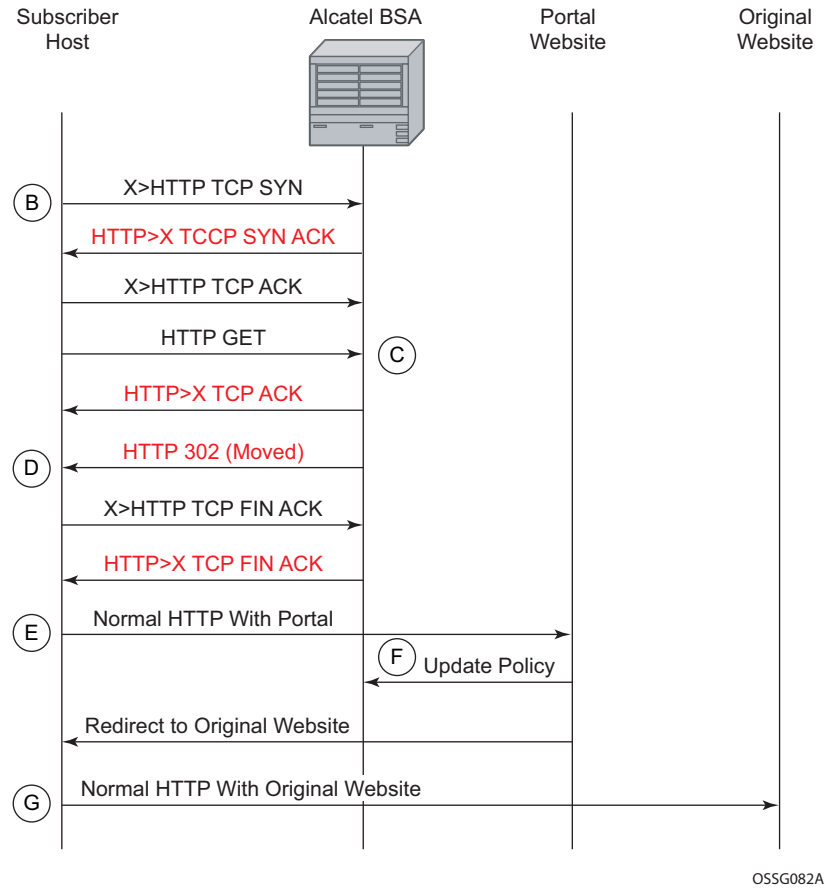
This section describes the Alcatel-Lucent 7450 ESS acting as a Broadband Subscriber Aggregator (BSA).

In a Triple Play network it may not be desired (or feasible) to perform manual provisioning of new services and service changes. The ideal way of working is automatic provisioning, with the end-user supplying his details at a retailer, or (if physical connectivity is already present through an on-line customer portal).

The 7450 ESS supports a special ACL that automatically redirects subscribers to a predefined URL. This is done by sending a HTTP 302 (moved) message to the subscriber, regardless of the requested URL.

The message flow is as follows (see [Figure 25](#) below):

- a. The subscriber gets an IP address using DHCP (if the customer is trying to use a static IP he will be blocked by the anti-spoofing filter).
- b. The subscriber tries to connect to a website (TCP SYN, TCP ACK, HTTP GET).
- c. The 7450 ESS intercepts the HTTP GET request and discards it.
- d. The 7450 ESS then responds to the subscriber with a HTTP 302 message (service temporarily unavailable/moved), containing a new target URL (that of the portal) configured by the operator. This target URL can include the subscriber's IP and MAC addresses as part of the portal's URL string.
- e. The subscriber's web browser will close the original TCP connection and open a new connection to the web portal, where the subscriber can sign up or change his/her service Profile.
- f. After approving the changes, the web portal updates the ACL (directly or through another system such as the Alcatel-Lucent 5750 SSC) to remove the redirection policy.
- g. The subscriber can now connect to the original site.



**Figure 25: IP Illustration of Message Flow in Web Portal Redirect**

The items in red text in [Figure 25](#) are messages the 7450 ESS will send (masquerading as the destination), regardless of the destination IP address or type of service.

The following displays information that can optionally be added as variables in the portal URL (http-redirect url):

- \$IP – Customer's IP address
- \$MAC – Customer's MAC address
- \$URL – Original requested URL
- \$SAP – Customer's SAP
- \$SUB – Customer's subscriber identification string
- \$CID – string that represents the circuit-id or interface-id of the subscriber host (hexadecimal format)
- \$RID – string that represents the remote-id of the subscriber host (hexadecimal format)

Note that the subscriber's IP and MAC address variables are populated from the anti-spoofing list, and thus anti-spoofing must be enabled (see section [Anti-Spoofing Filters on page 546](#)).

Since most web sites are accessed using the domain name, the 7450 ESS will need to allow DNS queries, and an ACL entry to this effect should be included in the filter (see example in section [Configuring Web Portal Redirect on page 571](#)).



## Configuring Triple Play Security with CLI

This section provides information to configure Residential Broadband Aggregation services using the command line interface. It is assumed that the reader is familiar with VPLS and IES services.

Topics in this section include:

- [Common Configuration Tasks on page 558](#)
  - [Configuring Anti-Spoofing Filters on page 558](#)
  - [Configuring Triple Play Security features on page 559](#)
  - [Configuring ARP Handling on page 565](#)
  - [Configuring Web Portal Redirect on page 571](#)

## Common Configuration Tasks

Topics in this section are:

- [Configuring Anti-Spoofing Filters on page 558](#)
  - [Configuring Triple Play Security features on page 559](#)
    - [Configuring MAC Pinning on page 559](#)
    - [Configuring MAC Protection on page 560](#)
    - [Configuring VPLS Redirect Policy on page 562](#)
    - [Configuring VPLS Redirect Policy on page 562](#)
  - [Configuring ARP Handling on page 565](#)
  - [Configuring Web Portal Redirect on page 571](#)
- 

### Configuring Anti-Spoofing Filters

Anti-spoofing filters are used to prevent malicious subscribers sending IP packets with a forged IP or MAC address, and thus mis-directing traffic. The anti-spoofing filter is populated from the DHCP lease state table, and DHCP snooping must be enabled on the SAP.

There are three types of filters (MAC, IP and IP+MAC), one type is allowed per SAP.

The following displays an IES service interface configuration with anti-spoofing.

```
A:ALA-48>config>service>ies# info
-----
      interface "test123" create
        address 10.10.42.41/24
        local-proxy-arp
        proxy-arp
          policy-statement "ProxyARP"
        exit
        sap 1/1/7:0 create
          anti-spoof ip
        exit
        arp-populate
        dhcp
          lease-populate 1
          no shutdown
        exit
      exit
    no shutdown
-----
A:ALA-48>config>service>ies#
```

## Configuring Triple Play Security features

Topics in this section are:

- [Configuring MAC Pinning on page 559](#)
  - [Configuring MAC Protection on page 560](#)
  - [Configuring VPLS Redirect Policy on page 562](#)
  - [Configuring VPLS Redirect Policy on page 562](#)
- 

### Configuring MAC Pinning

The following example displays a partial BSA configuration with MAC pinning enabled on a SAP:

```
A:ALA-48>config>service# info
-----
vpls 800 customer 6001 create
    description "VPLS with residential split horizon for DSL"
    stp
        shutdown
    exit
    sap 2/1/4:100 split-horizon-group "DSL-group2" create
        description "SAP for RSHG"
        mac-pinning
    exit
    no shutdown
-----
A:ALA-48>config>service#
```

## Configuring MAC Protection

---

### Preventing Access By Residential Subscribers Using Protected (Gateway) MAC Addresses

The first step is to create a list of MAC addresses to be protected, the second step is to prevent access using these source addresses inside an SHG or a SAP.

The following example displays a partial BSA configuration with some protected MAC addresses on any SAP created inside the SHG:

```
A:ALA-48>config>service# info
-----
vpls 800 customer 6001 create
    no shutdown
    split-horizon-group "mygroup" create
        restrict-protected-src
    exit
    description "VPLS with residential split horizon for DSL"
    mac-protect
        mac 00:00:17:FE:82:D8
        mac 93:33:00:00:BF:92
    exit
-----
A:ALA-48>config>service#
```

**Restricting Access By Residential Subscribers To a Small List Of Upstream MAC Addresses:**

The first step is to create a list of MAC addresses to be protected, the second step is to restrict access to these addresses only from an SHG or a SAP. (If the MAC address of an upstream server is not known, it can be discovered using e.g. the cpe-ping OAM tool.)

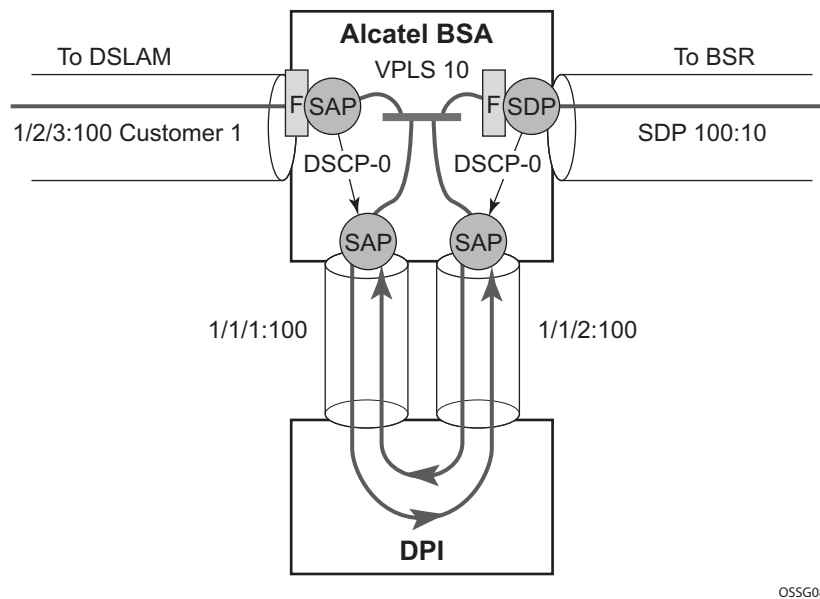
The following example displays a partial BSA configuration with restricted access to some MAC addresses from a specified SAP (an unrestricted access from any other SAP within the VPLS):

```
A:ALA-48>config>service# info
-----
vpls 800 customer 6001 create
    no shutdown
    description "VPLS with restricted access on a SAP"
    mac-protect
        mac 00:00:17:FE:82:D8
        mac 93:33:00:00:BF:92
    exit
    sap 1/1/4:30 create
        restrict-unprotected-dst
    exit
-----
A:ALA-48>config>service#
```

## Configuring VPLS Redirect Policy

- [Creating the Filter on page 563](#)
- [Applying the Filter to a VPLS Service on page 564](#)

Figure 26 displays an IP filter entry configuration for VPLS redirect policy.



**Figure 26: VPLS Redirect Policy Example**

Information about defining and applying IP and MAC filters is described in the 7450 ESS Router Configuration Guide.

## Creating the Filter

The following displays a redirect filter entry:

```
A:ALA-A>config>filter# info
-----
    ip-filter 10
      default-action forward
      entry 10
        match
          dscp be
        exit
      action forward next-hop sap 1/1/1:100
      exit
    exit
  exit
ip-filter 11
  default-action forward
  entry 10
    match
      dscp be
    exit
    dscp be
  exit
  action forward next-hop sap 1/1/2:100
  exit
exit
-----
A:ALA-A>config>filter#
```

### Applying the Filter to a VPLS Service

The following displays how the redirection filter configured above is assigned to the ingress SAP from the DSLAM, and the ingress SDP from the BSR:

```
A:ALA-A>config>service>vpls# info
-----
      vpls 10 customer 1 create
        description "vpls10"
          sap 1/2/3:100 create
            ingress ip filter 10
          exit
          sap 1/1/1:100 create
            exit
          sap 1/1/2:100 create
            exit
          mesh-sdp 100:10 create
            ingress ip filter 11
          exit
        exit
      exit
-----
A:ALA-A>config>service>vpls#
```



## Configuring ARP Handling

Topics in this section are:

- [Configuring Proxy ARP on page 565](#)
  - [Configuring Local Proxy ARP on page 566](#)
  - [Configuring ARP Reply Agent in a VPLS Service on page 567](#)
  - [Configuring Automatic ARP Table Population in an IES Interface on page 569](#)
- 

### Configuring Proxy ARP

The implementation of proxy ARP with support for local proxy ARP allows the 7450 ESS to respond to ARP requests in the subnet assigned to an IES interface.

Configuring this command will allow multiple customers to share the same IP subnet.

The following example displays an IES proxy ARP configuration:

```
A:ALA-48>config>service>ies# info
-----
      interface "test123" create
          address 10.10.42.41/24
          local-proxy-arp
          proxy-arp-policy "ProxyARP"
          exit
      exit
      no shutdown
-----
A:ALA-48>config>service>ies#
```

### Configuring Local Proxy ARP

When local proxy ARP is enabled on an IP interface, the 7450 ESS responds to all ARP requests for IP addresses belonging to the subnet with its own MAC address, and forwards all traffic between hosts in that subnet. Local proxy ARP is disabled by default.

Note: When local-proxy-arp is enabled under a IES service, all ICMP redirects on the ports associated with the service are automatically blocked. This prevents users from learning each other's MAC address (from ICMP redirects).

The following example displays a local proxy ARP IES configuration:

```
A:ALA-A>config>service>ies# info
-----
      interface "test" create
      shutdown
      address 10.10.36.2/24
      local-proxy-arp
      exit
-----
A:ALA-A>config>service>ies#
```

## Configuring ARP Reply Agent in a VPLS Service

When ARP reply agent is enabled, the ESS-Series will respond to ARP requests from the network, with information from the DHCP lease state table.

In the upstream direction (towards the network), the ARP reply agent intercepts ARP requests on subscriber SAPs, and checks them against the DHCP lease state table. The purpose is to prevent a malicious subscriber spoofing ARP request or ARP reply messages and thus populating the upstream router's ARP table with incorrect entries.

The following example displays a partial BSA configuration with ARP Reply Agent enabled on a SAP:

```
A:ALA-48>config>service# info
-----
...
    vpls 800 customer 6001 create
        description "VPLS with ARP Reply Agent active"
        sap 2/1/4:100 split-horizon-group "DSL-group2" create
            arp-reply-agent sub-ident
        exit
        sap 3/1/4:200 split-horizon-group "DSL-group2" create
            arp-reply-agent sub-ident
        exit
        no shutdown
    ...
-----
A:ALA-48>config>service#
```

### Configuring Remote Proxy ARP

The following example displays the IES configuration to enable remote proxy ARP:

```
A:ALA-49>config>service>ies# info
-----
      interface "test-1A" create
          address 10.10.26.3/24
          remote-proxy-arp
      exit
      no shutdown
-----
A:ALA-49>config>service>ies#
```

## Configuring Automatic ARP Table Population in an IES Interface

The following example displays the IES DHCP configuration to enable automatic population of the ARP table using snooped DHCP information on an IES interface:

```
A:ALA-1>config>service>ies>if# info
-----
      arp-populate
      dhcp
        description "snooping_only"
        lease-populate 1
        no shutdown
      exit
-----
A:ALA-1>config>service>ies>if#
```

## **Configuring CPU Protection**

CPU Protection can be used to protect the SR OS router in subscriber management scenarios. Refer to the SR OS System Management Guide for information about CPU Protection operation and configuration.

## Configuring Web Portal Redirect

The generic CLI structure for defining and applying IP and MAC filters is described in the 7450 ESS Router Configuration Guide.

The following example displays an IP filter entry configuration for web-portal redirect:

```
A:ALA-A>config>filter# info
-----
    ip-filter 10 create
        description "filter to forward DNS and web traffic to my portal; redirect all
other web traffic to the portal and drop everything else"
        default-action drop
        entry 10 create
            description "allows DNS traffic"
            match protocol 17
            dst-port 53
            exit
            action forward
        exit
        entry 20 create
            description "allows web traffic destined to portal (IP address 10.0.0.1)"
            match protocol 6
            dst-port eq 80
            dst-ip 10.0.0.1
            exit
            action forward
        exit
        entry 30
            description "redirects all web traffic to portal"
            match protocol 6
            dst-port eq 80
            exit
            action http-redirect http://www.myportal.com/defaultportal
/login.cgi?ip=$IP&mac=$MAC&orig_url=$URL&usb=$SUB
            exit
        exit
-----
A:ALA-A>config>filter#
```

- Filter entry 10 in the example output allows the customer to access DNS to get the IP address of the original website they are trying to view.
- Entry 20 allows HTTP packets destined to the captive portal itself to be forwarded. Note that the actual IP address (a.b.c.d) needs to be entered, not the DNS name (“www.myportal.com”). The IP address can be easily resolved from the 7450 ESS CLI using the “ping” command.
- Entry 30 (which is the last option that does not drop the customer packets) checks for HTTP protocol and then starts the redirection process:
  - The 7450 ESS will intercept the HTTP GET from the subscriber and respond with an HTTP 302 (temporarily moved) with the URL configured in the filter entry. This URL can contain some variables, notably the customer IP and MAC addresses to allow the portal to create an entry for the customer. The original requested URL is also included to redirect the client site back to the original requested site when the process is done.
  - The client will then close the connection with the original IP address and open a connection to the redirected server. Entry 20 will allow this connection.

The following displays how the redirection filter configured above is assigned to an ingress SAP:

```
A:ALA-A>config>service>vpls# info
-----
vpls 3 customer 6 create
  description "VPLS with web portal redirection filter applied"
  sap 2/1/5:0 create
    ingress
      filter ip 10
    exit
  exit
  no shutdown
exit
-----
A:ALA-A>config>service>vpls#
```



---

# Triple Play Security Command Reference

---

## Command Hierarchies

- [Anti-Spoofing Commands on page 573](#)
- [Layer 2 Security Commands on page 574](#)
- [ARP Handling Commands on page 574](#)

Note: This command tree is limited to those commands specific to Triple Play security. For the full command trees of a particular service type, consult the 7450 ESS Services Overview Guide.

## Anti-Spoofing Commands

```

config
  — service
    — vpls service-id [customer customer-id] [vpn vpn-id] [m-vpls]
      — sap sap-id [split-horizon-group group-name]
        — anti-spoof {ip | mac | ip-mac}
        — no anti-spoof

config
  — service
    — ies service-id [customer customer-id] [vpn vpn-id]
      — [no] interface ip-int-name
        — sap sap-id
          — anti-spoof {ip | mac | ip-mac}
          — no anti-spoof

```

## Layer 2 Security Commands

```

config
— service
— [no] vpls service-id [customer customer-id] [vpn vpn-id] [m-vpls]
— mac-protect
— [no] mac [ieee-mac-address]
— [no] split-horizon-group [group-name] [residential-group]
— [no] restrict-protected-src
— [no] restrict-unprotected-dst
— mesh-sdp sdp-id[:vc-id] [vc-type {ether | vlan}]
— [no] mac-pinning
— [no] restrict-protected-src
— [no] restrict-unprotected-dst
— spoke-sdp sdp-id[:vc-id] [vc-type {ether | vlan | vpls}] [split-horizon-group
group-name]
— [no] mac-pinning

```

## ARP Handling Commands

```

config
— service
— [no] vpls service-id [customer customer-id] [vpn vpn-id] [m-vpls]
— sap sap-id
— arp-reply-agent [sub-ident]
— no arp-reply-agent

config
— service
— ies service-id [customer customer-id] [vpn vpn-id]
— [no] interface ip-int-name
— [no] arp-populate
— arp-timeout seconds
— no arp-timeout
— [no] local-proxy-arp
— proxy-arp-policy policy-name [policy-name...(up to 5 max)]
— no proxy-arp-policy
— [no] remote-proxy-arp
— static-arp ip-address ieee-address
— no static-arp ip-address [ieee-address]

```

---

## Triple Play Security Configuration Commands

Topics in this section include:

- [Triple Play Anti-Spoofing Commands on page 576](#)
- [Triple Play Layer 2 Security Commands on page 577](#)
- [ARP Handling Commands on page 581](#)
- [Show Commands on page 586](#)

---

## Triple Play Anti-Spoofing Commands

### anti-spoof

<b>Syntax</b>	<b>anti-spoof {ip   mac   ip-mac}</b> <b>no anti-spoof</b>
<b>Context</b>	config>service>vpls>sap config>service>ies>if>sap
<b>Description</b>	<p>This command enables anti-spoof filtering and optionally changes the anti-spoof matching type for the SAP.</p> <p>The type of anti-spoof filtering defines what information in the incoming packet is used to generate the criteria to lookup an entry in the anti-spoof filter table. The type parameter (<b>ip</b>, <b>mac</b>, <b>ip-mac</b>) defines the anti-spoof filter type enforced by the SAP when anti-spoof filtering is enabled.</p> <p>The <b>no</b> form of the command disables anti-spoof filtering on the SAP.</p>
<b>Default</b>	<b>no anti-spoof</b>
<b>Parameters</b>	<p><b>ip</b> — Configures SAP anti-spoof filtering to use only the source IP address in its lookup. If a static host exists on the SAP without an IP address specified, the <b>anti-spoof ip</b> command will fail.</p> <p><b>mac</b> — Configures SAP anti-spoof filtering to use only the source MAC address in its lookup. If a static host exists on the SAP without a specified MAC address, the <b>anti-spoof mac</b> command will fail.</p> <p><b>ip-mac</b> — Configures SAP anti-spoof filtering to use both the source IP address and the source MAC address in its lookup. If a static host exists on the SAP without both the IP address and MAC address specified, the <b>anti-spoof ip-mac</b> command will fail.</p>

## Triple Play Layer 2 Security Commands

### split-horizon-group

<b>Syntax</b>	<b>[no] split-horizon-group</b> [ <i>group-name</i> ] [ <i>residential-group</i> ]
<b>Context</b>	config>service>vpls
<b>Description</b>	<p>This command creates a new split horizon group for the VPLS instance. Traffic arriving on a SAP or spoke SDP within this split horizon group will not be copied to other SAPs or spoke SDPs in the same split horizon group.</p> <p>A split horizon group must be created before SAPs and spoke SDPs can be assigned to the group.</p> <p>The split horizon group is defined within the context of a single VPLS. The same group-name can be re-used in different VPLS instances.</p> <p>Up to 30 split horizon groups can be defined per VPLS instance.</p> <p>The <b>no</b> form of the command removes the group name from the configuration.</p>
<b>Parameters</b>	<p><i>group-name</i> — Specifies the name of the split horizon group to which the SDP belongs.</p> <p><i>residential-group</i> — Defines a split horizon group as a residential split horizon group (RSHG). Doing so entails that:</p> <ul style="list-style-type: none"> <li>a) SAPs which are members of this Residential Split Horizon Group will have: <ul style="list-style-type: none"> <li>– Double-pass queuing at ingress as default setting (can be disabled)</li> <li>– STP disabled (can <u>not</u> be enabled)</li> <li>– ARP reply agent enabled per default (can be disabled)</li> <li>– MAC pinning enabled per default (can be disabled)</li> <li>– Besides the multicast downstream also broadcast packets are discarded thus also blocking the unknown, flooded traffic.</li> </ul> </li> <li>b) Spoke SDPs which are members of this Residential Split Horizon Group will have: <ul style="list-style-type: none"> <li>– Downstream multicast traffic supported</li> <li>– Double-pass queuing is not applicable</li> <li>– STP is disabled (can be enabled)</li> <li>– ARP reply agent is not applicable (dhcp-lease-states are not supported on spoke SDPs)</li> <li>– MAC pinning enabled per default (can be disabled)</li> </ul> </li> </ul>
<b>Default</b>	A split horizon group is by default not created as a residential-group.

## mac-protect

<b>Syntax</b>	<b>mac-protect</b>
<b>Context</b>	config>service>vpls
<b>Description</b>	This command indicates whether or not this MAC is protected on the MAC protect list. When enabled, the agent will protect the MAC from being learned or re-learned on a SAP, spoke SDP or mesh-SDP that has restricted learning enabled. The MAC protect list is used in conjunction with <b>restrict-protected-src</b> , <b>restrict-unprotected-dst</b> and <b>auto-learn-mac-protect</b> .
<b>Default</b>	disabled

## mac

<b>Syntax</b>	<b>[no] mac <i>ieee-address</i></b>
<b>Context</b>	config>service>vpls>mac-protect
<b>Description</b>	This command specifies the 48-bit IEEE 802.3 MAC address.
<b>Parameters</b>	<i>ieee-address</i> — Specifies the 48-bit MAC address in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee, and ff are hexadecimal numbers.

## auto-learn-mac-protect

<b>Syntax</b>	<b>[no] auto-learn-mac-protect</b>
<b>Context</b>	config>service>vpls>sap config>service>vpls>spoke-sdp config>service>vpls >mesh-sdp config>service>vpls>split-horizon-group config>service>vpls>endpoint config>service>pw-template config>service>pw-template>split-horizon-group
<b>Description</b>	<p>This command enables the automatic protection of source MAC addresses learned on the associated object. MAC protection is used in conjunction with restrict-protected-src, restrict-unprotected-dst and mac-protect. When this command is applied or removed, the MAC addresses are cleared from the related object.</p> <p>When the auto-learn-mac-protect is enabled on an SHG the action only applies to the associated SAPs (no action is taken by default for spoke SDPs in the SHG). In order to enable this function for spoke SDPs within a SHG, the auto-learn-mac-protect must be enabled explicitly under the spoke-SDP. If required, auto-learn-mac-protect can also be enabled explicitly under specific SAPs within the SHG.</p>
<b>Default</b>	no auto-learn-mac-protect

## restrict-protected-src

<b>Syntax</b>	<b>restrict-protected-src</b> [ <i>alarm-only</i>   <i>discard-frame</i> ] <b>no restrict-protected-src</b>
<b>Context</b>	config>service>vpls>sap config>service>vpls>spoke-sdp config>service>vpls>mesh-sdp config>service>vpls>split-horizon-group config>service>vpls>endpoint config>service>pw-template> config>service>pw-template>split-horizon-group
<b>Description</b>	<p>This command indicates how the agent will handle relearn requests for protected MAC addresses, either manually added using the mac-protect command or automatically added using the auto-learn-mac-protect command. While enabled all packets entering the configured SAP, spoke-SDP, mesh-SDP, or any SAP that is part of the configured split horizon group (SHG) will be verified not to contain a protected source MAC address. If the packet is found to contain such an address, the action taken depends on the parameter specified on the restrict-protected-src command, namely:</p> <ul style="list-style-type: none"> <li>• No parameter The packet will be discarded, an alarm will be generated and the SAP, spoke-SDP or mesh-SDP will be set operationally down. The SAP, spoke-SDP or mesh-SDP must be shutdown and enabled (no shutdown) for this state to be cleared.</li> <li>• alarm-only The packet will be forwarded, an alarm will be generated but the source MAC is not learned on the SAP/spoke-SDP/mesh-SDP.</li> <li>• discard-frame The packet will be discarded and an alarm generated. The frequency of alarm generation is fixed to be at most one alarm per MAC address per FP2 per 10 minutes in a given VPLS service. This parameter is only applicable to automatically protected MAC addresses.</li> </ul> <p>When the <b>restrict-protected-src</b> is enabled on an SHG the action only applies to the associated SAPs (no action is taken by default for spoke SDPs in the SHG). In order to enable this function for spoke SDPs within a SHG, the <b>restrict-protected-src</b> must be enabled explicitly under the spoke-SDP. If required, <b>restrict-protected-src</b> can also be enabled explicitly under specific SAPs within the SHG.</p> <p>When this command is applied or removed, with either the alarm-only or discard-frame parameters, the MAC addresses are cleared from the related object.</p> <p>The use of “<b>restrict-protected-src discard-frame</b>” is mutually exclusive with both the “<b>restrict-protected-src [alarm-only]</b>” command and with the configuration of manually protected MAC addresses within a given VPLS. “restrict-protected-src discard-frame” can only be enabled on SAPs on FP2 or later hardware or on SDPs where all network interfaces are on FP2 or later hardware.</p>
<b>Parameters</b>	<p><i>alarm-only</i> — Specifies that the packet will be forwarded, an alarm will be generated but the source MAC is not learned on the SAP/spoke-SDP/mesh-SDP.</p> <p><b>Default</b> no alarm-only</p>

*discard-frame* — Specifies that the packet will be discarded and an alarm generated. The frequency of alarm generation is fixed to be at most one alarm per FP2 per MAC address per 10 minutes within a given VPLS service.

**Default** no discard-frame

**Default** no restrict-protected-src

### restrict-unprotected-dst

<b>Syntax</b>	<b>restrict-unprotected-dst</b> <b>no restrict-unprotected-dst</b>
<b>Context</b>	config>service>pw-template>split-horizon-group config>service>vpls>split-horizon-group config>service>vpls>sap
<b>Description</b>	<p>This command indicates how the system will forward packets destined to an unprotected MAC address, either manually added using the <code>mac-protect</code> command or automatically added using the <code>auto-learn-mac-protect</code> command. While enabled all packets entering the configured SAP or SAPs within a split-horizon-group (but not spoke or mesh-SDPs) will be verified to contain a protected destination MAC address. If the packet is found to contain a non-protected destination MAC, it will be discarded. Detecting a non-protected destination MAC on the SAP will not cause the SAP to be placed in the operationally down state. No alarms are generated.</p> <p>If the destination MAC address is unknown, even if the packet is entering a restricted SAP, with <code>restrict-unprotected-dst</code> enabled, it will be flooded.</p>
<b>Default</b>	no restrict-unprotected-dst

### mac-pinning

<b>Syntax</b>	<b>[no] mac-pinning</b>
<b>Context</b>	config>service>vpls>sap config>service>vpls>spoke-sdp config>service>vpls>mesh-sdp
<b>Description</b>	<p>Enabling this command will disable re-learning of MAC addresses on other SAPs within the VPLS. The MAC address will remain attached to a given SAP for duration of its age-timer.</p> <p>The age of the MAC address entry in the FIB is set by the age timer. If <b>mac-aging</b> is disabled on a given VPLS service, any MAC address learned on a SAP/SDP with <b>mac-pinning</b> enabled will remain in the FIB on this SAP/SDP forever.</p> <p>Every event that would otherwise result in re-learning will be logged (MAC address; original-SAP; new-SAP).</p> <p>Note that MAC addresses learned during DHCP address assignment (DHCP snooping enabled) are not impacted by this command. MAC-pinning for such addresses is implicit.</p>
<b>Default</b>	When a SAP or spoke SDP is part of a Residential Split Horizon Group (RSHG), MAC pinning is activated at creation of the SAP. Otherwise MAC pinning is not enabled by default.



## ARP Handling Commands

### arp-reply-agent

<b>Syntax</b>	<b>arp-reply-agent [sub-ident]</b> <b>no arp-reply-agent</b>
<b>Context</b>	config>service>vpls>sap
<b>Description</b>	<p>This command enables a special ARP response mechanism in the system for ARP requests destined to static or dynamic hosts associated with the SAP. The system responds to each ARP request using the host's MAC address as the both the source MAC address in the Ethernet header and the target hardware address in the ARP header.</p> <p>ARP replies and requests received on a SAP with <b>arp-reply-agent</b> enabled will be evaluated by the system against the anti-spoof filter entries associated with the ingress SAP (if the SAP has anti-spoof filtering enabled). ARPs from unknown hosts on the SAP will be discarded when anti-spoof filtering is enabled.</p> <p>The ARP reply agent only responds if the ARP request enters an interface (SAP, spoke-SDP or mesh-SDP) associated with the VPLS instance of the SAP.</p> <p>A received ARP request that is not in the ARP reply agent table is flooded to all forwarding interfaces of the VPLS capable of broadcast except the ingress interface while honoring split-horizon constraints.</p> <p>Static hosts can be defined on the SAP using the <b>host</b> command. Dynamic hosts are enabled on the system by enabling the <b>lease-populate</b> command in the SAP's <b>dhcp</b> context. In the event that both a static host and a dynamic host share the same IP and MAC address, the VPLS ARP reply agent will retain the host information until both the static and dynamic information are removed. In the event that both a static and dynamic host share the same IP address, but different MAC addresses, the VPLS ARP reply agent is populated with the static host information.</p> <p>The <b>arp-reply-agent</b> command will fail if an existing static host on the SAP does not have both MAC and IP addresses specified. Once the ARP reply agent is enabled, creating a static host on the SAP without both an IP address and MAC address will fail.</p> <p>The ARP-reply-agent may only be enabled on SAPs supporting Ethernet encapsulation.</p> <p>The <b>no</b> form of the command disables ARP-reply-agent functions for static and dynamic hosts on the SAP.</p>
<b>Default</b>	not enabled
<b>Parameters</b>	<p><b>sub-ident</b> — Configures the arp-reply-agent to discard ARP requests received on the SAP that are targeted for a known host on the same SAP with the same subscriber identification.</p> <p>Hosts are identified by their subscriber information. For DHCP subscriber hosts, the subscriber hosts, the subscriber information is configured using the optional subscriber parameter string.</p> <p>When arp-reply-agent is enabled with <b>sub-ident</b>:</p> <ul style="list-style-type: none"> <li>• If the subscriber information for the destination host exactly matches the subscriber information for the originating host and the destination host is known on the same SAP as the source, the ARP request is silently discarded.</li> </ul>

- If the subscriber information for the destination host or originating host is unknown or undefined, the source and destination hosts are not considered to be the same subscriber. The ARP request is forwarded outside the SAP's Split Horizon Group.
- When **sub-ident** is not configured, the arp-reply-agent does not attempt to identify the subscriber information for the destination or originating host and will not discard an ARP request based on subscriber information.

### arp-populate

<b>Syntax</b>	<b>[no] arp-populate</b>
<b>Context</b>	config>service>ies>interface
<b>Description</b>	<p>This command, when enabled, disables dynamic learning of ARP entries. Instead, the ARP table is populated with dynamic entries from the DHCP lease state table (enabled with <b>lease-populate</b>), and optionally with static entries entered with the <b>host</b> command.</p> <p>Enabling the <b>arp-populate</b> command will remove any dynamic ARP entries learned on this interface from the ARP cache.</p> <p>The <b>arp-populate</b> command will fail if an existing static ARP entry exists for this interface.</p> <p>The <b>arp-populate</b> command will fail if an existing static subscriber host on the SAP does not have both MAC and IP addresses specified.</p> <p>Once <b>arp-populate</b> is enabled, creating a static subscriber host on the SAP without both an IP address and MAC address will fail.</p> <p>When <b>arp-populate</b> is enabled, the system will not send out ARP requests for hosts that are not in the ARP cache. Only statically configured and DHCP learned hosts are reachable through an IP interface with <b>arp-populate</b> enabled. The <b>arp-populate</b> command can only be enabled on IES interfaces supporting Ethernet encapsulation.</p> <p>Use the <b>no</b> form of the command to disable ARP cache population functions for static and dynamic hosts on the interface. All static and dynamic host information for this interface will be removed from the system's ARP cache.</p>
<b>Default</b>	not enabled

### arp-timeout

<b>Syntax</b>	<b>arp-timeout <i>seconds</i></b> <b>no arp-timeout</b>
<b>Context</b>	config>service>ies>interface
<b>Description</b>	<p>This command configures the minimum time in seconds an ARP entry learned on the IP interface will be stored in the ARP table. ARP entries are automatically refreshed when an ARP request or gratuitous ARP is seen from an IP host, otherwise, the ARP entry is aged from the ARP table. If <b>arp-timeout</b> is set to a value of zero seconds, ARP aging is disabled.</p>

When the **arp-populate** and **lease-populate** commands are enabled on an IES interface, the ARP table entries will no longer be dynamically learned, but instead by snooping DHCP ACK message from a DHCP server. In this case the configured **arp-timeout** value has no effect.

The default value for **arp-timeout** is 14400 seconds (4 hours).

The **no** form of this command restores **arp-timeout** to the default value.

<b>Default</b>	14400 seconds
<b>Parameters</b>	<i>seconds</i> — The minimum number of seconds a learned ARP entry will be stored in the ARP table, expressed as a decimal integer. A value of zero specifies that the timer is inoperative and learned ARP entries will not be aged.
<b>Values</b>	0 — 65535

## authentication-policy

<b>Syntax</b>	<b>authentication-policy</b> <i>name</i> <b>no authentication-policy</b>
<b>Context</b>	config>service>ies>sub-if>grp-if
<b>Description</b>	This command assigns a RADIUS authentication policy to the interface. The <b>no</b> form of this command removes the policy name from the group interface configuration.
<b>Default</b>	no authentication-policy
<b>Parameters</b>	<i>name</i> — Specifies the authentication policy name.

## local-proxy-arp

<b>Syntax</b>	<b>[no] local-proxy-arp</b>
<b>Context</b>	config>service>ies>interface
<b>Description</b>	This command enables local proxy ARP. When local proxy ARP is enabled on an IP interface, the system responds to all ARP requests for IP addresses belonging to the subnet with its own MAC address, and thus will become the forwarding point for all traffic between hosts in that subnet.  When local-proxy-arp is enabled, ICMP redirects on the ports associated with the service are automatically blocked.
<b>Default</b>	no local-proxy-arp

## remote-proxy-arp

<b>Syntax</b>	<b>[no] remote-proxy-arp</b>
<b>Context</b>	config>service>ies>interface

## ARP Handling Commands

<b>Description</b>	This command enables or disables remote proxy ARP on the interface.
<b>Default</b>	no remote-proxy-arp

## qos-route-lookup

<b>Syntax</b>	<b>[no] qos-route-lookup</b>
<b>Context</b>	config>service>ies>sub-if>grp-if>ipv6
<b>Description</b>	This command enables or disables Qos route lookup for the interface.

## proxy-arp-policy

<b>Syntax</b>	<b>proxy-arp-policy</b> <i>policy-name</i> [ <i>policy-name...</i> (up to 5 max)] <b>no proxy-arp-policy</b>
<b>Context</b>	config>service>ies>interface
<b>Description</b>	This command specifies an existing policy-statement to analyze match and action criteria that controls the flow of routing information to and from a given protocol, set of protocols, or a particular neighbor.
<b>Default</b>	none
<b>Parameters</b>	<i>policy-name</i> — The export route policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters excluding double quotes. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.  The specified name(s) must already be defined.

## static-arp

<b>Syntax</b>	<b>static-arp</b> <i>ip-address</i> <i>ieee-mac-address</i> <b>no static-arp</b> <i>ip-address</i> [ <i>ieee-mac-address</i> ]
<b>Context</b>	config>service>ies>interface
<b>Description</b>	This command configures a static address resolution protocol (ARP) entry associating a subscriber IP address with a MAC address for the core router instance. This static ARP appears in the core routing ARP table. A static ARP can only be configured if it exists on the network attached to the IP interface.  If an entry for a particular IP address already exists and a new MAC address is configured for the IP address, the existing MAC address will be replaced with the new MAC address.  The <b>no</b> form of the command removes a static ARP entry.
<b>Default</b>	None
<b>Parameters</b>	<i>ip-address</i> — Specifies the IP address for the static ARP in IP address dotted decimal notation.

*ieee-mac-address* — Specifies the 48-bit MAC address for the static ARP in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee and ff are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

## Show Commands

### arp

<b>Syntax</b>	<b>arp</b>
<b>Context</b>	show>service>id
<b>Description</b>	This command displays the ARP cache entries for this service.
<b>Output</b>	<b>Show All Service-ID Output</b> — The following table describes the show command output fields:

Label	Description
IP Address	Specifies the IP address of the ARP each entry.
MAC Address	Specifies the MAC address associated with the IP address.
Type	Other — Learned through normal ARP queries. Static — Configured by <b>static-arp</b> commands. Managed — Learned from DHCP snooping or configured by <b>host</b> commands.
Age	Indicates age of the ARP entry.
Interface	Indicates the name of the IP interface.
Port	Indicates the port that the entry was learned on.

### Sample Output

```
A:ALA-A# show service id 100 base
=====
ARP Table
=====
IP Address      MAC Address      Type    Age      Interface      Port
-----
101.1.0.1       00:00:66:66:66:01 Other    00h00m00s ies-100-101.1.0.1 1/1/4
200.1.1.2       00:00:5e:00:01:64 Other    00h00m00s ies-100-200.1.1.2 1/1/3
200.1.1.201     00:00:22:2e:a5:61 Static    00h00m00s ies-100-200.1.1.2 1/1/3
200.1.1.202     00:00:22:2e:a5:62 Static    00h00m00s ies-100-200.1.1.2 1/1/3
=====
A:ALA-A#
```

## Configuring Triple Play Multicast Services with CLI

This section provides information to configure multicast parameters in a Triple Play network using the command line interface.

Topics in this section include:

- [Configuring IGMP Snooping in the BSA on page 588](#)
  - [Enabling IGMP Snooping in a VPLS Service on page 588](#)
  - [Modifying IGMP Snooping Parameters on page 590](#)
  - [Configuring Static Multicast Groups on a SAP or SDP on page 593](#)
  - [Enabling IGMP Group Membership Report Filtering on page 594](#)
  - [Enabling IGMP Traffic Filtering on page 596](#)
  - [Configuring Multicast VPLS Registration \(MVR\) on page 597](#)

## Configuring IGMP Snooping in the BSA

- [Enabling IGMP Snooping in a VPLS Service on page 588](#)
  - [Configuring Static Multicast Groups on a SAP or SDP on page 593](#)
  - [Enabling IGMP Group Membership Report Filtering on page 594](#)
  - [Enabling IGMP Traffic Filtering on page 596](#)
  - [Configuring Multicast VPLS Registration \(MVR\) on page 597](#)
- 

### Enabling IGMP Snooping in a VPLS Service

- [With IGMPv3 Multicast Routers on page 588](#)
    - [With IGMPv3 Multicast Routers on page 588](#)
    - [With IGMPv1/2 Multicast Routers on page 589](#)
  - [Modifying IGMP Snooping Parameters on page 590](#)
  - [Modifying IGMP Snooping Parameters for a SAP or SDP on page 591](#)
- 

### With IGMPv3 Multicast Routers

When multicast routers use IGMPv3, it is sufficient to just enable IGMP snooping, without any further modification of parameters.

The following displays an example of an IGMP snooping configuration:

```
A:ALA-48>config>service>vpls# info
-----
      igmp-snooping
      no shutdown
      exit
      no shutdown
-----
A:ALA-48>config>service>vpls#
```



## With IGMPv1/2 Multicast Routers

When the multicast routers don't support IGMPv3, some timing parameters need to be configured locally in the Alcatel-Lucent ESS-Series. Note that all routers in the multicast network must use the same values for these parameters.

The following displays an example of a modified IGMP snooping configuration:

```
A:ALA-48>config>service>vpls# info
-----
      stp
        shutdown
      exit
      igmp-snooping
        query-interval 60
        robust-count 5
        no shutdown
      exit
      no shutdown
-----
A:ALA-48>config>service>vpls#
```

## Modifying IGMP Snooping Parameters

For interoperability with some multicast routers, the source IP address of IGMP group reports can be configured. Use the following CLI syntax to customize this IGMP snooping parameter:

The following displays an example of a modified IGMP snooping configuration:

```
A:ALA-48>config>service>vpls# info
-----
      stp
        shutdown
      exit
    igmp-snooping
      query-interval 60
      robust-count 5
      report-src-ip 10.20.20.20
      no shutdown
    exit
  no shutdown
-----
A:ALA-48>config>service>vpls#
```

## Modifying IGMP Snooping Parameters for a SAP or SDP

Use the following CLI syntax to customize IGMP snooping parameters on an existing SAP. Commands for spoke or mesh SDPs are identical.

**CLI Syntax:**

```
config>service# vpls service-id
    sap sap-id
        igmp-snooping
            fast-leave
            import policy-name
            last-member-query-interval interval
            max-num-groups max-num-groups
            mrouter-port
            query-interval interval
            query-response-interval interval
            robust-count count
            send-queries
```

To enable and customize sending of IGMP queries to the hosts:

**Example:**

```
config>service# vpls 1
config>service>vpls# sap 1/1/3:0
config>service>vpls>sap# igmp-snooping
config>service>vpls>sap>snooping# send-queries
config>service>vpls>sap>snooping# query-interval 100
config>service>vpls>sap>snooping# query-response-interval 60
config>service>vpls>sap>snooping# robust-count 5
config>service>vpls>sap>snooping# exit
config>service>vpls>sap# no shutdown
```

To customize the leave delay:

**Example:**

```
config>service# vpls 1
config>service>vpls# sap 1/1/1:1
config>service>vpls>sap# igmp-snooping
config>service>vpls>sap>snooping# last-member-query-interval 10
config>service>vpls>sap>snooping# no fast-leave
config>service>vpls>sap>snooping# exit
config>service>vpls>sap# exit
```

To enable Fast Leave:

**Example:** config>service# vpls 1  
config>service>vpls# sap 1/1/1:1  
config>service>vpls>sap# igmp-snooping  
config>service>vpls>sap>snooping# no last-member-query-interval  
config>service>vpls>sap>snooping# fast-leave  
config>service>vpls>sap>snooping# exit  
config>service>vpls>sap# exit

To limit the number of streams that a host can join:

**Example:** config>service# vpls 1  
config>service>vpls# sap 1/1/1:1  
config>service>vpls>sap# igmp-snooping  
config>service>vpls>sap>snooping# max-num-groups 4  
config>service>vpls>sap>snooping# exit  
config>service>vpls>sap# exit

To enable sending group reports on a SAP to standby multicast routers:

**Example:** config>service# vpls 1  
config>service>vpls# sap 1/1/1:1  
config>service>vpls>sap# igmp-snooping  
config>service>vpls>sap>snooping# mrouter-port  
config>service>vpls>sap>snooping# exit  
config>service>vpls>sap# exit

The following example displays the modified IGMP snooping configuration on a SAP:

```
A:ALA-48>config>service>vpls>sap>snooping# info detail
-----
no fast-leave
no import
max-num-groups 4
last-member-query-interval 10
no mrouter-port
query-interval 100
query-response-interval 60
robust-count 5
send-queries
-----
A:ALA-48>config>service>vpls>sap>snooping#
```

## Configuring Static Multicast Groups on a SAP or SDP

Use the following CLI syntax to add static group membership entries on an existing SAP (commands for spoke or mesh SDPs are identical):

The following displays an example of a static IGMP snooping configuration on a SAP:

```
A:ALA-48>config>service>vpls>sap# info
-----
      max-nbr-mac-addr 4
      igmp-snooping
        fast-leave
        mrouter-port
        static
          group 224.0.10.10
            source 10.10.10.1
            source 10.10.10.2
          exit
        exit
      exit
-----
A:ALA-48>config>service>vpls>sap#
```

## Enabling IGMP Group Membership Report Filtering

Routing policies can be defined to limit the multicast channels that can be joined by a host. For example, it is possible to define a policy listing a group of multicast streams (for example, 'basic' containing a basic set of TV channels or 'extended' containing a more extended set of TV channels), and to apply this policy to subscribers of IGMP snooping (SAPs and/or SDPs).

The following displays an example of a configuration to import a routing policy on a SAP:

```
A:ALA-48>config>service>vpls# info
-----
      stp
        shutdown
      exit
      igmp-snooping
        query-interval 60
        robust-count 5
        report-src-ip 10.20.20.20
        no shutdown
      exit
      sap 1/1/3:0 create
        igmp-snooping
          query-interval 100
          query-response-interval 60
          robust-count 5
          send-queries
        exit
      exit
      sap 1/1/3:22 create
        max-nbr-mac-addr 4
        igmp-snooping
          fast-leave
          import "test_policy"
          mrouter-port
          static
            group 224.0.10.10
            source 10.10.10.1
            source 10.10.10.2
          exit
        exit
      exit
      exit
      no shutdown
-----
A:ALA-48>config>service>vpls#
```

For details configuring a routing policy, see the Configuring Route Policies section in the 7450 ESS Router Configuration Guide.

The following shows a sample routing policy configuration accepting IGMP messages for only five multicast channels:

```
A:ALA-48>config>router>policy-options# info
-----
    prefix-list "basic_channels"
      prefix 224.10.0.1/32 exact
      prefix 224.10.0.2/32 exact
      prefix 224.10.0.3/32 exact
      prefix 224.10.0.4/32 exact
      prefix 224.10.0.5/32 exact
    exit
  policy-statement "test_policy"
    description "basic set of 5 multicast channels"
    entry 1
      from
        group-address "basic_channels"
      exit
      action accept
      exit
    exit
    default-action reject
  exit
-----
A:ALA-48>config>router>policy-options#
```

### Enabling IGMP Traffic Filtering

For security, it might be advisable to only allow multicast traffic into the ESS-Series from recognized multicast routers and servers. Multicast packets arriving on other interfaces (for example, customer-facing SAPs or spoke SDPs) can be filtered out by defining an appropriate IP filter policy.

For details on how to configure a filter policy, see section Creating an IP Filter Policy in the 7450 ESS Router Configuration Guide.

The following example shows a sample IP filter policy configuration dropping all multicast traffic:

```
A:ALA-48>config>filter>ip-filter# info
-----
ip-filter 1 create
  entry 1 create
    match
      dst-ip 224.0.0.0/24
    exit
    action accept
  exit
  entry 2 create
    match
      dst-ip 224.0.0.0/4
    exit
    action drop
  exit
exit
-----
A:ALA-48>config>filter>ip-filter#
```

The following example shows how to apply this sample IP filter policy to a SAP:

```
A:ALA-48>config>service>vpls # info
-----
sap 1/1/1:1
  ingress
    filter ip 1
  exit
exit
-----
A:ALA-48>config>service>vpls>snooping#
```



## Configuring Multicast VPLS Registration (MVR)

Use the following CLI syntax to configure Multicast VPLS Registration. The first step is to register a VPLS as a multicast VPLS.

**CLI Syntax:** config>service# vpls *service-id*  
                   igmp-snooping  
                   mvr  
                   no shutdown  
                   description *description*  
                   group-policy *policy-name*

**Example:** config>service# vpls 1000  
               config>service>vpls# igmp-snooping  
               config>service>vpls>snooping# mvr  
               config>service>vpls>snooping>mvr# no shutdown  
               config>service>vpls>snooping>mvr# description "MVR VPLS"  
               config>service>vpls>snooping>mvr# group-policy "basic\_channels\_policy"

The second step is to configure a SAP to take the multicast channels from the registered multicast VPLS.

**CLI Syntax:** config>service# vpls *service-id*  
                   sap *sap-id*  
                   igmp-snooping  
                   mvr  
                   from-vpls *vpls-id*

**Example:** config>service# vpls 1  
               config>service>vpls# sap 1/1/1:100  
               config>service>vpls>sap# igmp-snooping  
               config>service>vpls>snooping# mvr  
               config>service>vpls>snooping>mvr# from-vpls 1000

For MVR by proxy also the destination SAP for the multicast channels should be configured.

**CLI Syntax:** config>service# vpls *service-id*  
                   sap *sap-id*  
                   igmp-snooping  
                   mvr  
                   from-vpls *vpls-id*  
                   to-sap *sap-id*

**Example:** config>service# vpls 1  
               config>service>vpls# sap 1/1/1:100  
               config>service>vpls>sap# igmp-snooping  
               config>service>vpls>snooping# mvr  
               config>service>vpls>snooping>mvr# from-vpls 1000  
               config>service>vpls>snooping>mvr# to-sap 1/1/1:200



# Triple Play Multicast

---

## In This Chapter

This chapter provides information about Triple Play Multicast aspects, including configuration process overview, and implementation notes.

Topics in this chapter include:

- [Introduction to Multicast on page 601](#)
- [Multicast in the Broadband Service Router on page 602](#)
  - [Internet Group Management Protocol on page 602](#)
  - [Multicast Listener Discovery on page 604](#)
  - [Source Specific Multicast Groups on page 605](#)
  - [Protocol Independent Multicast Sparse Mode \(PIM-SM\) on page 606](#)
- [Multicast in the BSA on page 608](#)
  - [IGMP Snooping on page 608](#)
- [Multicast Support over Subscriber Interfaces in Routed CO Model on page 621](#)
  - [Hardware Support on page 623](#)
  - [Multicast Over IPoE on page 624](#)
  - [Multicast Over PPPoE on page 638](#)
  - [IGMP Flooding Containment on page 639](#)
  - [IGMP/MLD Timers on page 639](#)
  - [IGMP/MLD Query Intervals on page 640](#)
  - [HQoS Adjustment on page 640](#)
  - [Redirection on page 650](#)
  - [Hierarchical Multicast CAC \(H-MCAC\) on page 652](#)
  - [Determining MCAC Policy in Effect on page 660](#)
  - [Multicast Filtering on page 661](#)

- [Joining the Multicast Tree on page 662](#)
- [Wholesale/Retail Requirements on page 662](#)
- [QoS Considerations on page 664](#)
- [Redundancy Considerations on page 664](#)
- [Configuring Triple Play Multicast Services with CLI on page 587](#)

## Introduction to Multicast

IP multicast provides an effective method of many-to-many communication. Delivering unicast datagrams is fairly simple. Normally, IP packets are sent from a single source to a single recipient. The source inserts the address of the target host in the IP header destination field of an IP datagram, intermediate routers (if present) simply forward the datagram towards the target in accordance with their respective routing tables.

Sometimes distribution needs individual IP packets be delivered to multiple destinations (like audio or video streaming broadcasts). Multicast is a method of distributing datagrams sourced from one (or possibly more) host(s) to a set of receivers that may be distributed over different (sub) networks. This makes delivery of multicast datagrams significantly more complex.

Multicast sources can send a single copy of data using a single address for the entire group of recipients. The routers between the source and recipients route the data using the group address route. Multicast packets are delivered to a multicast group. A multicast group specifies a set of recipients who are interested in a particular data stream and is represented by an IP address from a specified range. Data addressed to the IP address is forwarded to the members of the group. A source host sends data to a multicast group by specifying the multicast group address in datagram's destination IP address. A source does not have to register in order to send data to a group nor do they need to be a member of the group.

Routers and Layer 3 switches use the Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) to manage membership for a multicast session. When a host wants to receive one or more multicast sessions, it will send a join message for each multicast group it wants to join. When a host wants to leave a multicast group, it will send a leave message.

## Multicast in the Broadband Service Router

This section describes the multicast protocols employed when an Alcatel-Lucent router is used as a Broadband Service Router (BSR) in a Triple Play aggregation network.

The protocols used are:

- Internet Group Management Protocol ([Internet Group Management Protocol on page 602](#))
  - Multicast Listener Discovery ([Multicast Listener Discovery on page 604](#))
  - Source Specific Multicast Groups ([Internet Group Management Protocol on page 602](#))
  - Protocol Independent Multicast (Sparse Mode) ([PIM-SM on page 606](#))
- 

### Internet Group Management Protocol

Internet Group Management Protocol (IGMP) is used by IPv4 hosts and routers to report their IP multicast group memberships to neighboring multicast routers. A multicast router keeps a list of multicast group memberships for each attached network, and a timer for each membership.

Multicast group memberships include at least one member of a multicast group on a given attached network, not a list of all of the members. With respect to each of its attached networks, a multicast router can assume one of two roles, querier or non-querier. There is normally only one querier per physical network.

A querier issues two types of queries, a general query and a group-specific query. General queries are issued to solicit membership information with regard to any multicast group. Group-specific queries are issued when a router receives a leave message from the node it perceives as the last group member remaining on that network segment.

Hosts wanting to receive a multicast session issue a multicast group membership report. These reports must be sent to all multicast enabled routers.

---

### IGMP Versions and Interoperability Requirements

If routers run different versions of IGMP, they will negotiate the lowest common version of IGMP that is supported on their subnet and operate in that version.

Version 1 — Specified in RFC-1112, *Host extensions for IP Multicasting*, was the first widely deployed version and the first version to become an Internet standard.

Version 2 — Specified in RFC-2236, *Internet Group Management Protocol*, added support for low leave latency, that is, a reduction in the time it takes for a multicast router to learn that there are no longer any members of a particular group present on an attached network.

Version 3 —Specified in RFC-3376, *Internet Group Management Protocol*, adds support for source filtering, that is, the ability for a system to report interest in receiving packets only from specific source addresses, as required to support Source-Specific Multicast (See Source Specific Multicast (SSM)), or from all but specific source addresses, sent to a particular multicast address.

IGMPv3 must keep state per group per attached network. This group state consists of a filter-mode, a list of sources, and various timers. For each attached network running IGMP, a multicast router records the desired reception state for that network.

---

## IGMP Version Transition

Alcatel-Lucent's SRs are capable of interoperating with routers and hosts running IGMPv1, IGMPv2, and/or IGMPv3. *Draft-ietf-magma-igmpv3-and-routing-0x.txt* explores some of the interoperability issues and how they affect the various routing protocols.

IGMP version 3 specifies that if at any point a router receives an older version query message on an interface that it must immediately switch into a compatibility mode with that earlier version. Since none of the previous versions of IGMP are source aware, should this occur and the interface switch to Version 1 or 2 compatibility mode, any previously learned group memberships with specific sources (learned from the IGMPv3 specific INCLUDE or EXCLUDE mechanisms) MUST be converted to non-source specific group memberships. The routing protocol will then treat this as if there is no EXCLUDE definition present.

## Multicast Listener Discovery

Multicast Listener Discovery (MLD) is used by IPv6 hosts and routers to report their IP multicast group memberships to neighboring multicast routers. A multicast router keeps a list of multicast group memberships for each attached network, and a timer for each membership. Multicast group memberships include at least one member of a multicast group on a given attached network, not a list of all of the members. With respect to each of its attached networks, a multicast router can assume one of two roles, querier or non-querier. There is normally only one querier per physical network.

A querier issues two types of queries, a general query and a group-specific query. General queries are issued to solicit membership information with regard to any multicast group. Group-specific queries are issued when a router receives a leave message from the node it perceives as the last group member remaining on that network segment.

Hosts wanting to receive a multicast session issue a multicast group membership report. These reports must be sent to all multicast enabled routers.

---

## MLD Versions and Interoperability Requirements

If routers run different versions of MLD, they will negotiate the lowest common version of MLD that is supported on their subnet and operate in that version.

Version 1 — Specified in RFC 2710, *Multicast Listener Discovery (MLD) for IPv6*, was the first deployed version and included low leave latency, that is, a reduction in the time it takes for a multicast router to learn that there are no longer any members of a particular group present on an attached network.

Version 2 — Specified in RFC 3810, *Multicast Listener Discovery Version 2 (MLDv2) for IPv6*, adds support for source filtering, that is, the ability for a system to report interest in receiving packets only from specific source addresses, as required to support Source-Specific Multicast.

Multicast (SSM)), or from all but specific source addresses, sent to a particular multicast address. MLDv2 must keep state per group per attached network. This group state consists of a filter mode, a list of sources, and various timers. For each attached network running MLD, a multicast router records the desired reception state for that network.



## Source Specific Multicast Groups

IGMPv3 and MLDv2 permits a receiver to join a group and specify that it only wants to receive traffic for a group if that traffic comes from a particular source. If a receiver does this, and no other receiver on the LAN requires all the traffic for the group, then the Designated Router (DR) can omit performing a (\*,G) join to set up the shared tree, and instead issue a source-specific (S,G) join only.

For IPv4, the range of multicast addresses from 232.0.0.0 to 232.255.255.255 is currently set aside for source-specific multicast. For groups in this range, receivers should only issue source specific IGMPv3 joins. If a PIM router receives a non-source-specific join for a group in this range, it should ignore it.

For IPv6, the multicast prefix FF3x::/32 is currently set aside for source-specific multicast. For groups in this range, receivers should only issue source specific MLDv2 joins. If a PIM router receives a non-source-specific join for a group in this range, it should ignore it.

An Alcatel-Lucent PIM router must silently ignore a received (\*, G) PIM join message where G is a multicast group address from the multicast address group range that has been explicitly configured for SSM. This occurrence should generate an event. If configured, the IGMPv2 (MLDv1 for IPv6) request can be translated into IGMPv3 (MLDv2 for IPv6). The SR allows for the conversion of an IGMPv2 (\*,G) request into a IGMPv3 (S,G) request based on manual entries. A maximum of 32 SSM ranges is supported.

IGMPv3 and MLDv2 also permits a receiver to join a group and specify that it only wants to receive traffic for a group if that traffic does not come from a specific source or sources. In this case, the DR will perform a (\*,G) join as normal, but can combine this with a prune for each of the sources the receiver does not wish to receive.

## Protocol Independent Multicast Sparse Mode (PIM-SM)

PIM-SM leverages the unicast routing protocols that are used to create the unicast routing table: OSPF, IS-IS, BGP, and static routes. Because PIM uses this unicast routing information to perform the multicast forwarding function it is effectively IP protocol independent. Unlike DVMRP, PIM does not send multicast routing tables updates to its neighbors.

PIM-SM uses the unicast routing table to perform the Reverse Path Forwarding (RPF) check function instead of building up a completely independent multicast routing table.

PIM-SM only forwards data to network segments with active receivers that have explicitly requested the multicast group. PIM-SM in the ASM model initially uses a shared tree to distribute information about active sources. Depending on the configuration options, the traffic can remain on the shared tree or switch over to an optimized source distribution tree. As multicast traffic starts to flow down the shared tree, routers along the path determine if there is a better path to the source. If a more direct path exists, then the router closest to the receiver sends a join message toward the source and then reroutes the traffic along this path.

As stated above, PIM-SM relies on an underlying topology-gathering protocol to populate a routing table with routes. This routing table is called the Multicast Routing Information Base (MRIB). The routes in this table can be taken directly from the unicast routing table, or it can be different and provided by a separate routing protocol such as MBGP. Regardless of how it is created, the primary role of the MRIB in the PIM-SM protocol is to provide the next hop router along a multicast-capable path to each destination subnet. The MRIB is used to determine the next hop neighbor to whom any PIM join/prune message is sent. Data flows along the reverse path of the join messages. Thus, in contrast to the unicast RIB that specifies the next hop that a data packet would take to get to some subnet, the MRIB gives reverse-path information, and indicates the path that a multicast data packet would take from its origin subnet to the router that has the MRIB.

## Ingress Multicast Path Management (IMPM) Enhancements

Refer to the SR OS Advanced Configuration Guide for more information on IMPM as well as detailed configuration examples.

Ingress multicast path management (IMPM) allows the system to dynamically manage Layer 2 and Layer 3 IP multicast flows into the available multicast paths through the switch fabric. The ingress multicast manager understands the amount of available multicast bandwidth per path and the amount of bandwidth used per IP multicast stream.

Two policies define how each path should be managed, the bandwidth policy, and how multicast channels compete for the available bandwidth, the multicast information policy.

Chassis multicast planes should not be confused with IOM/IMM multicast paths. The IOM/IMM uses multicast paths to reach multicast planes on the switch fabric. An IOM/IMM may have less or more multicast paths than the number of multicast planes available in the chassis.

Each IOM/IMM multicast path is either a primary or secondary path type. The path type indicates the multicast scheduling priority within the switch fabric. Multicast flows sent on primary paths are scheduled at multicast high priority while secondary paths are associated with multicast low priority.

The system determines the number of primary and secondary paths from each IOM/IMM forwarding plane and distributes them as equally as possible between the available switch fabric multicast planes. Each multicast plane may terminate multiple paths of both the primary and secondary types.

The system ingress multicast management module evaluates the ingress multicast flows from each ingress forwarding plane and determines the best multicast path for the flow. A particular path may be used until the terminating multicast plane is “maxed” out (based on the rate limit defined in the **per-mcast-plane-capacity** commands) at which time either flows are moved to other paths or potentially blackholed (flows with the lowest preference are dropped first). In this way, the system makes the best use of the available multicast capacity without congesting individual multicast planes.

The switch fabric is simultaneously handling both unicast and multicast flows. The switch fabric uses a weighted scheduling scheme between multicast high, unicast high, multicast low and unicast low when deciding which cell to forward to the egress forwarding plane next. The weighted mechanism allows some amount of unicast and lower priority multicast (secondary) to drain on the egress switch fabric links used by each multicast plane. The amount is variable based on the number of switch fabric planes available on the amount of traffic attempting to use the fabric planes. The **per-mcast-plane-capacity** commands allows the amount of managed multicast traffic to be tuned to compensate for the expected available egress multicast bandwidth per multicast plane. In conditions where it is highly desirable to prevent multicast plane congestion, the **per-mcast-plane-capacity** commands should be used to compensate for the non-multicast or secondary multicast switch fabric traffic.

## Multicast in the BSA

IP Multicast is normally not a function of the Broadband Service Aggregator (BSA) in a Triple Play aggregation network being a Layer 2 device. However, the BSA does use IGMP snooping to optimize bandwidth utilization.

---

### IGMP Snooping

For most Layer 2 switches, multicast traffic is treated like an unknown MAC address or broadcast frame, which causes the incoming frame to be flooded out (broadcast) on every port within a VLAN. While this is acceptable behavior for unknowns and broadcasts, as IP Multicast hosts may join and be interested in only specific multicast groups, all this flooded traffic results in wasted bandwidth on network segments and end stations.

IGMP snooping entails using information in layer 3 protocol headers of multicast control messages to determine the processing at layer 2. By doing so, an IGMP snooping switch provides the benefit of conserving bandwidth on those segments of the network where no node has expressed interest in receiving packets addressed to the group address.

On the Alcatel-Lucent 7450 ESS, IGMP snooping can be enabled in the context of VPLS services. The IGMP snooping feature allows for optimization of the multicast data flow for a group within a service to only those Service Access Points (SAPs) and Service Distribution Points (SDPs) that are members of the group. In fact, the Alcatel-Lucent 7450 ESS implementation performs more than pure snooping of IGMP data, since it also summarizes upstream IGMP reports and responds to downstream queries.

The Alcatel-Lucent 7450 ESS maintains a number of multicast databases:

- A port database on each SAP and SDP lists the multicast groups that are active on this SAP or SDP.
- All port databases are compiled into a central proxy database. Towards the multicast routers, summarized group membership reports are sent based on the information in the proxy database.
- The information in the different port databases is also used to compile the multicast forwarding information base (MFIB). This contains the active SAPs and SDPs for every combination of source router and group address (S,G), and is used for the actual multicast replication and forwarding.

When the router receives a join report from a host for a particular multicast group, it adds the group to the port database and (if it is a new group) to the proxy database. It also adds the SAP or SDP to existing (S,G) in the MFIB, or builds a new MFIB entry.

When the router receives a leave report from a host, it first checks if other devices on the SAP or SDP still want to receive the group (unless fast leave is enabled). Then it removes the group from

the port database, and from the proxy database if it was the only receiver of the group. The router also deletes entries if it does not receive periodic membership confirmations from the hosts.

The fast leave feature finds its use in multicast TV delivery systems, for example. Fast Leave speeds up the membership leave process by terminating the multicast session immediately, rather than the standard procedure of issuing a group specific query to check if other group members are present on the SAP or SDP.

## IGMP/MLD Message Processing

Figure 27 illustrates the basic IGMP message processing by the 7450 ESS in several situations.

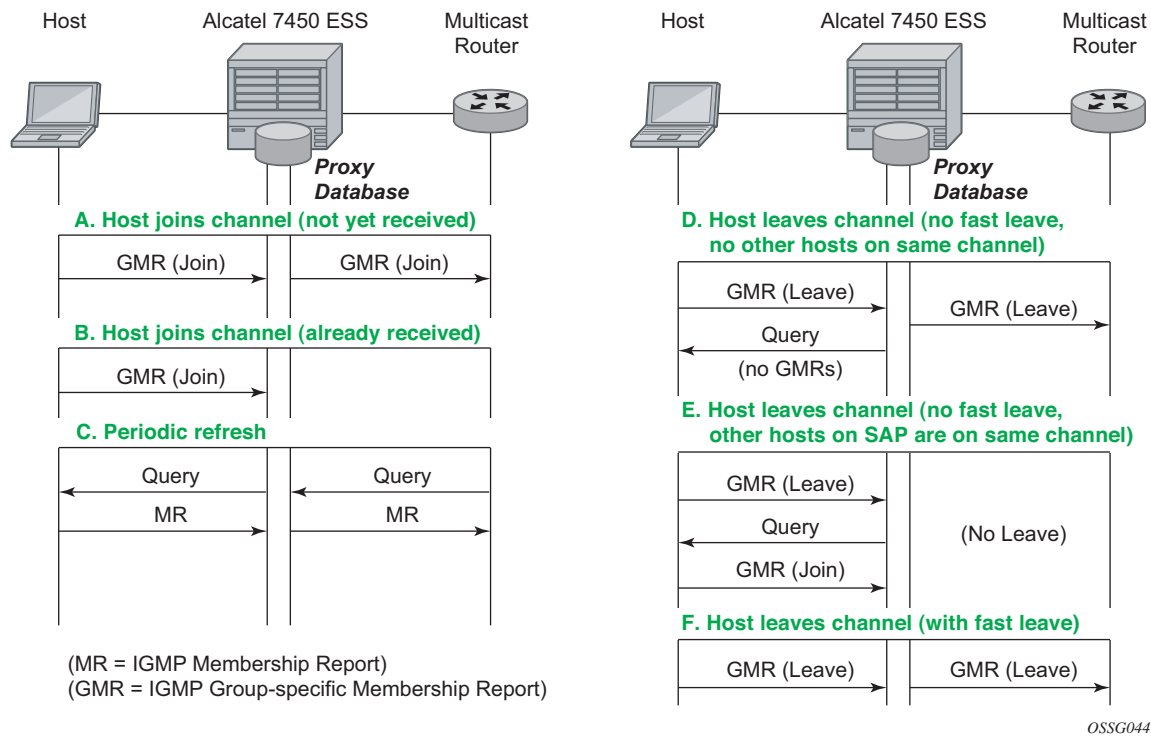


Figure 27: IGMP Message Processing

## IGMP Message Processing

Scenario A: A host joins a multicast group (TV channel) which is not yet being received by other hosts on the router, and thus is not yet present in the proxy database. The 7450 ESS adds the group

to the proxy database and sends a new IGMP Join group-specific membership report upstream to the multicast router.

Scenario B: A host joins a channel which is already being received by one or more hosts on the 7450 ESS, and thus is already present in the proxy database. No upstream IGMP report is generated by the router.

Scenario C: The multicast router will periodically send IGMP queries to the router, requesting it to respond with generic membership reports. Upon receiving such a query, the 7450 ESS will compile a report from its proxy database and send it back to the multicast router.

In addition, the router will flood the received IGMP query to all hosts (on SAPs and spoke SDPs), and will update its proxy database based on the membership reports received back.

Scenario D: A host leaves a channel by sending an IGMP leave message. If fast-leave is not enabled, the router will first check whether there are other hosts on the same SAP or spoke SDP by sending a query. If no other host responds, the 7450 ESS removes the channel from the SAP. In addition, if there are no other SAPs or spoke SDPs with hosts subscribing to the same channel, the channel is removed from the proxy database and an IGMP leave report is sent to the upstream Multicast Router.

Scenario E: A host leaves a channel by sending an IGMP leave message. If fast-leave is not enabled, the router will check whether there are other hosts on the same SAP or spoke SDP by sending a query. Another device on the same SAP or spoke SDP still wishes to receive the channel and responds with a membership report. Thus the 7450 ESS does not remove the channel from the SAP.

Scenario F: A host leaves a channel by sending an IGMP leave report. Fast-leave is enabled, so the 7450 ESS will not check whether there are other hosts on the same SAP or spoke SDP but immediately removes the group from the SAP. In addition, if there are no other SAPs or spoke SDPs with hosts subscribing to the same group, the group is removed from the proxy database and an IGMP leave report is sent to the upstream multicast router.

---

## MLD Message Processing

MLD message processing differs from IGMP. An IPv6 host can have two WAN IPv6 addresses and an IPv6 prefix. MLD messages source address are link local addresses. This makes it difficult to know if the originating host is a WAN host or a PD host. By default, all requested IPv6 (s,g) are first associated with a WAN host. If this particular WAN host disconnects or ends its IPv6 session, the (s,g) is then associated with the remaining WAN host. If there are no more WAN hosts, the (s,g) is then associated with the remaining PD host. The (s,g) is always transferred to the remaining IPv6 host until there are no more report replies to corresponding to the queries. Scenarios A — F will not differ for IPv6 hosts.

## IGMP/MLD Filtering

A provider may want to block receive or transmit permission to individual hosts or a range of hosts. To this end, the Alcatel-Lucent 7450 ESS supports IGMP/MLD filtering. Two types of filter can be defined:

- Filter IGMP/MLD membership reports from a particular host or range of hosts. This is performed by importing an appropriately defined routing policy into the SAP or spoke SDP.
- Filter to prevent a host from transmitting multicast streams into the network. The operator can define a data-plane filter (ACL) which drops all multicast traffic, and apply this filter to a SAP or spoke SDP.

## Multicast VPLS Registration (MVR)

Multicast VPLS Registration (MVR) is a bandwidth optimization method for multicast in a broadband services network. MVR allows a subscriber on a port to subscribe and unsubscribe to a multicast stream on one or more network-wide multicast VPLS instances.

MVR assumes that subscribers join and leave multicast streams by sending IGMP join and leave messages. The IGMP leave and join message are sent inside the VPLS to which the subscriber port is assigned. The multicast VPLS is shared in the network while the subscribers remain in separate VPLS services. Using MVR, users on different VPLS cannot exchange any information between them, but still multicast services are provided.

On the MVR VPLS, IGMP snooping must be enabled. On the user VPLS, IGMP snooping and MVR work independently. If IGMP snooping and MVR are both enabled, MVR reacts only to join and leave messages from multicast groups configured under MVR. Join and leave messages from all other multicast groups are managed by IGMP snooping in the local VPLS. This way, potentially several MVR VPLS instances could be configured, each with its own set of multicast channels.

MVR by proxy — In some situations, the multicast traffic should not be copied from the MVR VPLS to the SAP on which the IGMP message was received (standard MVR behavior) but to another SAP. This is called MVR by proxy.

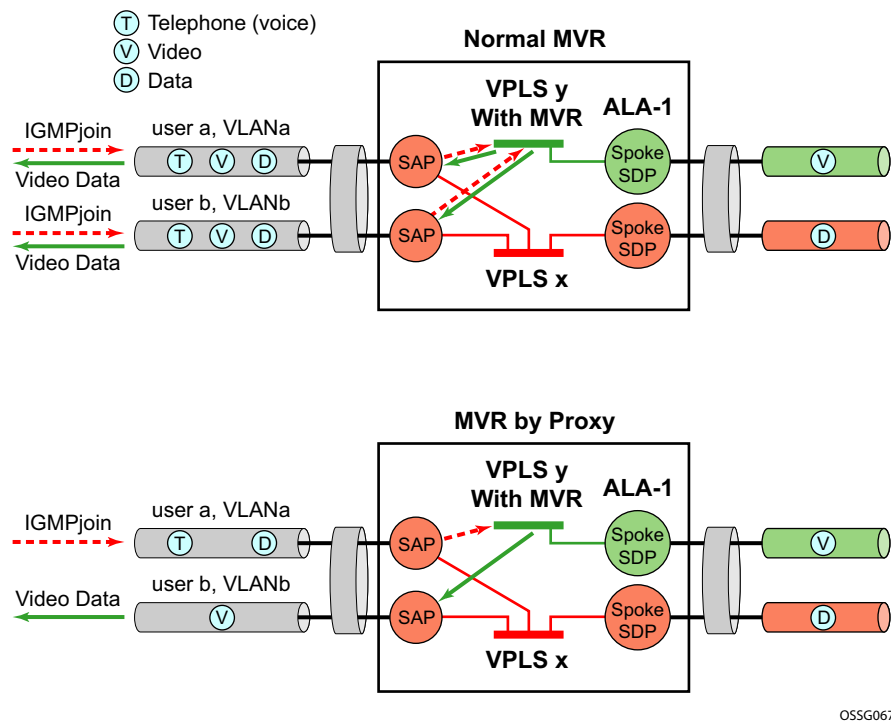


Figure 28: MVR and MVR by Proxy



## Layer 3 Multicast Load Balancing

Layer 3 multicast load balancing establishes a more efficient distribution of Layer 3 multicast data over ECMP and LAG links. Operators have the option to redistribute multicast groups over ECMP and/or LAG links if the number of links changes either up or down.

When implementing this feature, there are several considerations. When multicast load balancing is not configured, the distribution remains as is. Multicast load balancing is based on the number of “s,g” groups. This means that bandwidth considerations are not taken into account. The multicast groups are distributed over the available links as joins are processed. When link failure occurs, the load is distributed on the failed channel to the remaining channels so multicast groups are evenly distributed over the remaining links. When a link is added (or failed link returned) all multicast joins on the added link(s) are allocated until a balance is achieved.

When multicast load balancing is configured, but the channels are not found in the multicast-info-policy, then multicast load balancing is based on the number of “s,g” groups. This means that bandwidth considerations are not taken into account. The multicast groups are distributed over the available links as joins are processed. The multicast groups are evenly distributed over the remaining links. When link failure occurs, the load is distributed on the failed channel to the remaining channels. When a link is added (or failed link returned) all multicast joins on the added link(s) are allocated until a balance is achieved. A manual redistribute command enables the operator to re-evaluate the current balance and, if required, move channels to different links to achieve a balance. A timed redistribute parameter allows the system to automatically, at regular intervals, redistribute multicast groups over available links. If no links have been added or removed from the ECMP/LAG interface, then no redistribution is attempted.

When multicast load balancing is configured, multicast groups are distributed over the available links as joins are processed based on bandwidth configured for the specified group address. If the bandwidth is not configured for the multicast stream then the configured default value is used.

If link failure occurs, the load is distributed on the failed channel to the remaining channels. The bandwidth required over each individual link is evenly distributed over the remaining links.

When an additional link is available for a given multicast stream, then it is considered in all multicast stream additions applied to the interface. This multicast stream is included in the next scheduled automatic rebalance run. A rebalance run re-evaluates the current balance with regard to the bandwidth utilization and if required, move multicast streams to different links to achieve a balance.

A rebalance, either timed or executing the **mc-ecmp-rebalance** command, should be administered gradually in order to minimize the effect of the rebalancing process on the different multicast streams. If multicast re-balancing is disabled and subsequently (re)enabled, keeping with the rebalance process, the gradual and least invasive method is used to minimize the effect of the changes to the customer.

By default multicast load balancing over ECMP links is enabled and set at 30 minutes.

The rebalance process can be executed as a low priority background task while control of the console is returned to the operator. When multicast load rebalancing is not enabled, then ECMP changes will not be optimized, however, when a link is added occurs an attempt is made to balance the number of multicast streams on the available ECMP links. This however may not result in balanced utilization of ECMP links.

Only a single **mc-ecmp-rebalance** command can be executed at any given time, if a rebalance is in progress and the command is entered, it is rejected with the message saying that a rebalance is already in progress. A low priority event is generated when an actual change for a given multicast stream occurs as a result of the rebalance process.

## IGMP State Reporter

The target application for this feature is linear TV delivery. In some countries, wholesale Service Providers are obligated by the government regulation to provide information about channel viewership per subscriber to retailers.

A service provider (wholesaler or retailer) may use this information for:

- billing purposes
- market research/data mining to gain view into the most frequently watched channels, duration of the channel viewing, frequency of channel zapping by the time of the day, etc.

The information about channel viewership is based on IGMP states maintained per each subscriber host. Each event related to the IGMP state creation is recorded and formatted by the IGMP process. The formatted event is then sent to another task in the system (Exporter), which allocates a TX buffer and starts a timer.

The event is then written by the Exporter into the buffer. The buffer in essence corresponds to the packet that will contain a single event or a set of events. Those events are transported as data records over UDP transport to an external collector node. The packet itself has a header followed by a set of TLV type data structures, each describing a unique field within the IGMP event.

The packet is transmitted when it reaches a preconfigured size (1400bytes), or when the timer expires, whichever comes first. Note that the timer started when the buffer was initially created.

The receiving end (collector node) accepts the data on the destination UDP port. It must be aware of the data format so that it can interpret incoming data accordingly. The implementation details of the receiving node are outside of the scope of this description and are left to the network operator.

The IGMP state recording per subscriber host must be supported for hosts which are replicating multicast traffic directly as well as for those hosts that are only keeping track of IGMP states for the HQoS Adjustment purpose. The latter will be implemented via redirection and not the Host Tracking (HT) feature as originally proposed. The IGMP reporting must differentiate events between direct replication and redirection.

It further distinguishes events that are related to denial of IGMP state creation (due to filters, MCAC failure, etc.) and the ones that are related to removal of an already existing IGMP state in the system.

## IGMP Data Records

Each IGMP state change generates a data record that is formatted by the IGMP task and written into the buffer. IGMP state transitions configured statically through CLI are not reported.

In order to minimize the size of the records when transported over the network, most fields in the data record are HEX coded (as opposed to ASCII descriptive strings).

Each data record has a common header as shown in [Figure 29](#):



**Figure 29: Common IGMP Data Record Header**

Application:

- 0x01 - IGMP
- 0x02 - IGMP Host Tracking Event:

Event:

- Related to denial of a new state creation:
  - 0x01 – Join
  - 0x02 – (Join\_Deny\_Filter) Join denied due to filtering via an import policy
  - 0x03 – (Join\_Deny\_CAC) Join denied due to MCAC
  - 0x04 – (Join\_Deny\_MaxGrps) Join denied due to maximum groups per host limit reached
  - 0x05 – (Join\_Deny\_MaxSrcs) Join denied due to maximum sources limit reached
  - 0x06 - Join (Join\_Deny\_SysErr) Join denied due to an internal error (for example: out of memory)
- Related to removal of an existing IGMP state:
  - 0x07 - (Drop\_Leave\_Rx) IGMP state is removed due to the Leave message
  - 0x08 - (Drop\_Expiry) IGMP state is removed due to time out (by default  $2 * \text{query\_interval} + \text{query\_response\_interval} = 260\text{sec}$ )
  - 0x09 - (Drop\_Filter) IGMP state is removed due to filter (import policy) change
  - 0x0A - (Drop\_CfgChange) IGMP state is removed due to configuration change (clear grp, intf shutdown, PPPoE session goes unexpectedly down)

→ 0x0B - (Drop\_CAC) an existing stream is stopped due to configuration change in MCAC

Length:

- The length of the entire data record (including the header and TLVs) in octets.


16 bit Sequence Number

- Since IGMP Reporting is based on connectionless transport (UDP), a 16 bit sequence numbers are used in each data record so that data loss in the network can be tracked.
- The 16 bit sequence number is located after the timestamp field. The sequence numbers will increase sequentially from 0 — 65535 and then rollover back to 0.

Timestamp:

- Timestamp is in Unix format (32 bit integer in seconds since 01/01/1970) plus an extra 8 bits for 10msec resolution.

TLVs describing the IGMP state record will have the following structure:



**Figure 30: Data Record Field TLV Structure**

**Table 10: Data Record Field Description**

Type	Description	Encoding/Length	Mandatory/Optional
0x02	Subscriber ID	ASCII	M
0x03	Sub Host IP	4 Bytes IPv4	M
0x04	Mcast Group IP	4 Bytes IPv4	M
0x05	Mcast Source IP	4 Bytes IPv4	M
0x06	Host MAC	6 Bytes	M
0x07	PPPoE Session-ID	2 Bytes	M
0x08	Service ID	4 Bytes	M
0x09	SAP ID	ASCII	M
0x0A	Redirection vRtrId	4 Bytes	M
0x0B	Redirection ifIndex	4 Bytes	M

The *redirection destination* TLV is a mandatory TLV that is sent only in cases where redirection is enabled. It contains two 32 bit integer numbers. The first number identifies the VRF where IGMPs are redirected; the second number identifies the interface index.

Optional fields can be included in the data records according to the configuration.

In IGMPv3, if an IGMP message (Join or Leave) contains multiple multicast groups or a multicast group contains multiple IP sources, only a single event is generated per group-source combination. In other words, data records are transmitted with a single source IP address and multiple mcast group addresses or a single multicast group address with multiple source IP addresses, depending on the content of the IGMP message. (\*,G)

## Transport Mechanism

Data is transported via UDP socket. Destination IP address, the destination port and the source IP address are configurable. The default UDP source and destination port number is 1037.

Upon the arrival of an IGMP event, the Exporter allocates a buffer for the packet (if not already allocated) and starts writing the events into the buffer (packet). Along with the initial buffer creation, a timer is started. The trigger for the transmission of the packet is either the TX buffer being filled up to 1400B (hard coded value), or the timer expiry, whichever comes first.

The source IP address is configurable within GRT (by default system IP), and the destination IP address can be reachable only via GRT. The source IP address is modified via **system>security>source-address>application** CLI hierarchy.

The receiving end (the collector node) collects the data and process them according to the formatting rules defined in this document. The capturing and processing of the data on the collector node is outside of the context of this description.

It should be noted that the processing node will need to have sufficient resources to accept and process packets that contain information about every IGMP state change for every host from a set of network BRASes that are transporting data to this particular collector node.

Multicast Reporter traffic will be marked as BE (all 6 DSCP bits are set to 0) exiting our system.

---

## HA Compliance

IGMP Events are synchronized between two CPMs before they are transported out of the system.

---

## QoS Awareness

IGMP Reporter is a client of sgt-qos so that DSCP/dot1p bits can be appropriately be marked when egressing the system.

---

## Hardware Support

The following hardware is supported on the 7750 platform.

IOM support: IOM3, HSMDAv2, Ethernet based non HS-MDAs

Chassis mode: B, C, and D.

## **IGMP Reporting Caveats**

The following are not supported:

- Regular (non-subscriber) interfaces
- SAM support as the collector device



## Multicast Support over Subscriber Interfaces in Routed CO Model

Applications for multicast over Subscriber Interfaces in Routed CO ESM model can be divided in two main categories:

Residential customers where the driver applications are:

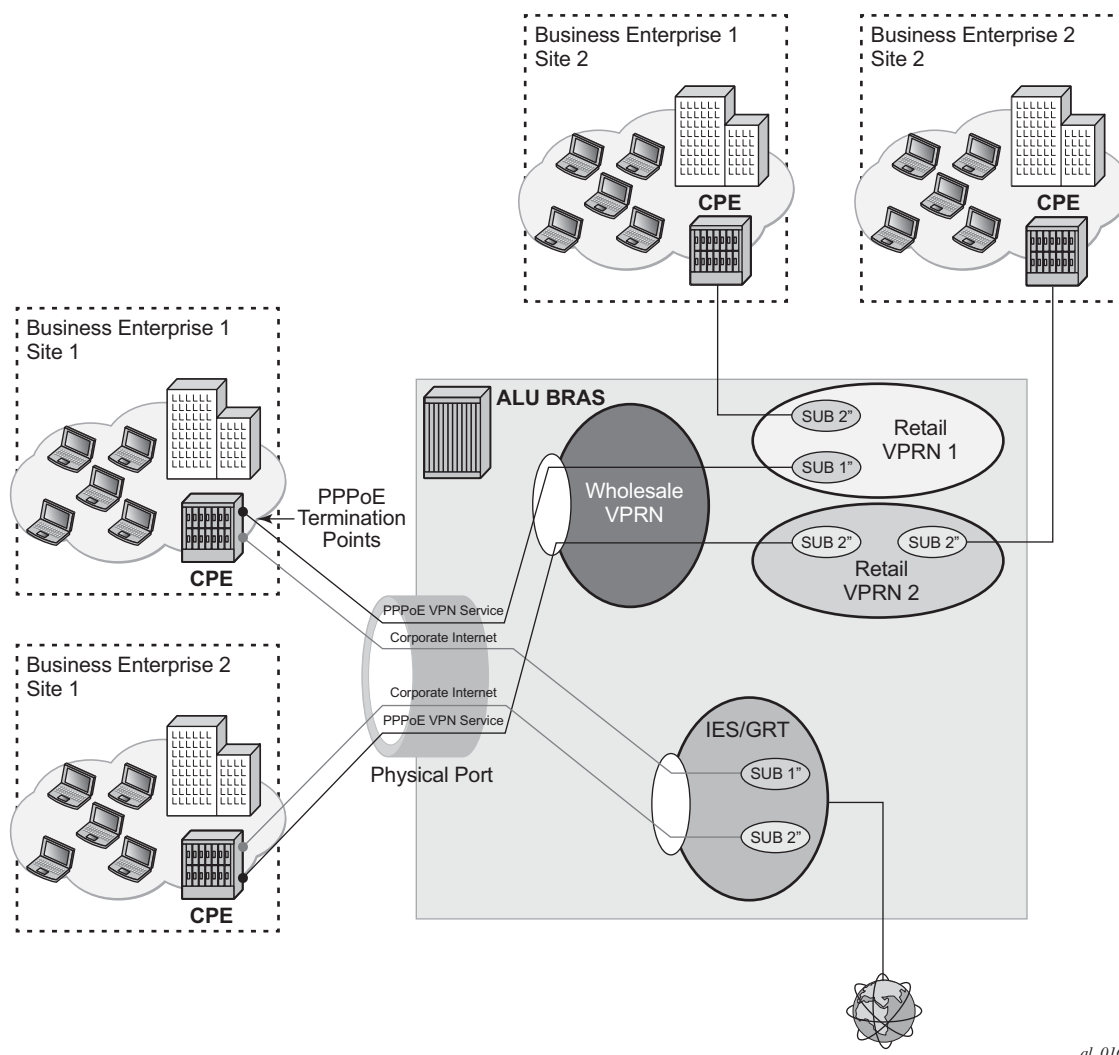
- IPTV in an environment with legacy non-multicasting DSLAMs
- Internet multicast where users connect to a multicast stream sourced from the Internet.

For the business customers, the main drivers are enterprise multicast and Internet multicast applications.

On multicast-capable ANs, a single copy of each multicast stream is delivered over a separate regular IP interface. AN would then perform the replication. This is how multicast would be deployed in Routed CO environment with 7x50s.

On legacy, non-multicast ANs, or in environments with low volume multicast traffic where it is not worth setting up a separate multicast topology (from BNG to AN), multicast replication is performed via subscriber-interfaces in 7x50. There are differences in replicating multicast traffic on IPoE vs PPPoX which will be described in subsequent sessions.

An example of a business connectivity model is shown in [Figure 31](#).



**Figure 31: A Typical Business Connectivity Model**

In this example, HSI is terminated in a Global Routing Table (GRT) whereas VRPN services are terminated in Wholesale/Retail VRPN fashion, with each customer using a separate VRPN.

The actual connectivity model that will be deployed depends on many operational aspects that are present in the customer environment.

Multicast over subscriber-interfaces in a Routed CO model is supported for both types of hosts, IPoE and PPPoE which can be simultaneously enabled on a shared SAP.

There are some fundamental differences in multicast behavior between two host types (IPoE and PPPoX). The differences will be discussed further in the next sections.

## Hardware Support

Multicast over subscriber interfaces is supported on all FP2 based hardware that supports Routed CO model. This includes:

- 7750 SR-7/12
- 7750-c4/12
- 7450 in mixed mode

Chassis modes B, C and D are supported.

## Multicast Over IPoE

There are several deployment scenarios for delivering multicast directly over subscriber hosts:

- 1:1 model (subscriber per VLAN/SAP) with the Access Node (AN) that is not IGMP/MLD aware.
- N:1 model (service per VLAN/SAP) with the AN in the Snooping mode.
- N:1 model with the AN in the Proxy mode.
- N:1 model with the AN that is not IGMP/MLD aware.

There are two modes of operation for subscriber multicast that can be chosen to address the above mentioned deployment scenarios:

1. Per SAP replication — A single multicast stream per group is forwarded on any given SAP. Even if the SAP has a multicast group (channel) that is registered to multiple hosts, only a single copy of the multicast stream is forwarded over this SAP. The multicast stream will have a multicast destination MAC address (as opposed to unicast). IGMP/MLD states will be maintained per host. This is the default mode of operation.
2. Per subscriber host replication in this mode of operation, multicast is replicated per subscriber host even if this means that multiple copies of the same stream will be forwarded over the same SAP. For example, if two hosts on the same SAP are registered to receive the same multicast group (channel), then this multicast channel will be replicated twice on the same SAP. The streams will have a unique unicast destination MAC address (otherwise it would not make sense to replicate the streams twice).

In all deployment scenarios and modes of operation the IGMP/MLD states per source IP address of the incoming IGMP/MLD message is maintained. This source IP address might represent a subscriber hosts or the AN (proxy mode).

For MLD, the source IP address is the host link local address. Therefore, the MLD message is associated with all IP address/prefix of the host.

---

### Per SAP Replication Mode

In the per SAP replication mode a single copy of the multicast channel is forwarded per SAP. In other words, if a subscriber (in 1:1 mode) or a group of subscribers (in N:1 mode) have multiple hosts and all of them are subscribed to the same multicast group (watching the same channel), then only a single copy of the multicast stream for that group will be sent. The destination MAC address will always be a multicast MAC (there will be no conversion to unicast mac address).

IGMP/MLD states are maintained per subscriber host and per SAP.

## Per SAP Queue

Multicast traffic over subscribers in a per SAP replication mode is flowing via a SAP queue which is outside of the subscriber queues context. Sending the multicast traffic over the default SAP queue is characterized by:

- The inability to classify multicast traffic into separate subscriber queues and therefore include it natively in HQoS. However, multicast traffic can be classified into a specific SAP queues, assuming that such queues are enabled via SAP based QoS policy. While multiple SAP queues can be defined under static SAPs, the dynamic SAPs (MSAPs) are limited to a single SAP queue defined in the default egress-sap policy. This default egress-sap policy under MSAP cannot be replaced or modified.
- Redirection of multicast traffic via internal queues in case that the SAP queue in subscriber environment is disabled (**sub-sla-mgmt>single-sub-parameters>profiled-traffic-only**). This is applicable only to 1:1 subscriber model.
- A possible necessity for HQoS Adjustment as multicast traffic is flowing outside of the subscriber queues.
- De-coupling of the multicast forwarding statistics from the overall subscriber forwarding statistics obtained via subscriber specific show commands.


---

## IPoE 1:1 Model (Subscriber per VLAN/SAP) — No IGMP/MLD in AN

This model is shown in [Figure 32](#). The AN is not IGMP/MLD aware, all replications are performed in the BNG. From the BNG perspective this deployment model has the following characteristics:

- IGMP/MLD states are kept per hosts and SAPs. Each host can be registered to more than one group.
- MLD utilize link-local as source address, which makes it difficult to associate with the originating host. MLD (S,G) are always first associated with a IPv6 WAN host if available. If this WAN host terminate its IPv6 session (via lease expiry, session terminate, etc.), the (S,G) is then associated to any remaining WAN host. Only when there are no more WAN hosts available will the (S,G) be associated to the PD host, if any.
- IGMP/MLD Joins will be accepted only from the active subscriber hosts as dictated by antispoofing.
- IGMP/MLD statistics can be displayed per host or per group.
- Multicast traffic for the subscriber is forwarded through the egress SAP queue. In case that the SAP queue is disabled (profiled-traffic-only command), multicast traffic will flow via internal queues outside of the subscriber context.
- A single copy of any multicast stream is generated per SAP. This can be viewed as replication per unique multicast group per SAP, rather than the replication per host. In other words, the number of multicast streams on this SAP is equal to the number of unique groups across all hosts on this SAP (subscriber).

- Traffic statistics are kept per the SAP queue. Consequently multicast traffic stats will be shown outside of the subscriber context.
- HQoS Adjustment might be necessary.
- Traffic cannot be explicitly classified (forwarding classes and queue mappings) inside of the subscriber queues.
- Redirection to the common multicast VLAN (or Layer 3 interface) is supported.
- Multicast streams have multicast destination MAC.
- MLD utilize link-local as source address, which makes it difficult to associate with the originating host. MLD (s,g) is associated with the IPv6 host. Therefore, if any WAN host or PD host end their IPv6 session (via lease expire, etc.), the (s,g) is associated with the remaining host address/prefix. The (s,g) will be delivered to the subscriber as long as a IPv6 address or prefix remains.




### Figure 32: 1:1 Model

## IPoE N:1 Model (Service per VLAN/SAP) — IGMP/MLD Snooping in the AN

This model is shown in [Figure 33](#). The AN is IGMP/MLD aware and is participating in multicast replication. From the BNG perspective this deployment model has the following characteristics:

- IGMP/MLD states are kept per hosts and SAPs. Each host can be registered to more than one group.
- IGMP/MLD Joins are accepted only from the active subscriber hosts as dictated by antispoofing.
- IGMP/MLD statistics are displayed per host, per group or per subscriber.
- Multicast traffic for ALL subscribers on this SAP is forwarded through the egress SAP queues.
- A single copy of any multicast stream is generated per SAP. This can be viewed as the replication per unique multicast group per SAP, rather than the replication per host or subscriber. In other words, the number of multicast streams on this SAP is equal to the number of unique groups across all hosts and subscribers on this SAP.
- The AN will receive a single multicast stream and based on its own (AN) IGMP/MLD snooping information, it will replicate the mcast stream to the appropriate subscribers.
- Traffic statistics are kept per the SAP queue. Consequently multicast traffic stats will be shown on a per SAP basis (aggregate of all subscribers on this SAP).
- Traffic cannot be explicitly classified (forwarding classes and queue mappings) inside of the subscriber queues.
- Redirection to the common multicast VLAN is supported.
- Multicast streams have multicast destination MAC.
- IGMP Joins are accepted (src IP address) only for the sub hosts that are already created in the system. IGMP Joins coming from the hosts that are nonexistent in the system will be rejected, unless this functionality is explicitly enabled by the sub-hosts-only command under the IGMP group-int CLI hierarchy level.
- MLD utilize link-local as source address, which makes it difficult to associate with the originating host. MLD (S,G) are always first associated with a IPv6 WAN host, if available. If this WAN host terminate its IPv6 session (via lease expiry, session terminate, etc), the (S,G) is then associated to any remaining WAN host. Only when there are no more WAN hosts available, will the (S,G) be associated to the PD host, if any.
- MLD join are only accepted if it matches the subscriber host link local address. MLD Joins coming from the hosts that are nonexistent in the system will be rejected, unless this functionality is explicitly enabled by the sub-hosts-only command under the MLD group-int CLI hierarchy level.



al\_0012

Figure 33: - N:1 Model - AN in IGMP Snooping Mode



## IPoE N:1 Model (Service per VLAN/SAP) — IGMP/MLD Proxy in the AN

This model is shown in [Figure 34](#). The AN is configured as IGMP/MLD Proxy node and is participating in downstream multicast replication.

For IPv4, IGMP messages from multiple sources (subscribers hosts) for the same multicast group are consolidated in the AN into a single IGMP messages. This single IGMP message has the source IP address of the AN.

For IPv6, MLD messages from multiple sources (subscribers hosts) for the same multicast group are consolidated in the AN into a single MLD messages. This single MLD message has the link-local IP address of the AN.

From the BNG perspective this deployment model has the following characteristics:

- Subscriber IGMP/MLD states are maintained in the AN.
- IGMP Joins are accepted from the source IP address that is different from any of the subscriber' IP addresses already existing in the BNG. This will be controlled via an IGMP filter on a per group-interface level assuming that the IGMP processing for subscriber hosts is disabled with the no **sub-hosts-only** command under the router/service **vprn>igmp>group-interface** CLI hierarchy. In this case all IGMP messages that cannot be related to existing hosts will be treated in the context of the sap while IGMP messages from the existing hosts will be treated in the context of the subscriber hosts.
- MLD Joins are only accepted if the link-local address matches the subscriber' link local address. To allow processing of foreign link-local address such as the AN link local address, the MLD processing for subscriber hosts should be disabled with the no **sub-hosts-only** command under the **router/service vprn>mld>group-interface** CLI hierarchy. In this case all MLD messages that cannot be related to existing hosts will be treated in the context of the sap while MLD messages from the existing hosts will be treated in the context of the subscriber hosts.
- IGMP/MLD statistics can be displayed per group-interface.
- Multicast traffic for all subscribers on this SAP is forwarded through the egress SAP queue.
- A single copy of any multicast stream is generated per SAP.
- The AN will receive a single multicast stream. Based on the IGMP/MLD proxy information, the AN will replicate the mcast stream to the appropriate subscribers.
- Traffic statistics are maintained per SAP queue.
- HQoS Adjustment is not useful because the per host/subscriber IGMP/MLD granularity is lost. IGMP/MLD states are aggregated per AN.
- Traffic can be explicitly classified into a specific SAP queues via a QoS policy applied under the SAP.
- Multicast streams have multicast destination MAC.

In the following example, IGMPs from the source IP address <ip> is accepted even though there is no subscriber-host with that IP addresses present in the system. An IGMP state will be created under the sap context (service per vlan, or N:1 model) for the group <pref-definition>. All other IGMP messages originated from non-subscriber hosts will be rejected. IGMP messages for subscriber hosts will be processed according to the igmp-policy applied to each subscriber host.

```
configure
  service vprn <id>
    igmp
      group-interface <name>
        import <policy-name>

configure
  router
    policy-options
      begin
        prefix-list <pref-name>
          prefix <pref-definition>

        policy-statement proxy-policy
          entry 1
            from
              group-address <pref-name>
              source-address <ip>
              protocol igmp
            exit
          action accept
        exit
        exit
        default-action reject
```

This functionality (accepting IGMP from non-subscriber hosts) can be disabled with the following flag.

```
configure
  service vprn <id>
    igmp
      group-interface <name>
        sub-host-only
```

In this case only per host IGMP processing will be allowed.

In the following example, MLDs with foreign link-local-address is accepted even though there is no subscriber-host with that link local addresses present in the system. An MLD state will be created under the sap context (service per vlan, or N:1 model) for the group <pref-definition>. All other MLD messages originated from non-subscriber hosts will be rejected. MLD messages for subscriber hosts will be processed according to the igmp-policy applied to each subscriber host.

```

configure
  service vprn <id>
    mld
      group-interface <name>
        import <policy-name>

configure
  router
    policy-options
      begin
        prefix-list <pref-name>
          prefix <pref-definition>

        policy-statement proxy-policy
          entry 1
            from
              group-address <pref-name>
              source-address <ip>
              protocol igmp
            exit
          action accept
        exit
      exit
      default-action reject

```

This functionality (accepting MLD from non-subscriber hosts) can be disabled with the following flag.

```

configure
  service vprn <id>
    mld
      group-interface <name>
        sub-host-only

```

In this case only per host MLD processing will be allowed.




Figure 34: N:1 Model - AN in Proxy mode

## Per Subscriber Host Replication Mode

In this mode a multicast stream is transmitted per subscriber hosts for each registered multicast group (channel). As a result, multiple copies of the same multicast stream destined to different destinations can be transmitted over the same SAP. In this case traffic flows within the subscriber queues and consequently it is accounted in HQoS. As a result, HQoS Adjustment is not needed. Each copy of the same multicast stream have a unique unicast destination MAC addresses. The per host unicast MAC destination addresses are necessary to differentiate multiple copies between different receivers on the same SAP.


Per host replication mode can be enabled on a subscriber basis with the **per-host-replication** command in the **config>subscriber-management>igmp-policy** context.

For IPv6, the command is in the **config>subscriber-management>mld-policy** context.

### IPoE 1:1 Model (Subscriber per VLAN/SAP) — No IGMP/MLD in AN

This model is shown in [Figure 35](#). The AN is not IGMP/MLD aware and multicast replication is performed in the BNG. Multicast streams are sent directly to the hosts using their unicast MAC addresses. HQoS adjustment is not needed as multicast traffic is flowing through subscriber queues. From the BNG perspective this deployment model has the following characteristics:

- IGMP/MLD states are kept per hosts. Each host can be registered to multiple IGMP/MLD groups.
- MLD utilize link-local as source address, which makes it difficult to associate with the originating host. MLD (S,G) are always first associated with a IPv6 WAN host if available. If this WAN host terminate its IPv6 session (via lease expiry, session terminate, etc), the (S,G) is then associated to any remaining WAN host. Only when there are no more WAN hosts available, will the (S,G) be associated to the PD host, if any.
- IGMP/MLD Joins will be accepted only from the active subscriber hosts. In other words antispoofing is in effect for IGMP/MLD messages.
- IGMP/MLD statistics can be displayed per host, per group or per subscriber.
- Multicast traffic is forwarded through subscriber queues using unicast destination MAC address of the destination host.
- Multiple copies of the same multicast stream can be generated per SAP. The number of copies depends on the number of hosts on the SAP that are registered to the same multicast group (channel). In other words, the number of multicast streams on the SAP is equal to the number of groups registered across all hosts on this SAP.
- Traffic statistics are kept per the host queue. In case that multicast statistics need to be separated from unicast, the multicast traffic should be classified in a subscriber separate queue.
- HQoS Adjustment is not needed as traffic is flowing within the subscriber queues and is automatically accounted in HQoS.
- Multicast traffic can be explicitly classified into forwarding classes and consequently directed into desired queues.
- MCAC is supported.
- profiled-traffic-only mode defined under sub-sla-mgmt is supported. This mode (profiled-traffic-only) is used to save the number of queues in 1:1 model (sub-sla-mgmt-> no multi-sub-SAP) by preventing the creation of the SAP queues. Since multicast traffic is not using the SAP queue, enabling this feature will not have any effect on the multicast operation.



al\_0011

Figure 35: 1:1 Model

### **IPoE N:1 Model (Service per VLAN/SAP) — No IGMP/MLD in the AN**

This model is shown in [Figure 36](#). The AN is not IGMP/MLD aware and is not participating in multicast replication. From the BNG perspective this deployment model has the following characteristics:

- IGMP/MLD states are kept per hosts. Each host can be registered to multiple multicast groups.
- MLD utilize link-local as source address, which makes it difficult to associate with the originating host. MLD (S,G) are always first associated with a IPv6 WAN host if available. If this WAN host terminate its IPv6 session (via lease expiry, session terminate, etc), the (S,G) is then associated to any remaining WAN host. Only when there are no more WAN hosts available, will the (S,G) be associated to the PD host, if any.
- IGMP/MLD Joins will be accepted only from the active subscriber hosts, subject to antispoofing.
- IGMP/MLD statistics can be displayed per host, per group or per subscriber.
- Multicast traffic is forwarded through subscriber queues using unicast destination MAC address of the destination host.
- Multiple copies of the same multicast stream can be generated per SAP. The number of copies depends on the number of hosts on the SAP that are registered to the same multicast group (channel). In other words, the number of multicast streams on the SAP is equal to the number of groups registered across all hosts on this SAP.
- Traffic statistics are kept per the host queue. In case that multicast statistics need to be separated from unicast, the multicast traffic should be classified in a separate subscriber queue.
- HQoS Adjustment is NOT needed as traffic is flowing within the subscriber queues and is automatically accounted in HQoS.
- Multicast traffic can be explicitly classified into forwarding classes and consequently directed into desired queues.
- MCAC is supported.




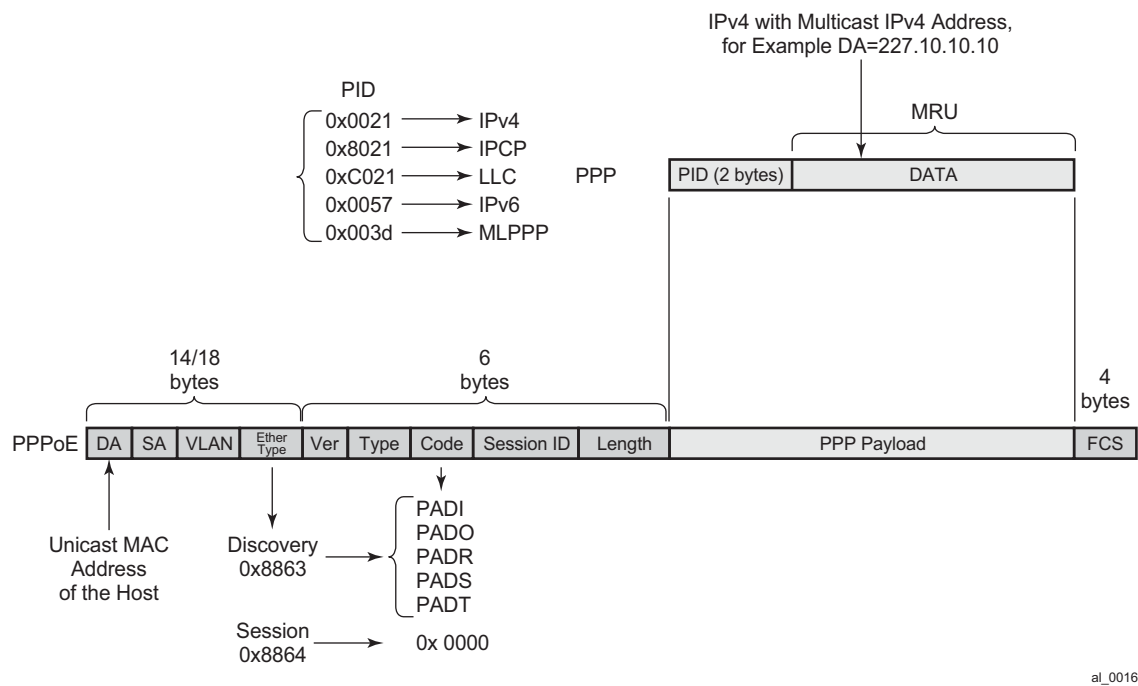


Figure 36: N:1 Model — No IGMP/MLD in the AN

Multicast Over PPPoE

In a PPPoE environment, multicast replication is performed per session (host) regardless of whether those sessions are shared per SAP or they reside on individual SAPs. This is due to the point-to-point nature of PPPoE sessions. There will be no need for HQoS adjustment as multicast is part of the PPPoE session traffic that is flowing via subscriber queues. Multicast packets are sent with unicast MAC address to each CPE. PPP protocol field is set to IP and the destination IP address is the multicast group address for each unique session ID (Figure 37).



al\_0016

Figure 37: Multicast IPv4 Address and Unicast MAC Address in PPPoE Subscriber Multicast

## IGMP Flooding Containment

The query function in IGMP can cause some unintended flooding in N:1 IPoE deployment model with AN in the IGMP snooping mode. By maintaining IGMP session states per host, it is assumed that the IGMP interaction between multicast receivers and the BNG will be on a one-to-one basis. Upon arrival of an IGMP leave from a host for a specific multicast group, the IGMP querier would normally multicast a group-specific query (fast-leave). In N:1 model with sap-replication mode enabled, 7x50 will send a group-specific query (fast-leave) only when it receives the IGMP leave message for the last group shared amongst all subscribers on this SAP.

---

## IGMP/MLD Timers

IGMP/MLD timers are maintained under the following hierarchy:

IPv4:

```
configure>router>igmp
```

```
configure>service vprn>igmp
```

IPv6:

```
configure>router>mld
```

```
configure>service vprn>mld
```

As it can be seen, the IGMP/MLD timers are controlled on a per routing instance (VRF or GRT) level.

The timer values are used to:

- Determine the interval at which queries are transmitted (query-interval).
- To determine the amount of time after which a join will time out.

However, the timers can be different for hosts and redirected interface in case that redirection between VRFs is enabled.

### IGMP/MLD Query Intervals

IGMP/MLD query related intervals (query-interval, query-last-member-interval, query-response-interval, robust-count) are configured on a global router/vprn IGMP/MLD level. They are used to determine the IGMP/MLD timeout states and the rates at which queries are transmitted.

In case of redirection, the subscriber-host IGMP/MLD state will determine the IGMP/MLD state on the redirected interface, assuming that IGMP/MLD messages are not directly received on the redirected interface (for example from the AN performing IGMP/MLD forking). For example if the redirected interface is not receiving IGMP/MLD messages from the downstream node, then the IGMP/MLD state under redirected interface will be removed simultaneously with the removal of the IGMP/MLD state for the subscriber host (due to leave or a timeout).

In case that the redirected interface is receiving IGMP/MLD message directly from the downstream node, the IGMP/MLD states on that redirected interface will be driven by those direct IGMP/MLD messages.

For example, an IGMP/MLD host in VRF1 has an expiry time of 60 seconds and the expiry time defined under the VRF2 where multicast traffic is redirected is set to 90 seconds. The IGMP/MLD state will time out for the host in VRF1 after 60s, and if no host has joined the same multicast group in VRF2 (where redirected interface resides), the IGMP state will be removed there too.

If a join was received directly on the redirection interface in VRF2, the IGMP/MLD state for that group will be maintained for 90s, regardless of the IGMP/MLD state for the same group in VRF1.

---

### HQoS Adjustment

HQoS Adjustment is required in the scenarios where subscriber multicast traffic flow is disassociated from subscriber queues. In other words, the unicast traffic for the subscriber is flowing through the subscriber queues while at the same time multicast traffic for the same subscriber is explicitly (through redirection) or implicitly (per-sap replication mode) redirected through a separate non-subscriber queue. In this case HQoS Adjustment can be deployed where preconfigured multicast bandwidth per channel is artificially included in HQoS. For example, bandwidth consumption per multicast group must be known in advance and configured within the 7x50. By keeping the IGMP state per host, the bandwidth for the multicast group (channel) to which the host is registered is known and is deducted as consumed from the aggregate subscriber bandwidth.

The multicast bandwidth per channel must be known (this is always an approximation) and provisioned in the BNG node in advance.

In PPPoE and in IPoE per host replication environment, HQoS Adjustment is not needed as multicast traffic is unicasted to each subscriber and therefore is flowing through subscriber queues.

For HQoS Adjustment, the channel bandwidth definition and association with an interface is the same as in the MCAC case. This is a departure from the legacy HT channel bandwidth definition which is done via multicast-info-policy.

Example of HQoS adjustment:

Channel definition:

```
configure
  router
    mcac
      policy <name>
        <channel definition>
```

Channel bandwidth definition policy can be applied under:

- group-interface

```
configure
  service vprn <id>
    igmp
      group-interface <grp-if-name>
        mcac
          policy <mcac-policy-name>
```

- plain interface

```
configure
  router/service vprn
    igmp
      interface <name>
        mcac
          policy <mcac-policy-name>
```

- retailer group-interface:

```
configure
  service vprn <id>
    igmp
      group-interface fwd-service <svc-id> <grp-if-name>
        mcac
      policy <mcac-policy-name>
```

Enabling HQoS adjustment:

```
configure
  subscriber-management
    igmp-policy <name>
```

## HQoS Adjustment

```
egress-rate-modify [egress-aggregate-rate-limit | scheduler <name>]
```

Applying HQoS adjustment to the subscriber:

```
configure
 subscriber-management
   sub-profile <name>
     igmp-policy <name>
```

In order to activate HQoS adjustment on the subscriber level, the sub-mcac-policy must be enabled under the subscriber via the following CLI:

```
configure
 subscriber-management
   sub-mcac-policy <pol-name>
     no shutdown

configure
 subscriber-management
   sub-profile <name>
     sub-mcac-policy <pol-name>
```

The adjusted bandwidth during operation can be verified with the following commands (depending whether agg-rate-limit or scheduler-policy is used):

```
*B:BNG-1# show service active-subscribers subscriber "sub-1" detail
=====
Active Subscribers
=====
-----
Subscriber sub-1
-----
I. Sched. Policy : up-silver
E. Sched. Policy : N/A                               E. Agg Rate Limit: 4000
I. Policer Ctrl. : N/A
E. Policer Ctrl. : N/A
Q Frame-Based Ac*: Disabled
Acct. Policy      : N/A                               Collect Stats      : Enabled
Rad. Acct. Pol.   : sub-1-acct
Dupl. Acct. Pol.  : N/A
ANCP Pol.         : N/A
HostTrk Pol.      : N/A
IGMP Policy       : sub-1-IGMP-Pol
Sub. MCAC Policy  : sub-1-MCAC
NAT Policy        : N/A
Def. Encap Offset: none                               Encap Offset Mode: none
Avg Frame Size    : N/A
Preference        : 5
Sub. ANCP-String  : "sub-1"
Sub. Int Dest Id  : ""
Igmp Rate Adj     : -2000
RADIUS Rate-Limit: N/A
Oper-Rate-Limit   : 2000
...
-----
*B:BNG-1#
```

Consider a different example with a scheduler instead of agg-rate-limit:

```
*A:Dut-C>config>subscr-mgmt>sub-prof# info
-----
    igmp-policy "poll"
    sub-mcac-policy "smp"
    egress
        scheduler-policy "h1"
        scheduler "t2" rate 30000
    exit
exit
-----

*A:Dut-C>config>subscr-mgmt>igmp-policy# info
-----
    egress-rate-modify scheduler "t2"
    redirection-policy "mc_redir1"
-----
```

Now, assume that the subscriber joins now a new channel with bandwidth of 1mbps (1000 kbps).

```
A:Dut-C>config>subscr-mgmt>sub-prof>egr>sched># show qos scheduler-hierarchy subscriber
"sub_1" detail
=====
Scheduler Hierarchy - Subscriber sub_1
=====
Ingress Scheduler Policy:
Egress Scheduler Policy : h1
-----
Legend :
(*) real-time dynamic value
(w) Wire rates
B   Bytes
-----

Root (Ing)
|
No Active Members Found on slot 1

Root (Egr)
| slot(1)
|--(S) : t1
|   |   AdminPIR:90000      AdminCIR:10000
|   |
|   |   [Within CIR Level 0 Weight 0]
|   |   Assigned:0         Offered:0
|   |   Consumed:0
|   |
|   |   [Above CIR Level 0 Weight 0]
|   |   Assigned:0         Offered:0
|   |   Consumed:0
|   |   TotalConsumed:0
|   |   OperPIR:90000
|   |
|   |   [As Parent]
|   |   Rate:90000
|   |   ConsumedByChildren:0
|   |
|   |--(S) : t2
```

```
| | | AdminPIR:29000 AdminCIR:10000(sum) <==== bw 1000 from igmp sub-
| | | stracted
| | |
| | | [Within CIR Level 0 Weight 1]
| | | Assigned:10000 Offered:0
| | | Consumed:0
| | |
| | | [Above CIR Level 1 Weight 1]
| | | Assigned:29000 Offered:0 <==== bw 1000 from igmp sub-
| | | stracted
| | | Consumed:0
| | |
| | | TotalConsumed:0
| | | OperPIR:29000 <==== bw 1000 from igmp subtracted
| | |
| | | [As Parent]
| | | Rate:29000 <==== bw 1000 from igmp subtracted
| | | ConsumedByChildren:0
| | |
| | | --(S) : t3
| | | AdminPIR:70000 AdminCIR:10000
| | |
| | | [Within CIR Level 0 Weight 1]
| | | Assigned:10000 Offered:0
| | | Consumed:0
| | |
| | | [Above CIR Level 1 Weight 1]
| | | Assigned:29000 Offered:0
| | | Consumed:0
| | |
| | | TotalConsumed:0
| | | OperPIR:29000
| | |
| | | [As Parent]
| | | Rate:29000
| | | ConsumedByChildren:0
```

```
*A:Dut-C>config>subscr-mgmt>igmp-policy# show service active-subscribers sub-mcac
=====
Active Subscribers Sub-MCAC
=====
Subscriber : sub_1
MCAC-policy : smp (inService)
In use mandatory bandwidth : 1000
In use optional bandwidth : 0
Available mandatory bandwidth : 1147482647
Available optional bandwidth : 1000000000
-----
Subscriber : sub_2
MCAC-policy : smp (inService)
In use mandatory bandwidth : 0
In use optional bandwidth : 0
Available mandatory bandwidth : 1147483647
Available optional bandwidth : 1000000000
-----
```



```

-----
Number of Subscribers : 2
=====
*A:Dut-C#

*A:Dut-C# show service active-subscribers subscriber "sub_1" detail
=====
Active Subscribers
=====
-----
Subscriber sub_1 (1)
-----
I. Sched. Policy : N/A
E. Sched. Policy : h1                      E. Agg Rate Limit: Max
I. Policer Ctrl. : N/A
E. Policer Ctrl. : N/A
Q Frame-Based Ac*: Disabled
Acct. Policy      : N/A                      Collect Stats      : Disabled
Rad. Acct. Pol.   : N/A
Dupl. Acct. Pol.  : N/A
ANCP Pol.         : N/A
HostTrk Pol.      : N/A
IGMP Policy        : poll
Sub. MCAC Policy   : smp
NAT Policy         : N/A
Def. Encap Offset : none                      Encap Offset Mode: none
Avg Frame Size    : N/A
Preference        : 5
Sub. ANCP-String  : "sub_1"
Sub. Int Dest Id  : ""
Igm Rate Adj      : N/A
RADIUS Rate-Limit: N/A
Oper-Rate-Limit   : Maximum
...
=====
*A:Dut-C#

*A:Dut-C# show subscriber-mgmt igmp-policy "poll"
=====
IGMP Policy poll
=====
Import Policy      :
Admin Version      : 3
Num Subscribers    : 2
Host Max Group     : No Limit
Host Max Sources   : No Limit
Fast Leave         : yes
Redirection Policy : mc_redir1
Per Host Replication : no
Egress Rate Modify : "t2"
Mcast Reporting Destination Name :
Mcast Reporting Admin State : Disabled
=====
*A:Dut-C#

```

### Host Tracking (HT) Considerations

HT is a light version of HQoS Adjustment feature. The use of HQoS Adjustment functionality in place of HT is strongly encouraged.

When HT is enabled, the AN will fork off (duplicate) the IGMP messages on the common mcast SAP to the subscriber SAP. IGMP states will not be fully maintained per sub-host in the BNG, instead they will be only tracked (less overhead) for bandwidth adjustment purposes.

#### Example of HT

Channel Definition:

```
configure
  mcast-management
    multicast-info-policy <name>
      <channel to b/w mapping definition>
```

Applying channel definition policy on a router/VRPN global level:

```
configure>router>multicast-info-policy <name>
configure>service>vrpn>multicast-info-policy <name>
```

Defining the rate object on which HT will be applied:

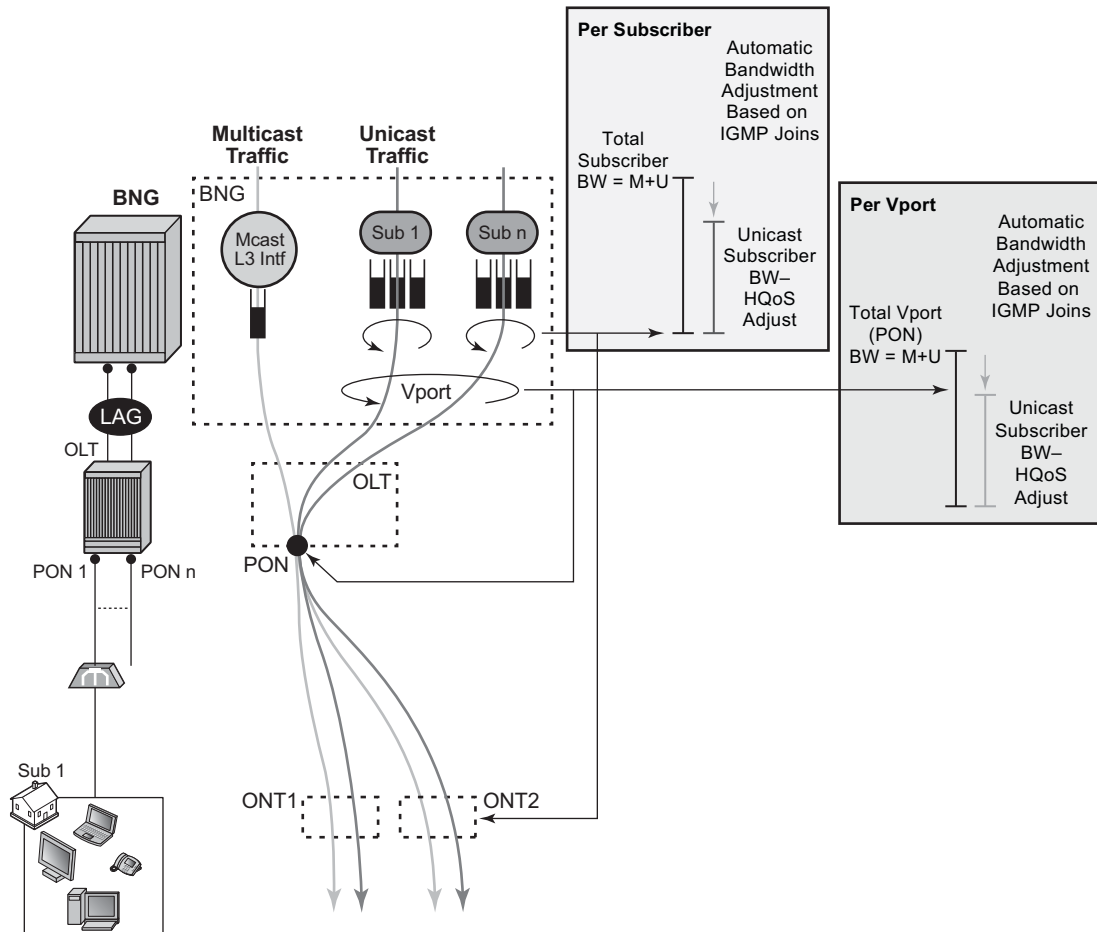
```
configure
  subscriber-management
    host-tracking-policy <name>
      egress-rate-modify [agg-rate-limit | scheduler <sch-name>]
```

Applying the HT to the subscriber:

```
configure
  subscriber-management
    sub-profile <name>
      host-tracking-policy <name> => mutually exclusive with igmp-policy
```

## HQoS Adjust Per Vport

HQoS adjust per Vport can be used in environments where Vport represents a physical medium over which traffic for multiple subscribers is shared. Typical example of this scenario is shown in [Figure 38](#). Multicast traffic within 7x50 is taking a separate path from unicast traffic, only for the two traffic flows to merge later in the PON (represented by Vport in 7x50) and ONT (represented by subscriber in 7x50).



al\_0166

**Figure 38: HQoS Adjustment per Subscriber and Vport**

A single copy of each channel is replicated on the PON as long as there is at least one subscriber on that PON interested in this channel (has joined the IGMP/MLD group).

7x50 monitors IGMP/MLD Joins at the subscriber level and consequently the channel bandwidth is subtracted from the current Vport rate limit only in the case that this is the first channel flowing through the corresponding PON. Otherwise, the Vport bandwidth is not modified. Similarly, when

the channel is removed from the last subscriber on the PON, the channel bandwidth is returned to the VPort.

Association between the Vport and the subscriber is performed via inter-destination-string or svlan during the subscriber setup phase. Inter-destination-string can be obtained either via Radius or LUDB. In case that the association between the Vport and the subscriber is performed based on the svlan (as specified in sub-sla-mgmt under the sap/msap), then the destination string under the Vport must be a number matching the svlan.

The mcac-policy (channel definition bandwidth) can be applied on the group interface under which the subscribers are instantiated or in case of redirection under the redirected-interface.

In a LAG environment, the Vport instance is instantiated per member LAG link on the IOM. For accurate bandwidth control, it is prerequisite for this feature that subscriber traffic hashing is performed per Vport.

The CLI structure is as follows.

```
configure
  port <port-id>
    ethernet
      access
        egress
          vport <name>
            egress-rate-modify
            agg-rate
            host-match <destination-string>
            port-scheduler-policy <port-scheduler-policy-name>

configure
  port <port-id>
    sonnet-sdh
      path [<sonnet-sdh-index>]
        access
          egress
            vport <name>
              egress-rate-modify
              agg-rate
              host-match <destination-string>
              port-scheduler-policy <port-scheduler-policy-name>
```

The Vport rate that will be affected by this functionality depends on the configuration:

- In case the agg-rate-limit within the Vport is configured, its value will be modified based on the IGMP activity associated with the subscriber under this Vport.
- In case that the port-scheduler-policy within the Vport is referenced, the max-rate defined in the corresponding port-scheduler-policy will be modified based on the IGMP activity associated with the subscriber under this Vport.

Note that HQoS adjust is not supported when a scheduler policy is configured under the VPORT.

The Vport rates can be displayed with the following two commands:

**show port 1/1/5 vport** *name*

**qos scheduler-hierarchy port** *port-id* **vport** *vport-name*

As an example:

```
*A:system-1# show port 1/1/7 vport
=====
Port 1/1/7 Access Egress vport
=====
VPort Name      : isam1
Description     : (Not Specified)
Sched Policy    : 1
Rate Limit      : Max
Rate Modify     : enabled
Modify delta    : -14000
```

In this case, the configured Vport aggregate-rate-limit max value has been reduced by 14Mbps.

Similarly, if the Vport had a port-scheduling-policy applied, the max-rate value configured in the port-scheduling-policy would have been modified by the amount shown in the Modify delta output in the above command.

---

## MULTI-CHASSIS REDUNDANCY

Modified Vport rate synchronization in multi-chassis environment relies on the synchronization of the subscriber IGMP/MLD states between the redundant nodes. Upon the switchover, the Vport rate on the newly active node is adjusted according to the current IGMP/MLD state of the subscribers associated with the Vport.

---

## SCALABILITY CONSIDERATIONS

It is assumed that the rate of the IGMP/MLD state change on the Vport level is substantially lower than on the subscriber level.

The reason for this is that the IGMP/MLD Join/Leaves are shared amongst subscribers on the same Vport (PON for example) and thus the IGMP/MLD state on the VPort level is changed only for the first IGMP/MLD Join per channel and the last IGMP/MLD leave per channel.

## Redirection

Two levels of MCAC can be enabled simultaneously and in such case this is referred as Hierarchical MCAC (H-MCAC). In case that redirection is enabled, H-MCAC per subscriber and the redirected interface is supported. However, mcac per group-interface in this case is not supported. Channel definition policy for the subscriber and the redirected interface is in this case referenced under the **igmp->interface** (redirected interface) CLI or for IPv6 **mld->interface**.

In case that redirection is disabled, H-MCAC for both, the subscriber and the group-interface is supported. The channel definition policy is in this case configured under the **config>router>igmp>group interface** context or for IPv6 **config>router>mld>group interface**.

There are two options in multicast redirection. The first option is to redirect all subscriber multicast traffic to a dedicated redirect interface.

Example:

Defining redirection action:

```
configure
  router
    policy-options
      begin
        policy-statement <name>
          default-action accept
          multicast-redirection [fwd-service <svc id>] <interface name>
        exit
      exit
    exit
  exit
```

The second option is to redirect only specific multicast groups to the redirect interface while the remaining groups remains on the subscriber SAP. This is applicable for both IPv4 and IPv6. For IPv6 host-ip for a policy statement is not supported.

Example:

Defining redirection action:

```
configure
  router
    policy-options
      begin
        prefix-list <name>
          prefix <IPv4 multicast groups>
          prefix <IPv4 multicast groups>
        exit
        policy-statement <name>
          entry 1
            from
              group-address <prefix-list name>
            action accept
```

```
        multicast-redirection [fwd-service <svc id>] <interface name>
    exit
exit
```

Applying redirection to the subscriber for IGMP and MLD respectively.

```
configure
  subscr-mgmt
    igmp-policy <name>
      redirection-policy <name>
    exit
  exit
  mld-policy <name>
    redirection-policy <name>
```

Redirection that cross-connects GRT and VPRN is not supported. Redirection can be only performed between interfaces in the GRT, or between the interfaces in any of the VPRN (cross connecting VPRNs is allowed).

Redirection is also supported in a wholesaler/retailer VPRN model where redirected Layer 3 interface resides in the retailer VPRN.

## Hierarchical Multicast CAC (H-MCAC)

MCAC is supported on three levels:

- per subscriber
- per group-interface
- per redirected interface

Two levels of MCAC can be enabled simultaneously and in such case this is referred as Hierarchical MCAC (H-MCAC). In case that redirection is enabled, H-MCAC per subscriber and the redirected interface is supported. However, MCAC per group-interface in this case is not supported. Channel definition policy for the subscriber and the redirected interface is in this case referenced under the **config>router>igmp->interface** (redirected interface) CLI hierarchy or IPv6 **config>router>mld->interface**.

In case that redirection is disabled, H-MCAC for both, the subscriber and the group-interface is supported. The channel definition policy is in this case configured under the **config>router>igmp>group-interface** CLI hierarchy or IPv6 **config>router>mld>group-interface**.

Examples

Note that the same channel definition and association with interfaces is used for MCAC/H-MCAC and HQoS Adjustment.

Channel definition:

```
configure
  router
    mcac
      policy <mcac-pol-name>
        bundle <bundle-name>
          bandwidth <kbps>
          channel <start-address> <end-address> bw <bw> [class {high|low}]
        [type {mandatory|optional}]
      :
      :
```

Channel bandwidth definition policy can be referenced under the:

- group-interface — This is used for subscribers when redirection is disabled.

```
configure
  service vprn <id>
    igmp/mld
      group-interface <grp-if-name>
        mcac
          policy <mcac-policy-name>
          policy <mcac-policy-name>
```

- plain interface



```

configure
  router
    igmp/mld
      interface <name>
        mcac
          policy <mcac-policy-name>

configure
  service vprn <id>
    igmp/mld
      interface <if-name>
        mcac
          unconstrained-bw <bandwidth> mandatory-bw <mandatory-bw>

```

- retailer's VPRN reference the group-interface in the wholesaler's VPRN

```

configure
  service vprn <id>
    igmp/mld
      group-interface fwd-service <svc-id> <grp-if-name>
        mcac
          policy <mcac-policy-name>

```

#### Enabling MCAC:

- per subscriber

```

configure
  subscr-mgmt
    sub-mcac-policy <name>
      unconstrained-bw <bandwidth> mandatory-bw <mandatory-bw>

configure
  subscr-mgmt
    sub-profile <name>
      sub-mcac-policy <name>

```

- per-group-interface

```

configure
  service vprn <id>
    igmp/mld
      group-interface <grp-if-name>
        mcac
          unconstrained-bw <bandwidth> mandatory-bw <mandatory-bw>

```

- per redirected interface

```

configure
  router
    igmp/mld
      interface <if-name>
        mcac
          unconstrained-bw <bandwidth> mandatory-bw <mandatory-bw>

configure
  service vprn <id>
    igmp/mld
      interface <if-name>

```

## Hierarchical Multicast CAC (H-MCAC)

```
mcac
    unconstrained-bw <bandwidth> mandatory-bw <mandatory-bw>
```

## MCAC Bundle Bandwidth Limit Considerations

In addition to multicast bandwidth limit that can be imposed on subscribers, group-interfaces or regular interfaces, there is another multicast bandwidth limit that can be imposed on a group of channels (channel bundle).

The MCAC policy, aside from the channel bandwidth definitions, could optionally contain this bandwidth cap for the group of channels:

```
config>router>mcac# info
-----
policy "test"
  bundle "test" create
    bandwidth 100000
    channel 225.0.0.10 225.0.0.10 bw 10000 type mandatory
    channel 225.0.0.11 225.0.0.15 bw 5000 type mandatory
    channel 225.0.0.20 225.0.0.30 bw 5000 type optional
  exit
exit
```

This can be used to prevent a single set of channels from monopolizing MCAC bandwidth allocated to the entire interface. The bandwidth of each individual bundle will be capped to some value below the interface MCAC bandwidth limit, allowing each bundle to have its own share of the interface MCAC bandwidth.

In most cases, the bandwidth limit per bundle is not necessary to configure. The aggregate limit per all channels as defined under the subscriber/interface will cover majority of scenarios. In case that one wants to explore the bundle bandwidth limits and how they affect MCAC behavior, the following text will help understanding this topic.

To further understand how various MCAC bandwidth limits are applied, one need to understand the concept of the mandatory bandwidth that is pre-allocated in the following way:

- Bandwidth of each mandatory channel in a bundle is pre-allocated. The artifacts of this are:
  - The total mandatory bandwidth in the bundle cannot exceed the bundle cap. For the sake of deterministic behavior, the configured bandwidth of each mandatory channel in the bundle is counted towards the total mandatory bandwidth only once. This means that only one replication of each mandatory channel is assumed. This is normal behavior on a regular interface with a single SAP under it. More than one replication of the same channel per regular interface (or sap) would lead to packet duplication.
  - Optional (non-mandatory) channels can use only the difference in bandwidth between the bundle cap and total pre-allocated mandatory bandwidth. They can NOT use more bandwidth than that even if the total pre-allocated mandatory bandwidth is not used up (mandatory channels are not being replicated).
- Mandatory bandwidth under the interface is pre-allocated and subtracted from the unconstrained bandwidth. In the configuration example below, 2mbps is pre-allocated

(guaranteed for mandatory channels) and the remaining 8mbps can be used by the optional channels on a first come first serve basis.

```
config>router>igmp# info
-----
      interface "ge-1/1/1"
          mcac
              unconstrained-bw 10000 mandatory-bw 2000
          exit
      exit
```

The bundle bandwidth limit poses a problem when the MCAC policy is applied under the group-interface. The reason is that the group-interface represents the aggregation point for the subscribers and their bandwidth. As such it is natural that the any aggregated bandwidth limit under the group interface be larger than the bandwidth limit applied to any individual subscriber under it. Since the MCAC policy, along with the bundle bandwidth limit, is inherited by all subscribers under the group-interface, the exhaustion of the bundle bandwidth limit under the group interface will coincide with the exhaustion of the bundle bandwidth limit of any individual subscriber. This will result in a single subscriber starving out of multicast bandwidth the remaining subscribers under the same group-interface. While it is perfectly acceptable for the subscribers to inherit the multicast channel definition from the group-interface, for the above reasons it is not acceptable that the subscriber inherit the bandwidth cap from the group-interface.

To remedy this situation, the MCAC bandwidth limits are independently configured under the group-interface level (aggregated level) and the subscriber level via the command `unconstrained-bw <kbps> mandatory-bw <kbps>`. The undesired bundle bandwidth cap in the MCAC policy will be ignored under the group-interface AND under the subscriber. However, the bundle bandwidth cap will be applied automatically to each SAP under the group interface. A SAP is a natural place for a bundle bandwidth limit since each channel on a SAP can be replicated only once and therefore the amount of pre-allocated mandatory bandwidth can be pre-calculated. This is obviously not the case for the group interface where single channel can be replicated multiple times (one per each SAP under the grp-if). Similarly, the same channel can be replicated multiple times for the same subscriber in per-host replication mode. Only subscribers in per-sap replication mode will warrant a single replication per channel. Therefore, if bundle cap is configured, it will be applied to limit the bandwidth of the bundle that is applied to a subscriber in a per-sap replication mode.

[Figure 39](#) depicts MCAC related inheritances and MCAC bandwidth allocation model in per-sap replication mode. The MCAC policy is applied to the group interface and inherited by each subscriber as well as each SAP under the same group interface. However, the bundle bandwidth limit in the MCAC policy is ignored on the group-interface and under the subscriber (denoted by the red X in the figure). The bundle limit is applied only to each sap under the group-interface.

Overall (non-bundle) MCAC bandwidth limits are independently applied to the group-interface and the subscribers. According to our example, 20mbps of multicast bandwidth in total is allocated per group-interface. 6mbps of the 20mbps is allocated for mandatory channels. This leaves 14mbps of multicast bandwidth for the optional channels combined served on a first come first serve basis. Each physical replication (multiple replications of the same channel can occur, one per each SAP), counts towards the respective group-interface bandwidth limits.

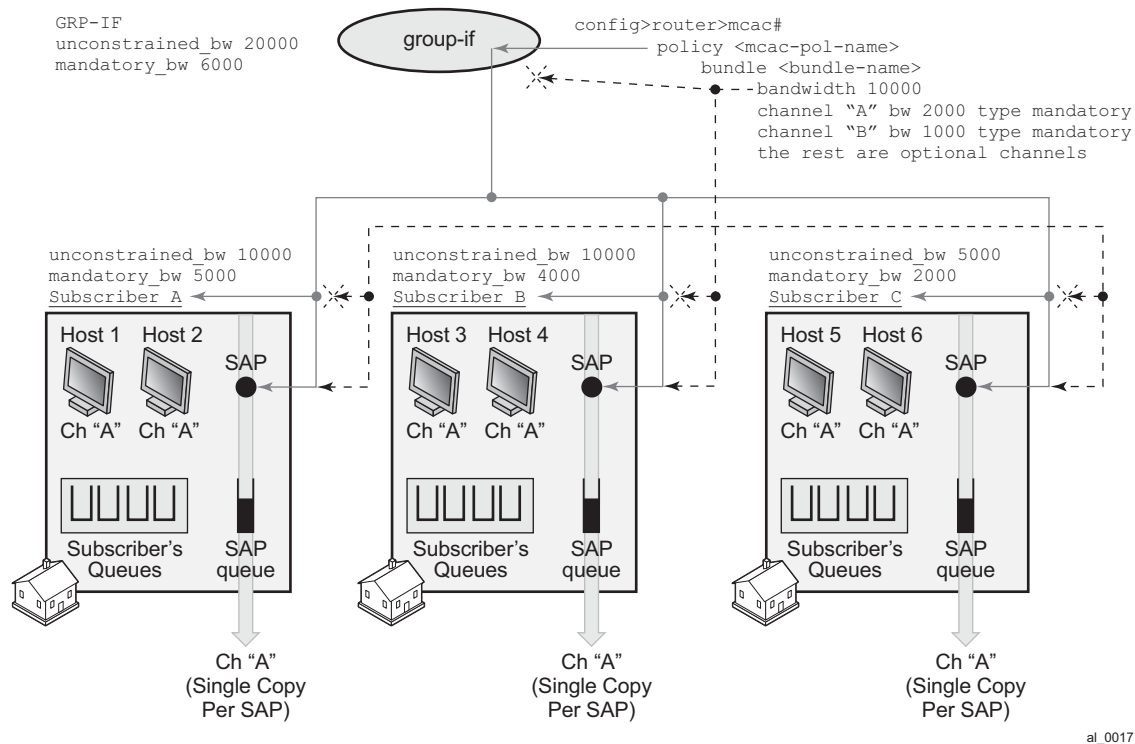
Similar logic applies to the subscriber MCAC bandwidth limits which are applied per sub-profile.

Finally, each SAP can optionally contain the bundle bandwidth limit. Note that in a hierarchical MCAC fashion, if either of the bandwidth checks fails (SAP, sub or grp-if) the channel admission for the subscriber also fails.

In our example, 6 subscriber hosts watch the same channel but there are only 3 active replications (one per SAP). This would yield:

- 14mbps of available multicast bandwidth under the group-interface. This bandwidth can be used for optional channels on a first come first serve basis. No reserved bandwidth is left.
- Subscriber A — 3mbps is still reserved for mandatory channels and 5mps is available for optional channels (first come first serve). All this assume that the SAP and the grp-if bandwidth checks pass.
- Subscriber B — 2 mbps is still reserved for mandatory channels and 6mps is available for optional channels (first come first serve). All this assume that the SAP and the grp-if bandwidth checks pass.
- Subscriber C — No reserved bandwidth for mandatory channels is left. 3 mbps is still left for optional channels. All this assume that the SAP and the grp-if bandwidth checks pass.
- SAPs — Considering that 2mbps are currently replicating (ch A, each SAP can still accept 1 mbps of the mandatory bandwidth (channel B and 7 mbps of the remaining optional channels.

## Hierarchical Multicast CAC (H-MCAC)



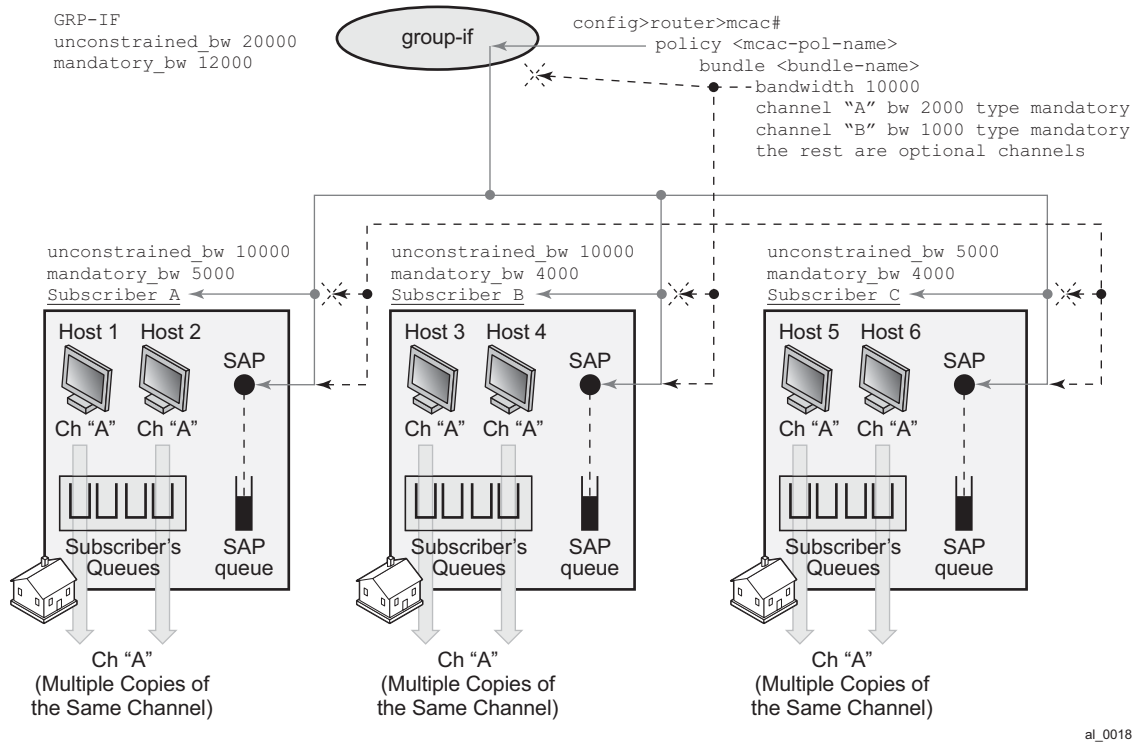
**Figure 39: MCAC Policy Inheritance in Per-SAP Replication Mode**

Figure 40 depicts behavior in per-host replication mode. MCAC policy inheritance flow is the same as in the previous example with the difference that the bundle limit has NO effect at all. Each host generates its own copy of the same multicast stream that is flowing via subscriber queues and not the SAP queue. Since each of the copies counts towards the subscriber or group-interface bandwidth limits, the multicast bandwidth consumption is higher in this example. This needs to be reflected in the configured multicast bandwidth limits. For example, the group-interface mandatory bandwidth limit is increased to 12mbps.

In our example, 6 subscriber hosts are still watching the same channel but now the number of replications is doubled from previous example. So the final tally for our MCAC bandwidth limit is as follows:

- Subscriber A - 1 mbps is still reserved for mandatory channels and 5mps for optional channels (first come first serve). All this assume that SAP and grp-if bandwidth checks pass.
- Subscriber B - No reserved mandatory bandwidth is left. 6 mps is still left for optional channels (first come first serve). All this assume that SAP and grp-if bandwidth checks pass.

- Subscriber C - No reserved mandatory bandwidth is left. 1 mps is still left for optional channels (first come first serve). All this assume that SAP and grp-if bandwidth checks pass.
- No reserved bandwidth is left under the group-interface. 8 mbps of available multicast bandwidth under the group-interface is still left for optional channels on a first come first serve basis
- Bundle limit on a SAP is irrelevant in this case.



**Figure 40: MCAC Policy Inheritance in Per-Host Replication Mode**

## Determining MCAC Policy in Effect

Channel bandwidth definition (via MCAC policy) can be applied under the interface level (group-interface or regular interface):

```
configure>router/service>igmp/mld>interface/grp-if
```

The following configuration options can lead to the confusion as to which MCAC policy is in effect:

- The MCAC policy (channel bandwidth definition) can be applied under two different places (grp-if and/or regular intf).
- The same policy is used for (H)MCAC and HQoS Adjust with redirection enabled/disabled.

The general, the rule is that the MCAC policy under the group-interface will always be in effect in cases where redirection is disabled. This is valid for subscriber or group-interface MCAC, hMCAC (subscriber and group-interface) and HQoS Adjust in per SAP replication mode.

If redirection is enabled, the MCAC policy under the group-interface will be ignored.

If redirection is enabled, but there is no MCAC policy applied under the redirected interface<sup>1</sup> (regardless of whether the MCAC policy under the group-interface is applied or not) then:

- HQoS Adjust will have no effect.
- MCAC will have no effect not only per redirected interface but also per subscriber.

---

1. Redirected interface is the interface to which IGMP/MLD Joins are redirected from subscriber hosts.



## Multicast Filtering

Multicast filtering must be done per session (host) for IPoE and PPPoE. There are two types of filters that are supported:

1. IGMP filters on access ingress. Those filters control the flow of IGMP messages between the host and the BNG. They are applied via the import statement in the igmp-policy. The same filters are used for multicast-redirection policy:

For IPv4

```
configure
  subscr-mgmt
    igmp-policy <name>
      import <policy-name>
```

For IPv6

```
configure
  subscr-mgmt
    mld-policy <name>
      import <policy-name>
```

An example of the filter definition is given below:

```
configure
  router
    policy-options
      begin
        prefix-list <pref-name>
          prefix <pref-definition>
        policy-statement <name>
          entry 1
            from
              group-address <pref-name>
              source-address <ip>
              protocol igmp
            exit
            action accept
            exit
          exit
          default-action reject
```

2. Regular traffic filters where control multicast traffic flow can be controlled in both directions (ingress/egress). This is supported through ip-filters under the SLA profile.

## Joining the Multicast Tree

The delivery of multicast to the subscribers-interface in a VPRN environment depends on the multicast deployment model (PIM, mBGP). In each model, a subscriber-interface is treated as a regular CE-PE interface that has registered v4 multicast listeners.

## Wholesale/Retail Requirements

Multicast support on subscriber interfaces is supported in both wholesale/retail models:

- Wholesale/retail VPRN (IPoE and PPPoE)
- LAC/LNS (PPPoE only)




Figure 41: Wholesale/Retail Multicast Support

The distinction between these two models is that in the case of LAC/LNS, the replication will be done further up in the network on an LNS node. This means that the traffic between LAC and LNS will be multiplied by the amount of replications.

## QoS Considerations

In per-sap replication mode (which is applicable only to IPoE subscribers), multicast traffic is forwarded through the SAP queue which is outside of the subscribers queues and therefore not accounted in subscriber aggregate rate limit. HQoS Adjust is used to remedy this situation.

In case that the SAP queue is removed from the static SAP in IPoE 1:1 model (with profiled-only-traffic command), multicast traffic will flow via internal queues which cannot be tied into a port-scheduler as part of HQoS. Consequently, the port-scheduler max-rate as defined in the port-scheduler-policy will be used only to rate limit unicast traffic. In other words, the max-rate value in port-scheduler-policy must be lowered for the amount of anticipated multicast traffic that will flow via the port where port-scheduler-policy is applied.

A similar logic applies to per-sap replication mode on dynamic SAPs (MSAPs) even if the SAP queue is not removed. Although the multicast traffic is flowing via the SAP queue in this case, the SAP qos policy on MSAP cannot be changed from the default one. The default QoS policy on a SAP contains a single queue that is not parented by the port-scheduler.

Those restrictions do not apply to static SAPs where the SAP QoS policy can be customized and its queues consequently tied to the port-scheduler.

---

## Redundancy Considerations

The subscriber can receive multicast content through the subscriber SAP, the redirected interface, or a combination of both.

Multicast redundancy is only supported on a MC-LAG topology. Multicast traffic can be delivered over the subscriber SAP, the redirected interface, or a combination of both.

Subscriber IGMP states can be synchronized across multiple 7x50 nodes in order to ensure minimal interruption of (video delivery) service during network outages. The IGMP/MLD state of a subscriber-host in a 7x50 node is tied to the state of the underlying MC-LAG protection mechanism. For example, IGMP states will be activated only for subscribers that are anchored under the group-interfaces with master SRRP state or under an active MC-LAG port.

For multicast redirection on a MC-LAG topology, it must be ensured that the redirected interface (the interface to which multicast forwarding is redirected) is under the same MC-LAG as the subscriber. Otherwise, IGMP states on the redirected interface will be derived independently of the IGMP states for the subscriber from which IGMP/MLD messages are redirected.

The IGMP/MLD synchronization process in conjunction with underlying access protection mechanisms will work as follows:

- IGMP/MLD states for the subscriber will be updated only if IGMP/MLD messages (Joins/Leaves/Reports, etc.) are received:

- Directly from the downstream node on an active MC-LAG link. This is valid irrespective of the IGMP querier status for the subscriber.

In all other cases, assuming that some protection mechanism in the access is present, the IGMP/MLD messages are discarded and consequently no IGMP/MLD state is updated. Similar logic applies to regular Layer 3 interfaces, where SRRP is replaced with VRRP.

- Once the subscriber IGMP/MLD state is updated as a result of directly received IGMP/MLD message on an active subscriber (SRRP master of active MC-LAG), the sync IGMP/MLD message is sent to the standby subscriber over the Multi-Chassis Synchronization protocol. Synchronized IGMP states will be populated in Multi-chassis Synchronization (MCS) DB in all pairing 7x50 nodes.
- In case that a IGMP/MLD sync (MCS) message is received from the peering node, the IGMP state for the standby subscriber is updated in the MCS DB but it is not downloaded into the forwarding plane unless there is a switchover. In case that the IGMP/MLD sync message is received for the active subscriber, the message will be discarded.
- IGMP/MLD queries are sent out only by IGMP/MLD querier. In a MC-LAG environment, it is the node with the active MC-LAG link. Note that MC-LAG is usually configured with SRRP and the SRRP state is derived from the MC-LAG.
- IGMP/MLD states from the MCS DB will be:
  - Activated on non-querier subscriber in case that neither SRRP nor MC-LAG is deployed. It is assumed that the querier subscriber has received the original IGMP/MLD message and consequently sent the IGMP/MLD MCS Sync to the non-querier (standby). Non-querier interface will accept the MCS sync message and also it will propagate the IGMP/MLD states to PIM.

The querier subscriber will not accept the IGMP/MLD update from the MCS database.

- Aware of the state of MC-LAG. As soon as the standby MC-LAG becomes active, the IGMP/MLD states will be activated and they will be propagated to PIM. Traffic will be forwarded as soon as multicast streams are delivered to the node and the IGMP states under the subscriber are activated. On a standby MC-LAG, IGMP states will not be propagated from the MCS DB to PIM and consequently subscribers.
- Aware of the SRRP state. Since the subscriber with SRRP Master state is considered active, the states will be propagated to PIM as well. On standby SRRP, IGMP states will not be propagated from MCS DB to PIM and consequently to subscribers.
- For MC-lag setups, once the switchover is triggered via MC-LAG or SRRP, the IGMP/MLD states from MCS DB on the newly active MC-LAG node or subscriber under the newly SRRP Master will be sent to PIM and consequently to the forwarding plane effectively turning on multicast forwarding.

An active and standby subscriber refers to the state of underlying protection mechanism (active MC-LAG). Note that the subscribers themselves are always instantiated (or active) on both nodes. However, traffic forwarding over those subscribers will be driven by the state of the underlying protection mechanism (MC-LAG). Hence the terms active and standby subscriber.

Note that in subscriber environment, SRRP should be always activated in dual-homing scenario. SRRP in subscriber environment will ensure that downstream traffic is forwarded via the same node that is forwarding upstream traffic. In this fashion, accounting and QoS for the subscriber are consolidated within a single node.

To summarize, in multi-chassis environment with subscribers, IGMP synchronization enabled and an access layer protection mechanism in place (MC-LAG), the behavior for is the following:

- IGMP/MLD states are synchronized between the chassis.
  - On a MC-LAG setup, only the SRRP master or active MC-LAG will forward downstream multicast traffic.
  - Length of outage during the switchover is determined by the detection and recovery of the underlying protection mechanism (MC-LAG or MCS) in addition to local propagation of IGMP/MLD states from MCS DB to PIM and consequently to forwarding plane. Note that IGMP/MLD states can be statically configured on both redundant nodes in order to attract multicast traffic from upstream and therefore minimize outage during the switchover.
- 

## Redirection Considerations

The redirection policy has two options wither to redirect only a certain set of multicast groups to the redirect interface or redirect all multicast to the redirect interface. The redirect policy is source agnostic.

On a MC-LAG setup, for redirection and MCS to work simultaneously in predictable manner, the redirected interface and the corresponding subscribers have to be protected by the same MC-LAG. This binds the redirected interfaces and the subscriber-hosts to the same physical port(s).

The following describe some guidelines for a MC-LAG setup:

- The active subscriber will replicate its received IGMP/MLD message to the redirected Layer 3 interface. The Layer 3 redirected interface will accept this message:
  - Independently of the corresponding VRRP state if MC-LAG is not used.
  - Only if the Layer 3 interface is IGMP querier.
  - MC-LAG is used and in active state.
- In all other cases the IGMP message under the Layer 3 redirected interface will be rejected. Note that Layer 3 redirected interface can also receive IGMP message directly from the downstream node in case that IGMP forking in the access node is activated.
- The Layer 3 redirected interface will NOT accept the IGMP state update from the MCS DB unless the Layer 3 interface is a non-querier.
- In case that the Layer 3 redirected interface is part of MC-LAG, the IGMP state update sent to it via MCS DB will be accepted only during the transitioning phase from standby to active MC-LAG state.

Briefly, IGMP states on Layer 3 interface are not VRRP aware. However, they are MC-LAG aware.





---

# Triple Play Multicast Command Reference

---

## Command Hierarchies

- [MCAST Management Commands on page 669](#)
- [Multicast Info Policy Bundle Commands on page 671](#)
- [Triple Play Multicast Service Commands on page 673](#)
- [Ingress Multicast Path Management Commands on page 676](#)
- [Show Commands on page 679](#)
- [Clear Commands on page 680](#)
- [Debug Commands on page 680](#)

## MCAST Management Commands

```

config
— mcast-management
    — bandwidth-policy policy-name [create]
    — no bandwidth-policy policy-name
        — admin-bw-threshold kilo-bits-per-second
        — no admin-bw-threshold
        — ancillary-path
            — path-limit megabits-per-second
            — no path-limit
            — queue-parameters
                — cbs percentage
                — no cbs
                — hi-priority-only percent-of-mbs
                — no hi-priority-only
                — mbs percentage
                — no mbs
        — description description
        — no description
        — falling-percent-reset percent-of-highest
        — no falling-percent-reset
        — mcast-pool percent-of-total percent-of-buffers resv-cbs percent-of-pool slope-policy policy-name
        — no mcast-pool
— chassis-level
    — [no] mmrp-imp-override
    — [no] per-mcast-plane-capacity
        — mcast-capacity primary-percentage secondary secondary-percentage
        — no mcast-capacity
  
```

- **redundant-mcast-capacity** *primary-percentage* **secondary** *secondary-percentage*
- **no redundant-mcast-capacity**
- **total-capacity** *capacity*
- **no total-capacity**
- **[no] round-robin-inactive-records**

## Multicast Info Policy Bundle Commands

```

config
— mcast-management
— multicast-info-policy policy-name [create]
— no multicast-info-policy policy-name
— bundle bundle-name [create]
— no bundle bundle-name
— admin-bw kbps
— no admin-bw
— bw-activity {use-admin-bw|dynamic [falling-delay seconds]} [black-hole-rate kbps]
— no bw-activity
— channel ip-address [ip-address] [create]
— no channel ip-address [ip-address]
— admin-bw kbps
— no admin-bw
— bw-activity {use-admin-bw|dynamic [falling-delay seconds]} [black-hole-rate kbps]
— no bw-activity
— explicit-sf-path {primary | secondary | ancillary}
— no explicit-sf-path
— keepalive-override keepalive-timer
— no keepalive-override
— preference preference-level
— no preference
— primary-tunnel-interface rsvp-p2mplsp-name sender ip-address
— no primary-tunnel-interface
— source-override ip-address [create]
— no source-override ip-address
— admin-bw kbps
— no admin-bw
— bw-activity {use-admin-bw|dynamic [falling-delay seconds]} [black-hole-rate kbps]
— no bw-activity
— cac-type {mandatory | optional}
— no cac-type
— explicit-sf-path {primary | secondary | ancillary}
— no explicit-sf-path
— keepalive-override keepalive-timer
— no keepalive-override
— preference preference-level
— no preference
— primary-tunnel-interface rsvp-p2mplsp-name sender ip-address
— no primary-tunnel-interface
— congestion-priority-threshold preference-level
— no congestion-priority-threshold
— description description
— no description
— ecmp-opt-threshold preference-level
— no ecmp-opt-threshold

```

- **explicit-sf-path** {primary | secondary | ancillary}
- **no explicit-sf-path**
- **keepalive-override** *keepalive-timer*
- **no keepalive-override**
- **preference** *preference-level*
- **no preference**
- **primary-tunnel-interface** **rsvp-p2mplsp-name** **sender** *ip-address*
- **no primary-tunnel-interface**
- **description** *description-string*
- **no description**

Refer to the SR OS Interface Configuration Guide for command descriptions.

config

- **card**
  - **fp** [*fp-number*]
    - **ingress**
      - **multicast-path-management**
        - **bandwidth-policy** *policy-name*
        - **no bandwidth-policy**
        - [no] **shutdown**
- **mda**
  - **ingress**
    - **multicast-path-management**
      - **ancillary-override**
        - **path-limit** *megabits-per-second*
        - **no path-limit**
      - **bandwidth-policy** *policy-name*
      - **no bandwidth-policy**
      - **primary-override**
        - **path-limit** *megabits-per-second*
        - **no path-limit**
      - **secondary-override**
        - **path-limit** *megabits-per-second*
        - **no path-limit**
    - [no] **shutdown**

## Triple Play Multicast Service Commands

```

config
  — service
    — [no] vpls service-id [customer customer-id] [vpn vpn-id] [m-vpls]
      — igmp-snooping
        — mvr
          — description description-string
          — no description
          — group-policy policy-name
          — no group-policy
          — [no] shutdown
        — query-interval seconds
        — no query-interval
        — query-src-ip ipv6-address
        — no query-src-ip
        — report-src-ip ipv6-address
        — no report-src-ip
        — robust-count robust-count
        — no robust-count
        — [no] shutdown
      — fdb-table-high-wmark high-water-mark
      — no fdb-table-high-wmark
      — fdb-table-low-wmark low-water-mark
      — no fdb-table-low-wmark
      — fdb-table-size table-size
      — no fdb-table-size
      — mld-snooping
        — mvr
          — description description-string
          — no description
          — group-policy policy-name
          — no group-policy
          — [no] shutdown
        — query-interval seconds
        — no query-interval
        — query-src-ip ipv6-address
        — no query-src-ip
        — report-src-ip ipv6-address
        — no report-src-ip
        — robust-count robust-count
        — no robust-count
        — [no] shutdown
      — multicast-info-policy policy-name
      — no multicast-info-policy
      — sap sap-id [split-horizon-group group-name]
      — no sap sap-id
        — igmp-snooping
          — [no] fast-leave
          — import policy-name
          — no import
          — last-member-query-interval interval
          — no last-member-query-interval

```

```

— max-num-groups max-num-groups
— no max-num-groups
— [no] mrouter-port
— mvr
    — from-vpls vpls-id
    — no from-vpls
    — to-sap sap-id
    — no to-sap
— query-interval interval
— no query-interval
— query-response-interval interval
— no query-response-interval
— robust-count count
— no robust-count
— [no] send-queries
— static
    — [no] group group-address
    — [no] source ip-addr
    — [no] starg
— mesh-sdp sdp-id[:vc-id] [vc-type {ether | vlan | vpls}]
— no mesh-sdp sdp-id[:vc-id]
    — igmp-snooping
        — [no] fast-leave
        — import policy-name
        — no import
        — last-member-query-interval interval
        — no last-member-query-interval
        — max-num-groups max-num-groups
        — no max-num-groups
        — mcac
            — policy policy-name
            — no policy
            — unconstrained-bw bandwidth mandatory-bw mandatory-bw
            — no unconstrained-bw
        — query-interval interval
        — no query-interval
        — query-response-interval interval
        — no query-response-interval
        — robust-count count
        — no robust-count
        — [no] send-queries
        — static
            — [no] group group-address
            — [no] source ip-addr
            — [no] starg
    — mld-snooping
        — [no] disable-router-alert-check
        — [no] fast-leave
        — import policy-name
        — no import
        — last-member-query-interval interval
        — no last-member-query-interval
        — max-num-groups max-num-groups
        — no max-num-groups

```

```

— mvr
  — fast-leave service-id
  — no fast-leave
  — to-sap sap-id
  — no to-sap
— query-interval seconds
— no query-interval
— query-response-interval seconds
— no query-response-interval
— robust-count robust-count
— no robust-count
— [no] send-queries
— spoke-sdp sdp-id[:vc-id] [vc-type {ether | vlan | vpls}] [split-horizon-group
group-name]
— no spoke-sdp sdp-id[:vc-id]
  — igmp-snooping
    — [no] fast-leave
    — import policy-name
    — no import
    — last-member-query-interval interval
    — no last-member-query-interval
    — max-num-groups max-num-groups
    — no max-num-groups
    — query-interval interval
    — no query-interval
    — query-response-interval interval
    — no query-response-interval
    — robust-count count
    — no robust-count
    — [no] send-queries
    — static
      — [no] group group-address
      — [no] source ip-addr
      — [no] starg
  — mvr
    — fast-leave service-id
    — no fast-leave
    — to-sap sap-id
    — no to-sap
  — query-interval seconds
  — no query-interval
  — query-response-interval seconds
  — no query-response-interval
  — robust-count robust-count
  — no robust-count
  — [no] send-queries

```

## Ingress Multicast Path Management Commands

```

config
— mcast-management
— bandwidth-policy policy-name [create]
— no bandwidth-policy policy-name
— admin-bw-threshold kilo-bits-per-second
— no admin-bw-threshold
— ancillary-path
— path-limit megabits-per-second
— no path-limit
— queue-parameters
— cbs percentage
— no cbs
— hi-priority-only percent-of-mbs
— no hi-priority-only
— mbs percentage
— no mbs
— description description-string
— no description
— falling-percent-reset percent-of-highest
— no falling-percent-reset
— mcast-pool percent-of-total percent-of-buffers resv-cbs percent-of-pool slope-policy policy-name
— no mcast-pool
— primary-path
— path-limit megabits-per-second
— no path-limit
— queue-parameters
— cbs percentage
— no cbs
— hi-priority-only percent-of-mbs
— no hi-priority-only
— mbs percentage
— no mbs
— secondary-path
— path-limit megabits-per-second
— no path-limit
— queue-parameters
— cbs percentage
— no cbs
— hi-priority-only percent-of-mbs
— no hi-priority-only
— mbs percentage
— no mbs
— t2-paths
— primary-path
— queue-parameters
— cbs percentage
— no cbs
— hi-priority-only percent-of-mbs
— no hi-priority-only
— mbs percentage
— no mbs
— secondary-path

```



- **number-paths** *number-of-paths* [**dual-sfm** *number-of-paths*]
- **queue-parameters**
  - **cbs** *percentage*
  - **no cbs**
  - **hi-priority-only** *percent-of-mbs*
  - **no hi-priority-only**
  - **mbs** *percentage*
  - **no mbs**

**config**

- **router**
  - **mcast-info-policy** *policy-name*
  - **no mcast-info-policy**
  - **pim**
    - [**no**] **mc-ecmp-balance**
    - **mc-ecmp-balance-hold** *minute*
    - **no mc-ecmp-balance-hold**

**tools**

- **perform**
  - **router**
    - **pim**
      - **mc-ecmp-rebalance** [*ecmp-opt-threshold*]

## Multicast Redirection

```
config
— router/service VPRN
    — policy-options
        — policy-statement
            — entry
                — action {accept|next-entry|next-policy|reject}
                — no action
        — igmp
            — interface port:tags
            — mld
                — interface port:tags
config
— service
    — ies service-id
    — vprn service-id
        — interface port:tags
```

## Show Commands

```

show
  — service
    — id service-id
      — igmp-snooping
        — all
        — mroute [detail]
        — mvr
        — port-db sap sap-id [detail]
        — port-db sap sap-id group grp-address
        — port-db sdp sdp-id:vc-id [detail]
        — port-db sdp sdp-id:vc-id group grp-address
        — proxy-db [detail]
        — proxy-db group grp-address
        — static [sap sap-id | sdp sdp-id:vc-id]
        — static [sap sap-id | sdp sdp-id:vc-id]
      — mfib brief
      — mfib [group grp-address]
      — mfib statistics [group grp-address]
      — mld-snooping
        — all
        — base
        — mroute [detail]
        — mvr
        — port-db sap sap-id
        — port-db sap sap-id detail
        — port-db sap sap-id group grp-ipv6-address
        — port-db sdp sdp-id:vc-id
        — port-db sdp sdp-id:vc-id detail
        — port-db sdp sdp-id:vc-id group grp-ipv6-address
        — proxy-db [detail]
        — proxy-db group grp-ipv6-address
        — querier
        — static [sap sap-id | sdp sdp-id:vc-id]
        — statistics [sap sap-id | sdp sdp-id:vc-id]

show
  — router
    — pim
      — mc-ecmp-balance [detail]

show
  — mcast-management
    — bandwidth-policy [policy-name] [detail]
    — channel [router router-instance | vpls service-id] [mda slot[/mda]] [group ip-address [source
      ip-address]] [path path-type] [detail]
    — mda [slot[/mda]] [path path-type]

show
  — router
    — igmp
      — group [grp-ip-address]
      — group summary
      — group-interface [fwd-service service-id] [ip-int-name] [detail]

```

```

— hosts [group grp-address] [detail] [fwd-service service-id] [grp-interface ip-int-
name]
— hosts [host ip-address] [group grp-address] [detail]
— hosts summary
— interface [ip-int-name|ip-address] [group] [grp-ip-address] [detail]
— mld
    — interface port:tags
— ssm-translate [interface-name]
— static [ip-int-name|ip-addr]
— statistics [ip-int-name|ip-address]
— statistics host [ip-address]
— status
— tunnel-interface

```

## Clear Commands

```

clear
  — service
    — id
      — igmp-snooping
          — port-db sap sap-id [group grp-address [source ip-address]]
          — port-db sdp sdp-id:vc-id [group grp-address [source ip-address]]
          — querier
          — querier [all | sap sap-id | sdp sdp-id:vc-id]
      — mfib
          — statistics [all | group grp-address ]
      — mld-snooping
          — port-db sap sap-id [group grp-ipv6-address]
          — port-db sap sap-id group grp-ipv6-address source src-ipv6-address
          — port-db sdp sdp-id:vc-id [group grp-ipv6-address]
          — port-db sdp sdp-id:vc-id group grp-ipv6-address source src-ipv6-address
          — querier
          — statistics all
          — statistics sap sap-id
          — statistics sdp sdp-id:vc-id

```

## Debug Commands

```

debug
  — service
    — id service-id
      — mld-snooping
          — detail-level {low|medium|high}
          — no detail-level
          — [no] mac ieee-address
          — mode {dropped-only|ingr-and-dropped|egr-ingr-and-dropped}
          — no mode
          — [no] sap sap-id
          — [no] sdp sdp-id:vc-id

```

---

# Multicast Management Configuration Commands

---

## Generic Commands

### description

<b>Syntax</b>	<b>description</b> <i>description-string</i> <b>no description</b>
<b>Context</b>	config>card>mda>ingress>mcast-mgmt config>card>mda>ingress>mcast-path-management config>mcast-management>bandwidth-policy config>mcast-management>multicast-info-policy config>mcast-management>multicast-info-policy>bundle config>mcast-mgmt>mcast-rprt-dest config>mcast-mgmt>mcast-info-plcy>bundle config>mcast-mgmt>mcast-info-plcy config>card>fp>ingress>mcast-mgmt config>card>mda>ingress
<b>Description</b>	This command creates a text description stored in the configuration file for a configuration context.  The <b>description</b> command associates a text string with a configuration context to help identify the context in the configuration file.  The <b>no</b> form of this command removes any description string from the context.
<b>Default</b>	No description is associated with the configuration context.
<b>Parameters</b>	<i>description-string</i> — A text string describing the entity. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters excluding double quotes. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

### shutdown

<b>Syntax</b>	<b>[no] shutdown</b>
<b>Context</b>	config>card>mda>ingress>mcast-mgmt config>card>mda>ingress>mcast-path-management config>mcast-mgmt>mcast-rprt-dest
<b>Description</b>	The shutdown command is used to disable ingress multicast path management for a forwarding plane or MDA. The default state is shutdown. When the no shutdown command is executed, ingress multicast path management is enabled and the system will evaluate all multicast channels on the MDA or forwarding plane based on the default or explicit bandwidth policy applied and multicast information policies associated with the channels.

## Generic Commands

The **no** form of the command is used to enable ingress multicast path management on the MDA or forwarding plane.

---

# Multicast Management Commands

## mcast-management

<b>Syntax</b>	<b>mcast-management</b>
<b>Context</b>	config
<b>Description</b>	<p>The mcast-management CLI node contains the bandwidth-policy and multicast-info-policy definitions. The bandwidth-policy is used to manage the ingress multicast paths into the switch fabric. The multicast-info-policy defines how each multicast channel is handled by the system. The policy may be used by the ingress multicast bandwidth manager, the ECMP path manager and the egress multicast CAC manager.</p> <p>The mcast-management node always exists and contains the default bandwidth-policy and the default multicast-info-policy. Enter the mcast-management node when editing, deleting or creating a bandwidth-policy or multicast-info-policy. The default bandwidth-policy and multicast-info-policy cannot be edited or deleted.</p> <p>A chassis-level node within multicast-management is used to control the switch fabric multicast planes replication limits. The switch fabric multicast planes are the individual multicast spatial replication contexts available in the system.</p>

## Bandwidth Policy Commands

### bandwidth-policy

<b>Syntax</b>	<b>bandwidth-policy</b> <i>policy-name</i> [create] <b>no bandwidth-policy</b> <i>policy-name</i>
<b>Context</b>	config>mcast-mgmt
<b>Description</b>	This command creates a multicast bandwidth policy. Bandwidth policies are used to manage the ingress multicast path bandwidth. Each forwarding plane supports multicast forwarding paths into the switch fabric. By default, two paths are available; the multicast high priority path and the multicast low priority path. Multicast packets are forwarded on either path based on the expedited or non-expedited (best-effort) nature of the queue the packets are scheduled from. The ingress forwarding plane uses the classification rules to determine the forwarding class of each multicast packet and uses the forwarding class to queue mapping to decide which ingress multipoint queue will forward the packet. When multicast path management has been enabled on an ingress forwarding plane, the multicast bandwidth manager adds a third path for ingress multicast forwarding (ancillary path) and changes the way multicast packets are mapped to the three paths. This new forwarding plane behavior only applies to Layer 2 snooped or Layer 3 routed IP multicast forwarding. VPLS broadcast and unknown or non-snooped flooding is not affected.

When multicast path management is enabled, the ingress forwarding plane allows IP multicast snooped or routed packets to be placed on to the three multicast paths independently of the ingress classification rules. The high priority multicast path is treated as the primary path and the low priority multicast path is treated as the secondary path. The ancillary path is the point-to-point bandwidth unused by switch fabric point-to-point traffic. The ingress bandwidth manager evaluates each multicast FIB (M-FIB) record to determine which path is best based on ingress bandwidth, number of switch fabric destinations and the fill level of each path. Explicit path association is also supported.

#### Dynamic Bandwidth Activity Monitoring

When ingress multicast path management is enabled on an MDA, the system monitors the in-use bandwidth associated with each Layer 2 and Layer 3 ingress multicast record. When records are first populated by static, snooping or routing protocols, they are first assumed to be inactive. An inactive record is not considered to be currently consuming ingress multicast path bandwidth.

Within the multicast-info-policy, the bandwidth activity of the new record was configured to be either managed based on an administrative bandwidth, or based on the dynamic bandwidth rate table. The bandwidth-policy associated with ingress MDA contains the configuration parameters for creating the dynamic bandwidth rate table. The purpose of the table is to allow for the system to monitor the bandwidth activity associated with a multicast record and compare the current rate against a number of rate thresholds. Rate thresholds are used to allow a multicast streams rate to fluctuate between a given range while keeping the managed rate at a certain level. Multiple dynamic managed rates are supported in the table to allow monitoring of different types of multicast traffic. Each rate threshold is associated with a rising and falling threshold that defines when the specified rate should be used and when the next lower rate should be used.

Once a record's monitored current rate rises to the first dynamic rising threshold, the record is considered to be active and the system will then manage the bandwidth the record represents



based on the parameters associated with the record in the records multicast-info-policy and the configured path information in the MDAs associated bandwidth-policy.

#### Ingress Multicast Path Parameters

The bandwidth-policy also contains the configuration parameters for each of the managed ingress multicast paths. Each path may be configured with a path-limit rate used to override a specific paths default rate. Also, forwarding on each path is managed through an ingress path queue. The queue default parameters may be overridden for each path.

#### Default Bandwidth Policy

A bandwidth policy with the name 'default' always exists and is used as the default bandwidth policy when ingress multicast path management is enabled without an explicit bandwidth policy defined on an MDA. The default policy cannot be deleted or edited.

The **no** form of the command removes the specified bandwidth policy from the system. The bandwidth policy associations must be removed from MDA configurations before it can be removed.

#### Parameters

- policy-name** — Specifies the name of the bandwidth policy, up to 32 characters in length. Each bandwidth policy must be uniquely named within the system. 32 policies can be configured per system.
- create** — The create keyword is required if creating a new bandwidth policy when the system is configured to require the explicit use of the keyword to prevent accidental object creation. Objects may be accidentally created when this protection is disabled and an object name is mistyped when attempting to edit the object. This keyword is not required when the protection is disabled. The keyword is ignored when the bandwidth policy name already exists.

## admin-bw-threshold

<b>Syntax</b>	<b>admin-bw-threshold</b> <i>kilo-bits-per-second</i> <b>no admin-bw-threshold</b>
<b>Context</b>	config>mcast-management>bandwidth-policy
<b>Description</b>	This command defines at which bandwidth rate a multicast channel configured to use an administrative rate will start and stop using that rate as the in-use ingress bandwidth when managing ingress multicast paths. This parameter only applies to channels that are configured to use the admin-bw rate with the bw-activity use-admin-bw command (both are configured in the multicast-info-policy associated with the channel context).

To be effective, the admin-bw-threshold must be less than the channels configured admin-bw. If the administrative bandwidth configured on the channel is less than the administrative bandwidth threshold defined in the bandwidth policy, the admin-bw value is ignored for ingress multicast path management and the system continually uses the dynamic ingress bandwidth associated with the channel. Since the admin-bw-threshold is defined in the bandwidth-policy and the channel admin-bw value is defined in the multicast-info-policy, it is not possible to pre-determine that a given administrative bandwidth value is less than an administrative bandwidth threshold. Since a typical administrative bandwidth threshold will be set significantly lower than any administrative bandwidth values, this corner case is not expected to be prevalent. However, if the case does arise in a production environment, no ill behavior is expected as the threshold is simply a tuning parameter used to detect when the bandwidth associated with a channel has risen above any OAM or background type traffic.

While a channel that is configured to use-admin-bw (in the bw-activity command) current bandwidth is less than the admin-bw-threshold, the system treats the channel as a dynamic type channel. Once the threshold is crossed, the system immediately allocates the full admin-bw value to the channel and manages the ingress multicast path accordingly. If the bandwidth monitored on the channel rises above the admin-bw value, the system reverts to dynamic bandwidth management operation. If the bandwidth drops below the admin-bw value, but is above the admin-bw-threshold, the system uses the admin-bw value. If the bandwidth drops below the admin-bw-threshold, the system goes back to dynamic bandwidth management operation.

This command has no effect on multicast ECMP or egress CAC management operations.

The no for of the command restores the default threshold value of 10 Kbps.

*kilobits-per-second* — The kilobits-per-second parameter must follow the admin-bw-threshold command and defines rate at which channels configured to use administrative bandwidths change from dynamic bandwidth management to using the channels configured administrative bandwidth. The parameter is expressed as an integer value and represents multiples of 1,000 bits per second. A value of 3000 indicates 3,000,000 bits per second.

<b>Values</b>	1 — 40,000,000
<b>Default</b>	10

primary-path

<b>Syntax</b>	<b>primary-path</b>
<b>Context</b>	config>mcast-mgmt>bandwidth-policy
<b>Description</b>	This command enables the context to configure primary path parameters.

ancillary-path

<b>Syntax</b>	<b>ancillary-path</b>
<b>Context</b>	config>mcast-mgmt>bandwidth-policy
<b>Description</b>	This command overrides the default path limit for the ancillary path, which is one of the three ingress multicast paths into the switch fabric.

t2-paths

<b>Syntax</b>	<b>t2-paths</b>
<b>Context</b>	config>mcast-management>bandwidth-policy
<b>Description</b>	The t2-paths CLI node contains the primary and secondary path CLI nodes for IOM-3s. The commands within this context are ignored when the policy is applied to an IOM-1 or IOM-2.

## secondary-path

<b>Syntax</b>	<b>secondary-path</b>
<b>Context</b>	config>mcast-mgmt>bandwidth-policy config>mcast-mgmt>bw-plcy>t2-paths
<b>Description</b>	This command overrides the default path limit for the secondary path, which is one of the three ingress multicast paths into the switch fabric.

## number-paths

<b>Syntax</b>	<b>number-paths</b> <i>number-of-paths</i> [ <b>dual-sfm</b> <i>number-of-paths</i> ] config>mcast-management>bandwidth-policy>t2-paths>secondary-paths>queue-parameters
<b>Description</b>	This command is used to explicitly provision the number of secondary paths (and imply the number of primary paths) supported by the T2 TChip based forwarding plane the bandwidth policy is managing. The default (and minimum) number of secondary paths is 1 and the maximum configurable is 15. The number of primary paths is total number of available paths minus the number of secondary paths.

Secondary paths are used by:

- Expedited VPLS, IES and VPRN service ingress multipoint queues
- Expedited network ingress multipoint queues
- 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)

Secondary paths are used by:

- Best-Effort VPLS, IES and VPRN service ingress multipoint queues
- Best-Effort network ingress multipoint queues
- 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)

The number of secondary paths should be increased from the default value of 1 when a single secondary path is insufficient for the amount of explicit secondary path managed traffic or the amount of best-effort multipoint non-managed queue traffic.

The **no** form of the command restores the default number of secondary paths.

<b>Parameters</b>	<i>number-of-paths</i> — 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.
<b>Values</b>	1 — 15
<b>Default</b>	1

### path-limit

<b>Syntax</b>	<b>path-limit</b> <i>megabits-per-second</i> <b>no path-limit</b>
<b>Context</b>	config>card>mda>ingress>mcast-mgmt>prim-override config>card>mda>ingress>mcast-mgmt>sec-override config>card>mda>ingress>mcast-mgmt>anc-override config>mcast-mgmt>bandwidth-policy >primary-path config>mcast-mgmt>bandwidth-policy >secondary-path config>mcast-mgmt>bandwidth-policy >ancillary-path
<b>Description</b>	<p>This command is used to individually override the default path limit for each of the three ingress multicast paths into the switch fabric on an IOM-1 or IOM-2. When the bandwidth policy is applied to an IOM-3, the path-limit commands are ignored. The configured path limit may be less than, equal to or greater than the terminating switch fabric multicast plane limit.</p> <p>The system will not violate the configured path limit or the multicast plane limits. The system uses all bandwidth limits and record preferences when determining which records should be allowed on which path, which records should be moved and which records should be placed in the black-hole state.</p> <p>Changing a paths bandwidth limit causes the system to immediately reevaluate each record on the path. For dynamic path records, this could cause records to be moved to other paths and records on other paths to be placed in the black-hole state.</p> <p>Changing a path-limit value within the bandwidth policy affects all MDAs where the policy is applied. The policy derived path limit may be overridden on each MDA using the primary-override, secondary-override or ancillary-override nodes in the MDA context.</p> <p>The <b>no</b> form of the command returns the default path bandwidth limit for the ingress multicast path. The command has no effect for MDAs with a path override in effect for the given path.</p>
<b>Default</b>	no path-limit
<b>Parameters</b>	<i>megabits-per-second</i> — Specifies the primary path limit in megabits per seconds.
<b>Values</b>	1 — 5000
<b>Default</b>	primary-path: 2,000 secondary-path 1,500 ancillary-path: 5,000

## queue-parameters

<b>Syntax</b>	<b>queue-parameters</b>
<b>Context</b>	<pre>config&gt;mcast-management&gt;bandwidth-policy&gt;primary-path config&gt;mcast-management&gt;bandwidth-policy&gt;secondary-path config&gt;mcast-management&gt;bandwidth-policy&gt;ancillary-path config&gt;mcast-management&gt;bandwidth-policy&gt;t2-paths</pre>
<b>Description</b>	<p>This command defines the individual parameters for the queues through which multicast packets are forwarded into the switch fabric on each path.</p> <p>The individual path queues may be viewed as shared queues. All multicast packets forwarded through the switch fabric associated with one of the paths traverses bypass the normal queuing behavior. Instead of being forwarded through the normal service or network multicast queue, a single queue associated with the multicast path is used. In order to retain billing and diagnostic information, the forwarding and discard statistics for the service or network queue the packet would have traversed without ingress multicast management is used to account for each packets behavior.</p> <p>Note that any ingress scheduling policy functions attempting to manage the service or network multicast queues will only be able to read the statistics of the multicast queues and will not be able to manage the queues dynamic rate since the packets are flowing through different, non-managed queues. Since this is the case, multicast queues parented to a scheduling policy should be parented to the root scheduler at the highest priority without any rate limitation. Any ingress rate limiting for multicast traffic will be preformed by the multicast path bandwidth manger based on each records priority and a possible “black-hole” rate threshold.</p> <p>All queues created for ingress multicast path management are automatically created by the system out of the system reserved queue space. Each queue is created as an expedited queue.</p> <p>When forwarding through the queues, each packets forwarding class is ignored. However, the forwarding class is retained for proper egress processing. The packets expressed or implied profile is also ignored within the ingress path queues. A packets congestion priority is derived from the records cong-priority-theshold evaluation result as indicated by the policy. The cong-priority-theshold sets the high or low congestion priority of a record based on the records preference value. Within each multicast information policy bundle the cong-priority-theshold is set with a value from 0 to 7 and defines the threshold at which all records with a preference equal to or higher than the defined preference will be treated as congestion priority high. Multicast records with a preference lower than the defined class threshold will be treated as congestion priority low. Low priority packets use the low priority MBS threshold of the queue while high priority packets use the standard MBS value. In the event of path congestion, low priority packets are discarded first, leaving room for the higher priority packets.</p> <p>For the primary and secondary paths, a single queue exists for each path and every packet forwarded through the path by the bandwidth manager uses that queue. For the ancillary path, a single queue exists for each switch fabric destination. Ancillary path packets are replicated to each switch fabric destination. The replication process places a copy of the packet in the correct ancillary path queue that forwards to that destination.</p>

cbs

<b>Syntax</b>	<b>cbs</b> <i>percent-of-resv-cbs</i> <b>no cbs</b>				
<b>Context</b>	config>mcast-management>bandwidth-policy>primary-path>queue-parameters config>mcast-management>bandwidth-policy>secondary-path>queue-parameters config>mcast-management>bandwidth-policy>ancillary-path>queue-parameters>ancillary-path config>mcast-management>bandwidth-policy>t2-paths>primary-paths>queue-parameters config>mcast-management>bandwidth-policy>t2-paths>secondary-paths>queue-parameters				
<b>Description</b>	<p>This command overrides the default Committed Buffer Size (CBS) for each individual path's queue. The queue's CBS threshold is used when requesting buffers from the system's ingress buffer pool to indicate whether the requested buffer should be removed from the reserved portion of the buffer pool or the shared portion. When the queue's fill depth is below or equal to the CBS threshold, the requested buffer comes from the reserved portion. Once the queue's depth exceeds the CBS threshold, buffers come from the shared portion.</p> <p>The <b>cbs percent-of-resv-cbs</b> parameter is defined as a percentage of the reserved portion of the pool. The system allows the sum of all CBS values to equal more than 100% allowing for oversubscription of the reserved portion of the pool. If the reserved portion is oversubscribed and the queues are currently using more reserved space than provisioned in the pool, the pool automatically starts using the shared portion of the pool for within-CBS buffer allocation. On the shared early detection slopes could assume more buffers exist within the shared portion than actually do which may cause the early detection function to fail.</p> <p>For the primary-path and secondary-path queues, the percentage is applied to a single queue for each path. For the ancillary-path, multiple queues are required. The <b>cbs percent-of-resv-cbs</b> is divided between all ancillary path queues. For example, if primary is given 10 percent, the primary path queue's CBS is set to 10% of the reserved portion of the buffer pool. If ancillary is given 50 percent and 10 ancillary path queues exist, each ancillary queue gets 5%. The number of path queues required depends on the available chassis slots and the provisioned IOM types. 7450 ESS-6 chassis have 4 slots, 7450 ESS-7 and 7750 SR-7 have 5 and the 7450 ESS-12 and 7750 SR-12 have 10. An IOM-1 or IOM-2 on a slot has two switch fabric destinations.</p> <p>The <b>no</b> form of the command restores the path queues' default CBS value.</p>				
<b>Parameters</b>	<p><i>percent-of-resv-cbs</i> — The percent of buffers reserved from the total buffer pool space, expressed as a decimal integer. If 10 MB is the total buffers in the buffer pool, a value of 10 would reserve 1 MB (10%) of buffer space for the forwarding class queue. The value 0 specifies that no reserved buffers are required by the queue (a minimal reserved size can be applied for scheduling purposes).</p> <table><tr><td><b>Values</b></td><td>0 — 100</td></tr><tr><td><b>Default</b></td><td><b>Primary:</b> 5 <b>Secondary:</b> 30 <b>Ancillary:</b> 65</td></tr></table>	<b>Values</b>	0 — 100	<b>Default</b>	<b>Primary:</b> 5 <b>Secondary:</b> 30 <b>Ancillary:</b> 65
<b>Values</b>	0 — 100				
<b>Default</b>	<b>Primary:</b> 5 <b>Secondary:</b> 30 <b>Ancillary:</b> 65				

## hi-priority-only

<b>Syntax</b>	<b>hi-priority-only</b> <i>percent-of-mbs</i> <b>no hi-priority-only</b>
<b>Context</b>	config>mcast-management>bandwidth-policy>primary-path>queue-parameters config>mcast-management>bandwidth-policy>secondary-path>queue-parameters config>mcast-management>bandwidth-policy>ancillary-path>queue-parameters>ancillary-path config>mcast-management>bandwidth-policy>t2-paths>primary-paths>queue-parameters config>mcast-management>bandwidth-policy>t2-paths>secondary-paths>queue-parameters
<b>Description</b>	<p>This command overrides the default percentage value used to determine the low priority MBS value for the queue. By default, 10 percent of the queue depth is reserved for high congestion priority traffic. When specified, the high-prior-only percentage value is applied to the queues MBS threshold. The resulting value is subtracted from the MBS to derive the low priority MBS threshold maintained by the queue. The low priority MBS threshold defines the point at which all low congestion priority packets destined for the queue will be discarded based on queue depth. Low and high congestion priority is derived from the multicast records preference value compared to the record's bundle priority-threshold.</p> <p>The <b>no</b> form of this command restores the default value.</p>
<b>Parameters</b>	<i>percent-of-mbs</i> — The percentage parameter is required when specifying high-prior-only.
<b>Values</b>	0 — 100
	0 specifies that the MBS and LP-MBS thresholds will be set to the same value resulting in high and low congestion priority packets being treated equally.
	100 specifies that the LP-MBS threshold will be set to 0, resulting in a discarded of all low congestion priority packets.
	Values in between 0 and 100 result in a corresponding differential between the MBS and LP-MBS threshold values.

## mbs

<b>Syntax</b>	<b>mbs</b> <i>percent-of-pool</i> <b>no mbs</b>
<b>Context</b>	config>mcast-management>bandwidth-policy>primary-path>queue-parameters config>mcast-management>bandwidth-policy>secondary-path>queue-parameters config>mcast-management>bandwidth-policy>ancillary-path>queue-parameters>ancillary-path config>mcast-management>bandwidth-policy>t2-paths>primary-paths>queue-parameters config>mcast-management>bandwidth-policy>t2-paths>secondary-paths>queue-parameters
<b>Description</b>	<p>This command is used to override the default Maximum Buffer Size (MBS) for each individual path in this queue. The queues MBS threshold defines the point at which all packets destined for the queue will be discarded based on queue depth. The defined threshold also provides context for the queues high-prior-only parameter.</p>

The *mbs percent-of-pool* parameter is defined as a percentage of the total pool size. The system allows the sum of all MBS values to equal more than 100% allowing for oversubscription of the pool.

For the primary-path and secondary-path queues, the mbs percent is applied to a single queue for each path. For the ancillary-path, multiple queues are required. The mbs percentage is divided between all ancillary path queues. For example, if primary is given 10 percent, the primary path queues MBS is set to 10% of the buffer pool. If ancillary is given 50 percent and 10 ancillary path queues exist, each ancillary queue gets 5%. The number of path queues required depends on the available chassis slots and the provisioned IOM types. ESS-6 chassis have 4 slots, ESS/SR-7 have 5 and ESS/SR-12 have 10. An IOM-1 or IOM-2 on a slot has two switch fabric destinations while an IOM-3 has one.

The **no** form of the command is used to restore the path queues default MBS value.

<b>Parameters</b>	<i>percent-of-pool</i> — The percent of buffers from the total buffer pool space for the maximum amount of buffers, expressed as a decimal integer. If 10 MB is the total buffers in the buffer pool, a value of 10 would limit the maximum queue size to 1MB (10%) of buffer space for the forwarding class queue. If the total size is increased to 20MB, the existing value of 10 would automatically increase the maximum size of the queue to 2MB.		
	<b>Values</b>	0 — 100	
	<b>Default</b>	<b>Primary Default:</b>	7
		<b>Secondary Default:</b>	40
		<b>Ancillary Default:</b>	80
<b>Default</b>	10		



## falling-percent-reset

<b>Syntax</b>	<b>falling-percent-reset</b> <i>percent-of-highest</i> <b>no falling-percent-reset</b>
<b>Context</b>	config>mcast-mgmt>bw-plcy
<b>Description</b>	<p>The falling-percent-reset command is used to configure the percentage of bandwidth decrease that must occur to reset the dynamic bandwidth monitoring function for a multicast channel. When a channel is configured to use the ingress dynamic bandwidth as the in-use bandwidth for ingress multicast path management, the system maintains a sliding window in time that defines how long the last highest bandwidth value associated with the channel should be used. The sliding window duration is derived from the channels bw-activity dynamic falling-delay parameter within the multicast information policy. Each time the system detects a current bandwidth for a channel that is equal to or greater than the current highest bandwidth for the channel, the sliding window is reset and the highest value is used when managing the ingress multicast paths. If the system does not detect a higher or equal bandwidth value for the channel within the window period, the system resets the sliding window and uses the next highest rate seen during the duration of the window period. In this way, the system delays relinquishing bandwidth for a dynamic bandwidth channel for a configurable period of time. If a momentary fluctuation (decrease) in ingress bandwidth occurs, the system ignores the bandwidth change.</p> <p>While this is useful for momentary fluctuations in bandwidth, it may be desirable to react faster when the current bandwidth monitored for a channel drops significantly relative to the currently in-use bandwidth. When the bandwidth decrease is equal to or greater than the falling-percent-reset value, the system immediately stops using the highest bandwidth and starts using the current bandwidth while resetting the sliding window.</p> <p>If falling-percent-reset is set to 50%, when the current ingress dynamic bandwidth is 50% of the current in-use highest bandwidth, the system will immediately use the current dynamic ingress bandwidth as the highest bandwidth for the channel.</p> <p>By default falling-percent-reset is 50% when a new bandwidth policy is created. The default bandwidth policy also has a hard configured value of 50%. Setting falling-percent-reset to 100 is equivalent to specifying no falling-percent-reset.</p> <p>The <b>no</b> form of the command restores the default value of 50%.</p>
<b>Parameters</b>	<p><i>percent-of-highest</i> — The percent-of-highest parameter is required and defines the percentage of decline between the current ingress dynamic bandwidth and the current in-use highest bandwidth at which the system will reset the dynamic ingress bandwidth monitoring for the channel. When reset in this case, the system uses the current ingress dynamic bandwidth as the highest rate and continues monitoring. The parameter must be defined as an integer value representing a percentage.</p> <p><b>Values</b>      1 — 100 percent</p> <p><b>Default</b>     100</p>

### Sample Output

The following output displays an example of bandwidth policy defaults.

```
*A:PE-1# configure mcast-management bandwidth-policy test create
```

## Bandwidth Policy Commands

```
*A:PE-1>config>mcast-mgmt>bw-plcy$ exit all
*A:PE-1# show mcast-management bandwidth-policy "test" detail
=====
Bandwidth Policy Details
=====
-----
Policy                : test
-----
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#
```

## mcast-pool

<b>Syntax</b>	<b>mcast-pool percent-of-total percent-of-buffers resv-cbs percent-of-pool slope-policy policy-name</b> <b>no mcast-pool</b>
<b>Context</b>	config>mcast-mgmt>bw-plcy
<b>Description</b>	This command configures the ingress multicast path management buffer pool. The pool is used by the primary, secondary and ancillary path queues through which all ingress managed multicast traffic must flow. The parameters may be used to configure the size of the pool relative to the total ingress buffer space, the amount of reserved CBS buffers within the pool and the slope policy used to manage early congestion detection functions in the shared portion of the pool.

Care should be taken when managing the buffer pool space as changes to the systems buffer pool behavior can have negative effects on multicast and unicast forwarding.

### Sizing the Pool

The percent-of-total command defines how much of the total ingress buffer pool space for the MDA is dedicated for multicast channels managed by the bandwidth policy. Since multicast typically has a higher scheduling priority through the switch fabric, the buffer pool does not need to be large. By default, the system reserves 10% of the buffers on the ingress side of the MDA once multicast path management is enabled.

### Reserved CBS Portion of the Pool

The multicast pool is divided into two portions; reserved and shared. The reserved portion is used by the multicast path queues until they cross their individual CBS thresholds. Since the CBS thresholds are configured as percents and the percents are allowed to oversubscribe the reserved portion of the pool, it is possible for some of the queues CBS buffer allocation to be met by the shared portion of the pool. By default, 50% of the pool is defined as reserved. This may be changed using the `resv-cbs` percentage parameter.

### Shared Portion WRED Slopes

The shared portion of the buffer pool is used by queues that have crossed over their CBS thresholds. Since the total MBS values for the multicast path queues may oversubscribe the pool size, a buffer congestion control mechanism is supported within the pool in the form of two WRED slopes. The `slope-policy` parameter defines how the slopes are configured and whether they are activated. Each packet entering a path queue is defined as high or low priority within the queue based on the channels preference value relative to the `cong-priority-threshold` command. When getting a shared buffer of a high priority packet, the high WRED slope is used. Low priority packets use the low WRED slope.

The **no** form of the command returns the managed multicast path pool to its default settings.

### Parameters

**percent-of-total** *percent-of-buffers* — The `percent-of-total` keyword is required when executing the `mcast-pool` command and must be followed by a `percent-of-buffers` parameter expressed as an integer and representing the percentage of ingress buffers that will be allocated to the multicast pool.

**Values** 1 — 50

**Default** 10

**resv-cbs** *percent-of-pool* — The `resv-cbs` keyword is required when executing the `mcast-pool` command and must be followed by a `percent-of-pool` parameter expressed as an integer and representing the percentage of the pool that will be reserved for multicast path queues within their CBS threshold.

**Values** 1 — 100

**Default** 50

**slope-policy** *slope-policy-name* — The `slope-policy` keyword is required when executing the `mcast-pool` command and must be followed by a valid `slope-policy-name`. The named policy will be used to configure the WRED slopes within the multicast pool. Once a slope policy is associated with a buffer pool, it cannot be deleted.

**Default** default

# Multicast Info Policy Commands

## multicast-info-policy

<b>Syntax</b>	<b>multicast-info-policy</b> <i>policy-name</i> [create] <b>no multicast-info-policy</b>
<b>Context</b>	config>mcast-management
<b>Description</b>	This command configures a multicast information policy. Multicast information policies are used to manage parameters associated with Layer 2 and Layer 3 multicast records. Multiple features use the configured information within the policy. The multicast ingress path manager uses the policy to decide the inactive and active state behavior for each multicast record using the ingress paths to the switch fabric. The system's multicast ECMP join decisions are influenced by the channel information contained within the policy.

### Multicast Bundles:

A multicast information policy consists of one or multiple named bundles. Multicast streams are mapped to a bundle based on matching the destination address of the multicast stream to configured channel ranges defined within the bundles. Each policy has a bundle named 'default' that is used when a destination address does not fall within any of the configured channel ranges.

Each bundle has a set of default parameters used as the starting point for multicast channels matching the bundle. The default parameters may be overridden by optional exception parameters defined under each channel range. Further optional parameter overrides are possible under explicit source address contexts within each channel range.

### Default Multicast Information Policy

A multicast information policy always exists with the name 'default' and cannot be edited or deleted. The following parameters are contained in the default multicast information policy:

Policy Description: Default policy, cannot be edited or deleted.  
 Bundle: default  
 Bundle Description: Default Bundle, cannot be edited or deleted.  
 Congestion-Priority-Threshold: 4  
 ECMP-Optimization-Limit-Threshold: 7

### Bundle Defaults:

Administrative Bandwidth: 0 (undefined)  
 Preference: 0  
 CAC-Type: Optional  
 Bandwidth Activity: Dynamic with no black-hole rate  
 Explicit Ingress SF Path: None (undefined)  
 Configured Channel Ranges: None

The default multicast information policy is applied to all VPLS and VPRN services and all routing contexts until an explicitly defined multicast information policy has been mapped.

### Explicit Multicast Information Policy Associations

Each VPLS service and each routing context (including VPRN routing contexts) supports an explicit association with an pre-existing multicast information policy. The policy may need to be unique per service or routing context due to the fact that each context has its own multicast

address space. The same multicast channels may be and most likely will be used for completely different multicast streams and applications in each forwarding context.

#### Interaction with Ingress Multicast Path Management

When ingress multicast path management is enabled on an MDA, the system automatically creates a bandwidth manager context that manages the multicast path bandwidth into the switch fabric used by the ingress ports on the MDA. As routing or snooping protocols generate L2 or L3 multicast FIB records that will be populated on the MDA's forwarding plane, they are processed through the multicast information policy that is associated with the service or routing context associated with the record. The policy will return the following information for the record to be used by the ingress bandwidth manager:

- The records administrative bandwidth ('0' if undefined)
- Preference level (0 to 7 with 7 being highest)
- Bandwidth activity monitoring setting (use admin bw or dynamic monitoring)  
If admin bw is indicated, will also return active and inactive thresholds
- Initial switch fabric multicast path (primary, secondary or ancillary)  
If ancillary path is indicated, will also return an SF destination threshold
- Explicit switch fabric multicast path (primary, secondary, ancillary or none)

#### Interaction with Multicast ECMP Optimization

The multicast information policy is used by the multicast ECMP optimization function to derive each channels administrative bandwidth. The ECMP function tallies all bandwidth information for channels joined and attempts to equalize the load between the various paths to the sender. The multicast information policy returns the following information to the ECMP path manager:

3. Administrative bandwidth ('0' if undefined)
4. Preference (0 to 7 with 7 the highest preference value)

- Parameters**
- policy-name* — Identifies the name of the policy to be either created or edited. Each multicast information policy must be uniquely named within the system. Names of up to 32 ASCII characters are supported with the normal character restrictions.
- create** — The create keyword is required if creating a new multicast information policy when the system is configured to require the explicit use of the keyword to prevent accidental object creation. Objects may be accidentally created when this protection is disabled and an object name is mistyped when attempting to edit the object. This keyword is not required when the protection is disabled. The keyword is ignored when the multicast information policy name already exists.

## multicast-info-policy

- |                    |   |
|--------------------|---|
| <b>Syntax</b>      | <b>multicast-info-policy</b> <i>policy-name</i><br><b>no multicast-info-policy</b>  |
| <b>Context</b>     | config>service>vppls<br>config>router   |
| <b>Description</b> | This command overrides the default multicast information policy on a VPLS or routing context. When the policy association is changed, all multicast channels in the service or routing context must be reevaluated. |

## Multicast Info Policy Commands

If a multicast information policy is not explicitly associated with the VPLS service or routing context, the default multicast information policy is used when ingress multicast path management is enabled.

While a multicast information policy is associated with a service or routing context, the policy cannot be deleted from the system.

The **no** form of the command removes an explicit multicast information policy from the VPLS or routing context and restores the default multicast information policy.

**Parameters**     *policy-name* — The *policy-name* parameter is required and specifies an existing multicast information policy that should be associated with the VPLS service or routing context.

**Default**        default

## bundle

**Syntax**        **bundle** *bundle-name* [**create**]  
**no bundle** *bundle-name*

**Context**        config>mcast-mgmt>mcast-info-plcy

**Description**    The **bundle** command is used to create or edit channel bundles within a multicast information policy. Bundles are used for two main purposes. First, bundles are used by the multicast CAC function to group multicast channels into a common bandwidth context. The CAC function limits the ability for downstream nodes to join multicast channels based on the egress interfaces ability to handle the multicast traffic. Bundling allows multicast channels with common preference or application to be managed into a certain percentage of the available bandwidth.

The second function of bundles is to provide a simple provisioning mechanism. Each bundle within a multicast information policy has a set of default channel parameters. If each channel provisioned in to the bundle is able to use the default parameters for the bundle, the provisioning and configuration storage requirements are minimized.

Up to 31 explicit bundles may be defined within a multicast information policy (32 including the default bundle).

Once a bundle is created, the default channel parameters should be configured and the individual channel ranges should be defined. Within each channel range, override parameters may be defined that override the default channel parameters. Further overrides are supported within the channel range based on explicit source overrides.

A bundle may be deleted at anytime (except for the default bundle). When a bundle is deleted, all configuration information within the bundle is removed including multicast channel ranges. Any multicast records using the bundle should be reevaluated. Multicast CAC and ECMP managers should also be updated.

### Default Bundle

Each multicast information policy contains a bundle named **default**. The default bundle cannot be deleted. Any multicast channel that fails to match a channel range within an explicit bundle is automatically associated with the default bundle.

The **no** form of the command removes a bundle from the multicast information policy. The default bundle cannot be removed from the policy.

**Default** default

*bundle-name* — Specifies bundle expressed as an ASCII string with up to 16 characters and must follow normal naming conventions. If bundle-name already exists, the system will enter the bundle context for editing purposes. If bundle-name does not exist, the system will create the defined bundle in the policy and enter the bundle context for editing purposes.

**create** — The create keyword is required if creating a new multicast information policy bundle when the system is configured to require the explicit use of the keyword to prevent accidental object creation. Objects may be accidentally created when this protection is disabled and an object name is mistyped when attempting to edit the object. This keyword is not required when the protection is disabled. The keyword is ignored when the bundle name already exists.

## channel

**Syntax** **channel** *ip-address* [*ip-address*] [**create**]  
**no channel** *ip-address* [*ip-address*]

**Context** config>mcast-mgmt>mcast-info-plcy>bundle

**Description** This command defines explicit channels or channel ranges that are associated with the containing bundle. A channel or channel range is defined by their destination IP addresses. A channel may be defined using either IPv4 or IPv6 addresses. If a channel range is being defined, both the start and ending addresses must be the same type.

A specific channel may only be defined within a single channel or channel range within the multicast information policy. A defined channel range cannot overlap with an existing channel range.

If a channel range is to be shortened, extended, split or moved to another bundle, it must first be removed from its existing bundle.

Each specified channel range creates a containing context for any override parameters for the channel range. By default, no override parameters exist.

The **no** form of the command removes the specified multicast channel from the containing bundle.

**Parameters** **start-channel-ip-address** [**end-channel-ip-address**] — The start-channel-ip-address parameter and optional end-channel-ip-address parameters define the starting and ending destination IP addresses for a channel range.

If only the start-channel-ip-address is given, the channel ranges comprises of a single multicast channel.

If both the starting and ending address are specified, all addresses within the range including the specified address are part of the channel range.

IPv4 or IPv6 addresses may be defined. All specified addresses must be valid multicast destination addresses. The starting IP address must be numerically lower then the ending IP address.

**Values** Any valid IP multicast destination address

**Default** None

**create** — The create keyword is required if creating a new multicast channel range when the system is configured to require the explicit use of the keyword to prevent accidental object creation.

Objects may be accidentally created when this protection is disabled and an object name is mistyped when attempting to edit the object. This keyword is not required when the protection is disabled. The keyword is ignored when the specified channel range already exists.

### admin-bw

<b>Syntax</b>	<b>admin-bw</b> <i>kbps</i> <b>no admin-bw</b>
<b>Context</b>	config>mcast-mgmt>mcast-info-plcy>bundle>channel config>mcast-mgmt>mcast-info-plcy>bundle>source-override
<b>Description</b>	<p>This command specifies an administrative bandwidth for multicast channels.</p> <p>The specified bandwidth rate may be used by the multicast ingress path manger, multicast CAC manager or multicast ECMP manager.</p> <p>The admin-bw value is closely tied to the bw-activity command. When the bw-activity command is set to use-admin-bw, the multicast ingress path manager uses the configured administrative bandwidth value as the managed ingress bandwidth. The admin-bw value must be defined for the <b>bw-activity use-admin-bw</b> command to succeed. Once the bw-activity command is set to use the admin-bw value, the value cannot be set to 0 and the no admin-bw command will fail. Setting the bw-activity command to dynamic (the default setting), breaks the association between the commands.</p> <p>The <b>no</b> form of the command restores the default value for admin-bw. If the command is executed in the <b>channel</b> context, the channels administrative bandwidth value is set to null. If the command is executed in the <b>source-override</b> context, the source override administrative bandwidth value is set to null.</p>
<b>Default</b>	none
<b>Parameters</b>	<p><i>kbps</i> — Specifies the administrative bandwidth for multicast channels.</p> <p><b>Values</b>      1 — 40000000 kbps</p> <p>Bundle default:                      0</p> <p>Channel default:                      Null (undefined)</p> <p>Source-override default:              Null (undefined)</p> <p>Override sequence — The channel setting overrides the bundle setting. The source-override setting overrides the channel and bundle settings.</p>

### bw-activity

<b>Syntax</b>	<b>bw-activity</b> { <b>use-admin-bw</b>   <b>dynamic</b> [ <b>falling-delay</b> <i>seconds</i> ]} [ <b>black-hole-rate</b> <i>kbps</i> ] <b>no bw-activity</b>
<b>Context</b>	config>mcast-mgmt>mcast-info-plcy>bundle>channel config>mcast-mgmt>mcast-info-plcy>bundle>source-override
<b>Description</b>	<p>This command defines how the multicast ingress path manager determines the amount of bandwidth required by a multicast channel. The default setting is dynamic which causes the bandwidth manager to use the bandwidth policies dynamic rate table entries to determine the current rate. The alternative setting is use-admin-bw which causes the bandwidth manager to use the configured admin-bw</p>



associated with the channel. The **use-admin-bw** setting also requires an active and inactive threshold to be defined which allows the bandwidth manager to determine when the channel is actively using ingress path bandwidth and when the channel is idle.

The **use-admin-bw** setting requires that the channel be configured with an admin-bw value that is not equal to 0 in the same context as the **bw-activity** command using the setting. Once a context has use-admin-bw configured, the context's admin-bw value cannot be set to 0 and the no admin-bw command will fail.

This command also supports an optional **black-hole-rate** *kbps* command that defines at which current rate a channel should be placed in the black-hole state. This is intended to provide a protection mechanism against multicast channels that exceed a reasonable rate and cause outages in other channels.

The **no** form of the command restores the default bandwidth activity monitoring setting (dynamic or null depending on the context).

<b>Default</b>	no bw-activity
<b>Parameters</b>	<p><b>use-admin-bw   dynamic</b> — The <b>use-admin-bw</b> and <b>dynamic</b> keywords are mutually exclusive and one must be specified when executing the <b>bw-activity</b> command. The <b>use-admin-bw</b> keyword indicates the channels current ingress bandwidth should be derived from the <b>admin-bw</b> setting. The <b>admin-bw</b> setting must not currently be set to 0 for the <b>use-admin-bw</b> setting to succeed. The <b>dynamic</b> keyword indicates that the multicast ingress path manager should use the dynamic rate table (as defined in the bandwidth-policy) to derive the channels current ingress rate.</p> <p><b>falling-delay</b> <i>seconds</i> — specifies the value the bandwidth manager uses the falling-delay threshold to hold on to the previous highest bandwidth until the delay time has expired while operating in dynamic bandwidth mode. This allows the bandwidth manager to ignore momentary drops in channel bandwidth.</p> <p><b>Values</b> 10 — 3600</p> <p><b>Default</b> 30</p> <p>Bundle default: dynamic  Channel default: Null (undefined)  Source-override default: Null (undefined)</p> <p><b>black-hole-rate</b> <i>kbps</i> — Specifies a rate at which a channel will be placed in the black-hole state. The kilobits-per-second parameter is expressed as an integer and represents multiples of 1,000 bits per second.</p> <p><b>Values</b> 0 — 40000000</p> <p><b>Default</b> None</p>

## cac-type

<b>Syntax</b>	<b>cac-type {mandatory   optional}</b>
<b>Context</b>	config>mcast-mgmt>mcast-info-plcy>bundle config>mcast-mgmt>mcast-info-plcy>bundle>channel config>mcast-mgmt>mcast-info-plcy>bundle>channel>source-override
<b>Description</b>	This command defines the channels multicast CAC channel type, either <b>mandatory</b> or <b>optional</b> . The cac-type command is supported for future interaction with the egress multicast CAC manager policy.

The multicast CAC manager always reserves egress bandwidth for mandatory channels within a bundle, while optional channels are only given bandwidth when a join is received.

- Parameters**
- mandatory** — This keyword is mutually exclusive with the **optional** keyword and specifies the channels multicast CAC channel type.
  - optional** — This keyword is mutually exclusive with the **mandatory** keyword and specifies the channels multicast CAC channel type.

### explicit-sf-path

- Syntax** **explicit-sf-path {primary|secondary|ancillary}**  
**no explicit-sf-path**
- Context** config>mcast-mgmt>mcast-info-plcy>bundle>channel  
config>mcast-mgmt>mcast-info-plcy>bundle>source-override
- Description** This command defines an explicit ingress switch fabric multicast path assigned to a multicast channel. When defined, the channel is setup with the explicit path as its inactive path. When an explicit path is not defined, all multicast channels are initialized on the secondary path and when they start to consume bandwidth, they are moved to the appropriate path based on the channel attributes and path limitations. Explicit path channels are not allowed to move from their defined path.
- The **explicit-sf-path** command in the bundle context defines the initial path for all channels associated with the bundle unless the channel has an overriding explicit-sw-path defined in the channel context. The channel context may also be overridden by the explicit-sf-path command in the source-override context. The channel and source-override explicit-sf-path settings default to null (undefined) and have no effect unless explicitly set.
- The **no** form of the command restores default path association behavior (dynamic or null depending on the context).
- Default** no explicit-sf-path
- primary** — The primary, secondary and ancillary keywords are mutually exclusive to one another. One keyword must be specified when executing the explicit-sf-path command. The primary keyword specifies that the primary ingress multicast path should be used as the explicit path for the channel.
- secondary** — The primary, secondary and ancillary keywords are mutually exclusive to one another. One keyword must be specified when executing the explicit-sf-path command. The secondary keyword specifies that the secondary ingress multicast path should be used as the explicit path for the channel.
- ancillary** — The primary, secondary and ancillary keywords are mutually exclusive to one another. One keyword must be specified when executing the explicit-sf-path command. The ancillary keyword specifies that the ancillary ingress multicast path should be used as the explicit path for the channel.
- Default**
- |                 |                            |
|-----------------|----------------------------|
| Bundle :        | None (no explicit-sf-path) |
| channel:        | Null (undefined)           |
| Source-override | Null (undefined)           |

Override sequence — The channel setting overrides the bundle setting. The source-override setting overrides the channel and bundle settings.

## keepalive-override

<b>Syntax</b>	<b>keepalive-override</b> <i>keepalive-timer</i> <b>no keepalive-override</b>
<b>Context</b>	config>mcast-mgmt>mcast-info-policy>bundle config>mcast-mgmt>mcast-info-policy>bundle>channel config>mcast-mgmt>mcast-info-policy>bundle>channel>source-override
<b>Description</b>	<p>This command configures the keepalive timer override. The PIM (S,G) Keepalive Timer (KAT) is used to maintain the (S,G) state when (S,G) join is not received. Expiry of the KAT causes the (S,G) entry to be removed.</p> <p>The KAT override configuration is performed with an multicast information policy, which must be applied to the related PIM routing instance. When a KAT override is configured under a channel (a group or a group range), it applies to all (S,G) entries that fall under it, except when the source-override is configured and a KAT override is also configured under the source-override. In this scenario, the specific KAT override must be used for the (S,G) entries that fall under the source-override, while other (S,G) entries under the bundle use the KAT override configured under the channel.</p>
<b>Parameters</b>	<i>keepalive-timer</i> — Specifies the keepalive timer override, in seconds.
<b>Values</b>	10 — 300

## preference

<b>Syntax</b>	<b>preference</b> <i>preference-level</i> <b>no preference</b>
<b>Context</b>	config>mcast-mgmt>mcast-info-plcy>bundle config>mcast-mgmt>mcast-info-plcy>bundle>channel config>mcast-mgmt>mcast-info-plcy>bundle>channel>source-override
<b>Description</b>	<p>This command sets the relative preference level for multicast channels. The preference of a channel specifies its relative importance over other multicast channels. Eight levels of preference are supported; 0 through 7. Preference value 7 indicates the highest preference level.</p> <p>When the multicast ingress path manager is congested on one or more of the switch fabric multicast paths, it uses the preference values associated with each multicast record to determine which records will be allowed on the path and which records should be placed in a black-hole state.</p> <p>The preference value is also compared to the bundles cong-priority-threshold setting to determine the congestion priority of the channel. The result also dictates the channels multicast CAC class level (high or low). When the channels preference value is less than the congestion priority threshold, it is considered to have a congestion priority and CAC class value equal to low. When the channels preference value is equal to or greater than the threshold, it is considered to have a congestion priority and a CAC class value equal to high.</p> <p>The preference value is also compared to the bundles ecmp-opt-threshold setting to determine whether the channel is eligible for ECMP path dynamic optimization. If the preference value is equal to or less than the threshold, the channel may be optimized. If the preference value is greater than the threshold, the channel will not be dynamically optimized.</p>

The preference command may be executed in three contexts; bundle, channel and source-override. The bundle default preference value is 0. The channel and source-override preference settings are considered overrides to the bundle setting and have a default value of null (undefined).

The **no** form of the command restores the default preference value (0 or null depending on the context).

**Parameters** *preference-level* — The preference-level parameter is required and defines the preference value of the channel. It is represented by an integer value between 0 and 7.

**Values** 0 — 7

Bundle default:	0
Channel default:	Null (undefined)
Source-override default:	Null (undefined)

Override sequence — The channel setting overrides the bundle setting. The source-override setting overrides the channel and bundle settings.

## primary-tunnel-interface

**Syntax** **primary-tunnel-interface** **rsvp-p2mp** *lsp-name* **sender** *ip-address*  
**no primary-tunnel-interface**

**Context** config>mcast-mgmt>mcast-info-plcy>bundle  
 config>mcast-mgmt>mcast-info-plcy>bundle>channel  
 config>mcast-mgmt>mcast-info-plcy>bundle>channel>source-override

**Description** This command allows the user to define a bundle in the multicast-info-policy and specify channels in the bundle that must be received from the primary tunnel interface associated with an RSVP P2MP LSP. The multicast info policy is applied to the base router instance.

The egress LER will be able to accept multicast packets via two different methods. The regular RPF check on unlabelled IP multicast packets, which is based on routing table lookup. The static assignment which specifies the receiving of a multicast group <\*,G> or a specific <S,G> from a primary tunnel-interface associated with an RSVP P2MP LSP.

One or more primary tunnel interfaces in the base router instance can be configured. That is, the user will be able to specify to receive different multicast groups, <\*,G> or specific <S,G>, from different P2MP LSPs. This assumes that there are static joins configured for the same multicast groups at the ingress LER to forward over a tunnel interface associated with the same P2MP LSP.

At any given time, packets of the same multicast group can be accepted from either the primary tunnel interface associated with a P2MP LSP or from a PIM interface. These are mutually exclusive options. As soon as a multicast group is configured against a primary tunnel interface in the multicast info policy, it is blocked from other PIM interfaces.

A multicast packet received on a tunnel interface associated with a P2MP LSP can be forwarded over a PIM or IGMP interface which can be an IES interface, a spoke SDP terminated IES interface, or a network interface.

The **no** form of the command removes the static RPF check.

**Default** none

**Parameters** **rsvp-p2mp** *lsp-name* — Species a string of up to 32 characters identifying the LSP name as configured at the ingress LER.

**sender** *ip-address* — Specifies a string of 15 characters representing the IP address of the ingress LER for the LSP.

## source-override

**Syntax** **source-override** *ip-address* [**create**]  
**no source-override** *ip-address*

**Context** config>mcast-mgmt>mcast-info-plcy>bundle>channel

**Description** This command defines a multicast channel parameter override context for a specific multicast sender within the channel range. The specified senders IP address must be of the same type (IPv4 or IPv6) as the containing channel range.

The **no** form of the command removes the specified sender override context from the channel range.

**Default** none

**Parameters** *ip-address* — Specifies either an IPv4 or IPv6 address and it must be the same type as the containing channel range.

**Values**

ipv4-address	a.b.c.d
ipv6-address	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

**create** — The create keyword is required if creating a new source override when the system is configured to require the explicit use of the keyword to prevent accidental object creation. Objects may be accidentally created when this protection is disabled and an object name is mistyped when attempting to edit the object. This keyword is not required when the protection is disabled. The keyword is ignored when the specified source override IP address already exists.

## cong-priority-threshold

**Syntax** **cong-priority-threshold** *preference-level*  
**no cong-priority-threshold**

**Context** config>mcast-mgmt>mcast-info-plcy>bundle

**Description** This command defines the preference level threshold where records change from low congestion priority to high congestion priority. Congestion priority is used by the ingress multicast path queues to map packets entering the queue to either the low priority MBS (LP-MBS) or the MBS tail-drop threshold. In the event that congestion happens on the queue, the queue depth increases. As the queue depth increases beyond the low priority MBS, packets with low priority congestion priority are discarded. This leaves room in the queue for packets with high congestion priority until the queue reaches the MBS threshold.

The default congestion priority threshold is 4. This means that multicast channels with a preference level of 0 to 3 will be treated as having low congestion priority and channels with preference level of 4 to 7 will be treated as having a high congestion priority. The **cong-priority-threshold** command

can be used to change the default threshold. Any multicast channel with a preference equal to or higher than the configured threshold will be treated with high congestion priority.

The **cong-priority-threshold** value is also used by the multicast CAC manager to derive the class of a channel matched by the multicast information policy. Channels with a preference less than the configured threshold will be treated as *low* class and channels with a preference equal to or greater than the threshold will be treated as *high* class.

Changing the **cong-priority-threshold** value causes all channels congestion priority to be reevaluated. Both the ingress multicast path managers and multicast CAC managers must be updated.

The **no** form of the command restores the default congestion priority preference threshold value.

<b>Default</b>	4
<b>Parameters</b>	<i>preference-level</i> — The preference-level parameter is required when specifying the cong-priority-threshold.
<b>Values</b>	0 — 7
<b>Default</b>	4

## ecmp-opt-threshold

<b>Syntax</b>	<b>ecmp-opt-threshold</b> <i>preference-level</i> <b>no ecmp-opt-threshold</b>
<b>Context</b>	config>mcast-mgmt>mcast-info-plcy>bundle
<b>Description</b>	<p>This command defines the preference level threshold where multicast ECMP path management is allowed to dynamically optimize channels based on topology or bandwidth events. If the channels preference is equal to or less than the ecmp-opt-threshold, ECMP is allowed to move the channel between ECMP paths when bandwidth distribution events happen. Channels with a preference level higher than the threshold will not be moved during these events.</p> <p>The default ECMP optimization limit threshold is 7. This means that multicast channels with a preference level of 0 to 7 (all channels) will be allowed to move between ECMP paths. The ecmp-opt-threshold command can be used to change the default threshold.</p> <p>Changing the threshold causes all channels ECMP optimization eligibility to be reevaluated.</p> <p>The <b>no</b> form of the command restores the default ECMP optimization preference threshold value.</p>
<b>Parameters</b>	<i>preference-level</i> — The preference-level parameter is required when specifying the ecmp-opt-threshold. An integer value from 0 to 7 must be specified.
<b>Values</b>	0 — 7
<b>Default</b>	7

## mc-ecmp-balance

<b>Syntax</b>	<b>[no] mc-ecmp-balance</b>
<b>Context</b>	config>router>pim

<b>Description</b>	<p>This command enables multicast balancing of traffic over ECMP links considering multicast bandwidth. When enabled, every multicast stream that needs to be forwarded over an ECMP link will be evaluated for the sum total multicast utilization currently using the ECMP interface in question.</p> <p>Note that a given interface can be shared between multiple (partially overlapping) ECMP sets. This is taken into consideration and a complete balance is attempted.</p> <p>ECMP load balancing helps to avoid loss of unicast traffic over ECMP links as it will load balance over ECMP links and if multicast is not balanced then it is possible that a given link does not have sufficient bandwidth to pass its allotted unicast traffic portion.</p> <p>In order to achieve a proper balance, multicast groups and their bandwidth should be configured in the config <b>mcast-management</b> context.</p> <p>If the bandwidth is not configured, then the default value applies, and for the purpose of ECMP load balancing, the net effect will be that the balance achieved reflects a balance of the number of multicast groups traversing over the ECMP link. The bandwidth used in this policy is the configured value, not the actual bandwidth.</p> <p>If <b>mc-ecmp-balance</b> is enabled, a redistribution may be triggered whenever a interface is added to an ECMP link.</p> <p>If <b>mc-ecmp-balance</b> is enabled, a period re-balance may be configured that re-optimizes the distribution as some multicast streams may have been removed from the ECMP link.</p> <p>If <b>mc-ecmp-balance</b> is enabled, then a threshold (<b>ecmp-opt-threshold</b>) can be configured to avoid moving multicast streams where interruption should be avoided.</p> <p>The <b>ecmp-opt-threshold</b> is used to define the preference level threshold where multicast ECMP path management is allowed to dynamically optimize channels based on topology or bandwidth events. If the channels preference is equal to or less than the <b>ecmp-opt-threshold</b>, ECMP is allowed to move the channel between ECMP paths when bandwidth distribution events happen. Channels with a preference level higher than the threshold will not be moved during these events.</p> <p>Changing the <b>ecmp-opt-threshold</b> causes all channels ECMP optimization eligibility to be reevaluated.</p> <p>The <b>no</b> form of the command removes the re-balancing capability from the configuration.</p>
--------------------	--

## mc-ecmp-balance-hold

<b>Syntax</b>	<b>mc-ecmp-balance-hold</b> <i>minutes</i> <b>no mc-ecmp-balance-hold</b>
<b>Context</b>	config>router>pim
<b>Description</b>	<p>This command defines a hold period that applies after an interface has been added to the ECMP link. It is also used periodically to rebalance if channels have been removed from the ECMP link.</p> <p>If the ECMP interface has not changed in the hold period and if no multicast streams have been removed, then no action will be taken when the timer triggers.</p> <p>This parameter should be used to avoid excessive changes to the multicast distribution.</p> <p>A rebalance will not occur to multicast streams that have a priority greater than the configured <b>ecmp-opt-threshold</b>.</p> <p>The <b>no</b> form of the command reverts to default.</p>
<b>Parameters</b>	<i>minutes</i> — Specifies the hold down time in minutes.

**Values** 2 — 600

## mc-ecmp-rebalance

<b>Syntax</b>	<b>mc-ecmp-rebalance [ecmp-opt-threshold]</b>
<b>Context</b>	tools>perform>router>pim
<b>Description</b>	This command triggers an immediate rebalance, regardless if the hold timer has triggered or if any changes have occurred.
<b>Parameters</b>	<b>ecmp-opt-treshold</b> — Forces a rebalance of all multicast streams with a priority equal or less then the specified value.  Specifying the value of 7 will force all multicast streams to be re-balanced regardless of the configured <b>ecmp-opt-threshold</b> value.  <b>Values</b> 1 — 7



---

## Chassis Level Commands

### chassis-level

<b>Syntax</b>	<b>chassis-level</b>
<b>Context</b>	config>mcast-management
<b>Description</b>	<p>The chassis-level CLI node contains the multicast plane replication limit for each switch fabric multicast plane.</p> <p>The chassis-level node always exists and contains the configuration command to define the total replication rates for primary and secondary associated ingress paths for each switch fabric multicast plane.</p>

### mmlrp-impmm-override

<b>Syntax</b>	<b>[no] mmlrp-impmm-override</b>
<b>Context</b>	config>mcast-mgmt>chassis-level
<b>Description</b>	<p>This command enables enables ingress Multicast Path Management (IMPM) from monitoring PIM and IGMP.</p> <p>The <b>no</b> form of the command disables the IMPM monitoring.</p>
<b>Default</b>	no mmlrp-impmm-override

### per-mcast-plane-capacity

<b>Syntax</b>	<b>[no] per-mcast-plane-capacity</b>
<b>Context</b>	config>mcast-mgmt>chassis-level
<b>Description</b>	<p>This CLI node contains the configuration of the overall multicast (primary plus secondary) and specific secondary rate limits for each switch fabric multicast plane.</p>

### mcast-capacity

<b>Syntax</b>	<b>mcast-capacity</b> <i>primary-percentage</i> <b>secondary</b> <i>secondary-percentage</i> <b>no mcast-capacity</b>
<b>Context</b>	config>mcast-mgmt>chassis-level>plane-capacity
<b>Description</b>	<p>This command configures the primary and secondary multicast plane capacities used when the full complement of possible switch fabrics in the system is not up (at least one possible switch fabric is not provisioned or is down). The rates are defined as a percentage of the total multicast plane capacity which is configured using the total-capacity command.</p>

The **no** form of the command reverts to the default values.

<b>Default</b>	primary-percentage	100.00
	secondary-percentage	90.00
with SFM5	primary-percentage	87.5
	secondary-percentage	87.5

**Parameters**

*primary-percentage* — Specifies the percentage of the total multicast plane capacity to be used for primary multicast planes.

**secondary** *secondary-percentage* — Specifies the percentage of the total multicast plane capacity to be used for secondary multicast planes.

**Values** 0.01 — 100

## redundant-mcast-capacity

<b>Syntax</b>	<b>redundant-mcast-capacity</b> <i>primary-percentage</i> <b>secondary</b> <i>secondary-percentage</i> <b>no redundant-mcast-capacity</b>
<b>Context</b>	config>mcast-mgmt>chassis-level>plane-capacity
<b>Description</b>	This command configures the primary and secondary multicast plane capacities used when the full complement of possible switch fabrics in the system are up. The rates are defined as a percentage of the total multicast plane capacity which is configured using the total-capacity command.  The <b>no</b> form of the command reverts to the default values.
<b>Default</b>	primary-percentage 100.00
	secondary-percentage 90.00
with SFM5	primary-percentage 87.5
	secondary-percentage 87.5
<b>Parameters</b>	<i>primary-percentage</i> — Specifies the percentage of the total multicast plane capacity to be used for primary multicast planes.
	<b>Values</b> 0.01 — 100
	<b>secondary</b> <i>secondary-percentage</i> — Specifies the percentage of the total multicast plane capacity to be used for secondary multicast planes.
	<b>Values</b> 0.01 — 100

## total-capacity

<b>Syntax</b>	<b>total-capacity</b> <i>capacity</i> <b>no total-capacity</b>
<b>Context</b>	config>mcast-mgmt>chassis-level>plane-capacity
<b>Description</b>	This command configures the total multicast plane capacity supported individually by all switch fabric multicast planes.  The <b>no</b> form of the command reverts to the default.
<b>Default</b>	SR/ESS: 2000 SR12e/XRS: dynamic
<b>Parameters</b>	<i>capacity</i> — Specifies the multicast plane capacity in Mbps.  <b>Values</b> <b>2000, 4000, 5250, 8250, dynamic</b> (Specifies that multicast plane capacity will be determined based on provisioned line cards and switch fabrics in the chassis.)

## round-robin-inactive-records

<b>Syntax</b>	<b>[no] round-robin-inactive-records</b>
<b>Context</b>	config>mcast-mgmt>chassis-level
<b>Description</b>	This command specifies whether initially inactive multicast records use the IOM default secondary multicast path or not. When enabled, the system redistributes newly populated inactive records among all available IOM multicast paths and multicast switch fabric planes. When disabled, the system continues to set all initially inactive multicast records to use the IOM default secondary multicast path.
<b>Default</b>	no round-robin-inactive-records

---

## Multicast Redirection Commands

### action

<b>Syntax</b>	<b>action</b> { <b>accept</b>   <b>next-entry</b>   <b>next-policy</b>   <b>reject</b> } <b>no action</b>
<b>Context</b>	config>router>policy-options>policy-statement>entry
<b>Description</b>	<p>This command creates the context to configure actions to take for routes matching a route policy statement entry.</p> <p>This command is required and must be entered for the entry to be active.</p> <p>Any route policy entry without the <b>action</b> command will be considered incomplete and will be inactive.</p> <p>The <b>no</b> form of the command deletes the action context from the entry.</p>
<b>Default</b>	<b>no action</b> — No action is defined.
<b>Parameters</b>	<p><b>accept</b> — Specifies routes matching the entry match criteria will be accepted and propagated.</p> <p><b>next-entry</b> — Specifies that the actions specified would be made to the route attributes and then policy evaluation would continue with next policy entry (if any others are specified).</p> <p><b>next-policy</b> — Specifies that the actions specified would be made to the route attributes and then policy evaluation would continue with next route policy (if any others are specified).</p> <p><b>reject</b> — Specifies routes matching the entry match criteria would be rejected.</p>

### interface

<b>Syntax</b>	<b>interface</b> <i>port:tags</i>
<b>Context</b>	config>router>igmp config>services>ies config>services>vpn
<b>Description</b>	This command enables IGMP (or MLD) on the multicast redirect interface.

---

## Forwarding Plane Commands

### fp

<b>Syntax</b>	<b>fp</b> [ <i>fp-number</i> ]				
<b>Context</b>	config>card				
<b>Description</b>	<p>The fp CLI node contains the multicast path management configuration commands for IOM-3 ingress multicast management. Ingress multicast management manages multicast switch fabric paths which are forwarding plane specific. On IOM-1 and IOM-2, each MDA has a dedicated forwarding plane and so have dedicated multicast paths to the switch fabric allowing the multicast management to be defined per MDA. IOM-3 has a single forwarding plane shared by two MDAs making the previous model of managing multicast at the MDA level problematic. The fp node has been added to simplify ingress multicast management on IOM-3.</p> <p>In subsequent releases, the fp node will be moved to IOM-1 and IOM-2 and the multicast path management commands will be consistent for all IOM types. Other forwarding plane resource configuration commands (i.e. buffer pool management) are expected to also move to the fp node.</p> <p>While IOM-3 only has a single forwarding plane, other IOMs that will use the node in the future will have multiple (i.e. IOM-1 and IOM-2). To accommodate multiple forwarding planes, each forwarding plane is assigned a number. The default forwarding plane is 1. When entering the fp node, if the forwarding plane number is omitted, the system will assume forwarding plane number 1. All show and save configuration output must include the forwarding plane number.</p>				
<b>Parameters</b>	<p><i>fp-number</i> — The fp-number parameter is optional following the fp command. If omitted, the system assumes forwarding plane number 1. More than a single forwarding plane will be supported in the future when IOMs other than IOM-3 are supported.</p> <table><tr><td><b>Values</b></td><td>1</td></tr><tr><td><b>Default</b></td><td>1</td></tr></table>	<b>Values</b>	1	<b>Default</b>	1
<b>Values</b>	1				
<b>Default</b>	1				

### ingress

<b>Syntax</b>	<b>ingress</b>
<b>Context</b>	config>card>fp
<b>Description</b>	<p>The ingress CLI node within the fp node contains the multicast path management configuration commands for IOM-3 ingress multicast management. For this release, on the bandwidth-policy command is supported within the ingress node.</p>

### multicast-path-management

<b>Syntax</b>	<b>multicast-path-management</b>
<b>Context</b>	config>card>fp>ingress config>card>mda>ingress
<b>Description</b>	The multicast-path-management CLI node contains the forwarding plane or MDA settings for ingress multicast path management. Enter the node to configure the bandwidth-policy, the individual path bandwidth overrides and the administrative state of ingress multicast path management.

### bandwidth-policy

<b>Syntax</b>	<b>bandwidth-policy</b> <i>policy-name</i> <b>no bandwidth-policy</b>				
<b>Context</b>	config>card>fp>ingress>mcast-path-management config>card>mda>ingress>mcast-path-management				
<b>Description</b>	<p>This command is used to explicitly associate a bandwidth policy to a forwarding plane or MDA. The bandwidth policy defines the dynamic rate table and the multicast paths bandwidth and queuing parameters.</p> <p>If a bandwidth policy is not explicitly associated with a forwarding plane or MDA, the default bandwidth policy is used when ingress multicast path management is enabled.</p> <p>The <b>no</b> form of the command removes an explicit bandwidth policy from a forwarding plane or MDA and restores the default bandwidth policy.</p>				
<b>Parameters</b>	<p><i>policy-name</i> — The policy-name parameter is required and defines the bandwidth policy that should be associated with the MDA or forwarding plane for ingress multicast path management. If the policy name does not exist, the bandwidth-policy command will fail.</p> <table><tr><td><b>Values</b></td><td>Any existing bandwidth policy name</td></tr><tr><td><b>Default</b></td><td>default</td></tr></table>	<b>Values</b>	Any existing bandwidth policy name	<b>Default</b>	default
<b>Values</b>	Any existing bandwidth policy name				
<b>Default</b>	default				

### primary-override

<b>Syntax</b>	<b>primary-override</b>
<b>Context</b>	config>card>mda>ingress>mcast-mgmt
<b>Description</b>	<p>This command enables the context to configure MDA ingress multicast path-limit overrides.</p> <p>The path override CLI nodes are not supported on IOM-3.</p>

## secondary-override

<b>Syntax</b>	<b>secondary-override</b>
<b>Context</b>	config>card>mda>ingress>mcast-mgmt
<b>Description</b>	This command enables the context to configure MDA ingress multicast path-limit overrides.  The path override CLI nodes are not supported on IOM-3.

## ancillary-override

<b>Syntax</b>	<b>ancillary-override</b>
<b>Context</b>	config>card>mda>ingress>mcast-mgmt
<b>Description</b>	This command enables the context to configure MDA ingress multicast path-limit overrides.

## path-limit

Syntax	path-limit <i>megabits-per-second</i> no path-limit		
Context	config>card>mda>ingress>mcast-mgmt>primary-override config>card>mda>ingress>mcast-mgmt>secondary-override config>card>mda>ingress>mcast-mgmt>ancillary-override		
Description	<p>The path-limit command is used to override the path limits contained in the bandwidth policy associated with the MDA. The path limits are used to give the upper limit that multicast channels may use on each path.</p> <p>The path-limit commands are not supported on IOM-3.</p> <p>The no form of the command removes a path limit override from an ingress multicast path and restore the path limit defined in the bandwidth policy associated with the MDA.</p>		
Parameters	<i>megabits-per-second</i> — The megabits-per-second parameter is required when executing the path-limit command and is expressed as an integer representing multiples of 1,000,000 bits per second.		
Values	Primary-override:	1 to 2000	
	Secondary-override:	1 to 2000	
	Ancillary-override:	1 to 5000	
Default	None		





# Configuring Enhanced Subscriber Management with CLI

This section provides information to configure subscriber management features using the command line interface. It is assumed that the reader is familiar with VPLS services.

Topics in this section include:

- [Configuring RADIUS Authentication of DHCP Sessions on page 718](#)
- [Configuring Enhanced Subscriber Management on page 719](#)
  - [Basic Configurations on page 719](#)
  - [Configuring Enhanced Subscriber Management Entities on page 720](#)
    - [Configuring a Subscriber Identification Policy on page 721](#)
    - [Configuring a Subscriber Profile on page 722](#)
    - [Configuring an SLA Profile on page 724](#)
    - [Configuring Explicit Mapping Entries on page 725](#)
  - [Applying the Profiles and Policies on page 726](#)
- [Python Script Support for ESM on page 1465](#)
- [Sample Python Scripts on page 1474](#)

## Configuring RADIUS Authentication of DHCP Sessions

When RADIUS authentication for subscriber sessions is enabled, DHCP messages from subscribers are temporarily held by the BSA, while the user's credentials are checked on a RADIUS server.

Configuring RADIUS authentication for subscriber sessions is done in two steps:

- First define an authentication-policy in the **config>subscriber-mgmt>authentication-policy** context.
- Then apply the policy to one or more SAPs in the **config>service>vpls>sap>authentication-policy *auth-plcy-name*** context (for a VPLS service).

Or apply the policy to one or more interfaces **config>service>ies>if>authentication-policy *auth-plcy-name*** context (for an IES service):

The following example displays a partial BSA configuration with RADIUS authentication:

```
A:ALA-1>config>service# info
-----
subscriber-management
  authentication-policy BSA_RADIUS create
    description "RADIUS policy for DHCP users Authentication"
    password "mysecretpassword"
    radius-authentication-server
      server 1 address 10.100.1.1 secret "radiuskey"
      retry 3
      timeout 10
    exit
    re-authentication
    user-name-format circuit-id
  exit
exit
...
vpls 800 customer 6001
  description "VPLS with RADIUS authentication"
  sap 2/1/4:100 split-horizon-group DSL-group create
    authentication-policy BSA_RADIUS
  exit
  sap 3/1/4:200 split-horizon-group DSL-group create
    authentication-policy BSA_RADIUS
  exit
  no shutdown
exit
...
-----
A:ALA-1>config>service#
```

## Configuring Enhanced Subscriber Management

---

### Basic Configurations

Configuring and applying the Enhanced Subscriber Management profiles and policies are optional. There are no default Profiles or policies.

The basic Enhanced Subscriber Management profiles and policies must conform to the following:

- Unique profile or policy names (IDs)
  - Profiles and/or policies must be associated with a VPLS or IES service to facilitate Enhanced Subscriber Management.
  - QoS and IP filter entries configured in Enhanced Subscriber Management profiles and policies override the defaults and/or modified parameters or the default policies.
  - The Enhanced Subscriber Management profiles and policies must be configured within the context of VPLS or IES.
- 

### Subscriber Interface Configuration

The following output displays a basic subscriber interface configuration.

```
*A:ALA-48>config>service>ies>sub-if# info
-----
description "Routed CO - Antwerp 2018"
address 192.168.2.254/24
address 192.168.3.254/24
address 192.168.4.254/24
address 192.168.5.254/24
address 192.168.6.254/24
group-interface "DSLAM_01" create
  description "Routed CO - vlan / subscriber"
  sap 1/1/2:1001 create
    static-host ip 192.168.2.2 create
    exit
  sap 1/1/2:1002 create
    static-host ip 192.168.2.2 create
    exit
  sap 1/1/2:1004 create
    static-host ip 192.168.2.4 create
    exit
  sap 1/1/2:1100 create
    static-host ip 192.168.2.100 create
    exit
  exit
exit
-----
*A:ALA-48>config>service>ies>sub-if#
```

## **Configuring Enhanced Subscriber Management Entities**

- [Configuring a Subscriber Identification Policy on page 721](#)
- [Configuring a Subscriber Profile on page 722](#)
- [Configuring a Subscriber Identification Policy on page 721](#)
- [Configuring Explicit Mapping Entries on page 725](#)
- [Applying the Profiles and Policies on page 726](#)

## Configuring a Subscriber Identification Policy

The following displays an example of a subscriber identification policy configuration:

```
A:ALA-48>config>subscr-mgmt# info
-----
...
    sub-ident-policy "Globocom" create
        description "Subscriber Identification Policy Id Globocom"
        sub-profile-map
            entry key "1/1/2" sub-profile "ADSL Business"
        exit
        sla-profile-map
            entry key "1/1/2" sla-profile "BE-Video"
        exit
        primary
            script-url "primaryscript.py"
            no shutdown
        exit
        secondary
            script-url "secondaryscript.py"
        exit
        tertiary
            script-url "tertiaryscript.py"
            no shutdown
        exit
    exit
...
-----
A:ALA-48>config>subscr-mgmt#
```

### Configuring a Subscriber Profile

Enhanced Subscriber Management subscriber profile configurations specify existing QoS scheduler profiles. In the following example, “BE-Video-max100M” is specified in the sub-profile “ADSL Business” for the ingress-scheduler-policy. “Upload” is specified in the sub-profile egress-scheduler-policy.

```
#-----
echo "QoS Policy Configuration"
#-----
    qos
        scheduler-policy "BE-Video-max100M" create
            description "Scheduler Policy Id BE-Video-max100M"
            tier 1
                scheduler "tier1" create
                    description "Scheduler Policy Id BE-Video-max100M Tier 1 tier1"
                exit
            exit
        exit
        scheduler-policy "Upload" create
            description "Scheduler Policy Id Upload"
            tier 3
                scheduler "tier3" create
                    description "Scheduler Policy Id Upload Tier 3 tier3"
                exit
            exit
        exit
        sap-ingress 2 create
            description "Description for Sap-Ingress Policy id # 2"
            queue 1 create
                parent "tier1"
            exit
            queue 11 multipoint create
                parent "tier1"
            exit
        exit
        sap-egress 3 create
            description "Description for Sap-Egress Policy id # 3"
            queue 1 create
                parent "tier3"
            exit
        exit
    exit
#-----
```

The following displays an example of a subscriber identification policy configuration:

```
A:ALA-48>config>subscr-mgmt# info
-----
...
    sub-profile "ADSL Business" create
        description "Subscriber Profile Id ADSL Business"
        ingress-scheduler-policy "BE-Video-max100M"
            scheduler "tier1" rate 99
        exit
        egress-scheduler-policy "Upload"
            scheduler "tier3" rate 1 cir 1
        exit
        sla-profile-map
            entry key "1/1/3" sla-profile "BE-Video"
        exit
    exit
-----
A:ALA-48>config>subscr-mgmt#
```

### Configuring an SLA Profile

The following displays an example of a SLA Profile configuration:

```
A:ALA-48>config>subscr-mgmt# info
-----
subscriber-mgmt
  sla-profile "BE-Video" create
    description "SLA Profile Id BE-Video"
    ingress
      qos 2
      queue 1
      exit
    exit
  exit
  egress
    qos 3
    queue 1
    exit
  exit
  exit
  exit
-----
A:ALA-48>config>subscr-mgmt#
```



## Configuring Explicit Mapping Entries

The following displays an example of a explicit subscriber mapping:

```
A:ALA-7>config>subscr-mgmt# info
-----
A:ALA-48>config>subscr-mgmt# info
-----
...
    explicit-subscriber-map
        entry key "1/1/1:1111" sub-profile "ADSL GO" alias "Sub-Ident-1/1/1:
1111" sla-profile "BE-Video"
        exit
...
-----
A:A:ALA-48>config>subscr-mgmt#
```

## Applying the Profiles and Policies

Apply the Enhanced Subscriber Management profiles and policies to the following entities:

- [SLA Profile on page 726](#)
- [Subscriber Identification Policy on page 727](#)
- [Subscriber Profile on page 727](#)

---

### SLA Profile

**CLI Syntax:** `configure>service>ies service-id  
interface ip-int-name  
sap sap-id  
host {[ip ip-address] [mac ieee-address] [subscriber  
sub-ident-string] [sub-profile sub-profile-name]  
[sla-profile sla-profile-name]}`

**CLI Syntax:** `configure>service>vpls service-id  
sap sap-id  
host {[ip ip-address] [mac ieee-address]} [subscriber sub-  
ident-string] [sub-profile sub-profile-name] [sla-pro-  
file sla-profile-name]  
sub-sla-mgmt  
def-sla-profile default-sla-profile-name  
single-sub-parameters  
non-sub-traffic sub-profile sub-profile-name sla-  
profile sla-profile-name [subscriber sub-ident-  
string]`

**CLI Syntax:** `configure>subscriber-mgmt  
explicit-subscriber-map  
entry key sub-ident-string [sub-profile sub-profile-name]  
[alias sub-alias-string] [sla-profile sla-profile-  
name]  
sub-ident-policy sub-ident-policy-name  
sla-profile-map  
entry key sla-profile-string sla-profile sla-profile-  
name  
sub-profile sla-profile-map  
sla-profile-map  
entry key sla-profile-string sla-profile sla-profile`

## Subscriber Identification Policy

**CLI Syntax:** `configure>service>ies service-id`  
`interface ip-int-name`  
`sap sap-id`  
`host {[ip ip-address] [mac ieee-address] [subscriber`  
`sub-ident-string] [sub-profile sub-profile-name]`  
`[sla-profile sla-profile-name]`

---

## Subscriber Profile

**CLI Syntax:** `configure>service>ies service-id`  
`interface ip-int-name`  
`sap sap-id`  
`host {[ip ip-address] [mac ieee-address] [subscriber`  
`sub-ident-string] [sub-profile sub-profile-name]`  
`[sla-profile sla-profile-name]`

**CLI Syntax:** `configure>service>vpls service-id`  
`sap sap-id`  
`host {[ip ip-address] [mac ieee-address]} [subscriber sub-`  
`ident-string] [sub-profile sub-profile-name] [sla-pro-`  
`file sla-profile-name]`  
`sub-sla-mgmt`  
`def-sub-profile default-sub-profile-name`  
`single-sub-parameters`  
`non-sub-traffic sub-profile sub-profile-name sla-`  
`profile sla-profile-name [subscriber sub-ident-`  
`string]`

**CLI Syntax:** configure>subscriber-mgmt  
sub-profile *subscriber-profile-name*  
explicit-subscriber-map  
    entry key *sub-ident-string* [**sub-profile** *sub-profile-name*]  
        [alias *sub-alias-string*] [sla-profile *sla-profile-name*]  
sub-ident-policy *sub-ident-policy-name*  
sub-profile-map  
    entry key *sub-profile-string* **sub-profile** *sub-profile-name*

# Triple Play Enhanced Subscriber Management

---

## In This Section

This section describes features which provide Enhanced Subscriber Management functions for Triple Play services.

Topics in this section include:

- [Uniform RADIUS Server Configuration on page 730](#)
- [RADIUS Authentication of Subscriber Sessions on page 735](#)
- [radius-server-policy retry Attempt Overview on page 744](#)
- [Enhanced Subscriber Management Overview on page 773](#)
- [L2TP Tunnel RADIUS Accounting on page 1007](#)
- [RADIUS Route Download on page 1014](#)
- [Managed SAP \(MSAP\) on page 1016](#)
- [Volume and Time Based Accounting on page 1022](#)
- [Subscriber Host Idle Timeout on page 1028](#)
- [Web Authentication Protocol \(WPP\) on page 1030](#)
- [One-time HTTP Redirection Overview on page 1035](#)
- [ESM over MPLS Pseudowires on page 1036](#)
- [On-Demand Subnet Allocation \(ODSA\) on page 1054](#)
- [Logical Link Identifier \(LLID\) on page 1058](#)
- [Open Authentication Model for DHCP and PPPoE Hosts on page 1059](#)
- [Flexible Subscriber-Interface Addressing \(Unnumbered Subscriber-Interfaces\) on page 1064](#)
- [uRPF for Subscriber Management on page 1073](#)
- [IPoE Sessions on page 1074](#)

# Uniform RADIUS Server Configuration

---

## RADIUS Server Configuration

The following two configuration methods co-exist but are mutually exclusive:

- [Uniform RADIUS Server Configuration \(Preferred\) on page 730](#)
  - [Legacy RADIUS Server Configuration on page 734](#)
- 

## Uniform RADIUS Server Configuration (Preferred)

This configuration method is preferred as it can be re-used amongst multiple applications (Subscriber authentication and accounting, L2TP tunnel accounting, WLAN gateway RADIUS proxy,) and enables additional functionality not available in the legacy configuration method. For example:

- A RADIUS server policy operational state can be controlled by reception of accounting on/off responses.
- Buffering of accounting messages: When all servers in a **radius-server-policy** are unreachable, it is possible to buffer the acct-stop and acct-interim-update messages for up to 25 hrs. When a RADIUS server becomes reachable again then the messages in the buffer are retransmitted.
- A configurable hold down time for accounting servers that are marked down and during which no new communication attempts will be made (hold-down-time).
- A configurable maximum number of outstanding RADIUS requests for accounting servers (pending-requests-limit).
- Increased retry and timeout values for unsuccessful RADIUS communication.
- Enhanced RADIUS server statistics
- IPv6 RADIUS server

**Note:** A RADIUS server is marked down if it detects a number of consecutive timeouts independent of transaction-id or origin of request.

Where a number of consecutive timeouts is defined by the number of retries configured below the radius-server-policy servers.

The default number of retries is 3, meaning 1 initial try and 2 retries.

If, for example, the RADIUS server has “2 timeouts, 1 reply, 1 timeouts”, whereby the timeouts are originated for the same host, the server is not marked down since intermediate replies were received.

To attach a RADIUS server policy to an authentication policy:

For example,

```
configure
  subscriber-mgmt
    authentication-policy "auth-policy-1" create
      radius-server-policy "aaa-server-policy-1"
    exit
  exit
```

**Notes:**

- To avoid conflicts, the following CLI commands are ignored in the authentication policy when a **radius-server-policy** is attached:
  - All commands in the **radius-authentication-server** context
  - **accept-authorization-change**
  - **coa-script-policy**
  - **accept-script-policy**
  - **request-script-policy**
- The **fallback-action** command specifies the action when no RADIUS server is available is configured direct in the **config>subscr-mgmt>auth-plcy** CLI context.

To attach a RADIUS server policy to a RADIUS accounting policy:

For example:

```
configure
  subscriber-mgmt
    radius-accounting-policy "acct-policy-1" create
      radius-server-policy "aaa-server-policy-1"
    exit
  exit
```

**Note:** To avoid conflicts, the following CLI commands are ignored in the RADIUS accounting policy when a **radius-server-policy** is attached:

- All commands in the **radius-accounting-server** context
- **acct-request-script-policy**

## RADIUS Server Configuration

To configure the RADIUS servers in a RADIUS server policy:

For example:

```
configure
aaa
    radius-server-policy "aaa-server-policy-1" create
        description "Radius AAA server policy"
        accept-script-policy "script-policy-2"
        acct-on-off oper-state-change
        acct-request-script-policy "script-policy-3"
        auth-request-script-policy "script-policy-1"
        no python-policy
        servers
            access-algorithm direct
            hold-down-time sec 30
            no ipv6-source-address
            retry 3
            router "Base"
            no source-address
            timeout sec 5
            buffering
                acct-interim min 60 max 3600 lifetime 5
                acct-stop min 60 max 3600 lifetime 5
        exit
        server 1 name "server-1"
        server 2 name "server-2"
    exit
exit
exit
```

To configure the RADIUS servers in the routing instance:

- In the Base routing instance: **configure>router>radius-server.**
- In a VPRN routing instance: **configure>service>vprn 10>radius-server.**
- In the management routing instance (out of band): **configure>router management>radius-server.**

For example:

```
configure
router
    radius-server
        server "server-1" address 172.16.1.1 secret <shared secret> hash2 create
            accept-coa
            coa-script-policy "script-policy-4"
            description "Radius server 1"
            pending-requests-limit 4096
            acct-port 1813
            auth-port 1812
        exit
        server "server-2" address 172.16.1.2 secret <shared secret> hash2 create
            accept-coa
            coa-script-policy "script-policy-4"
            description "Radius server 2"
```



```
        pending-requests-limit 4096
        acct-port 1813
        auth-port 1812
    exit
exit
exit
```

**Note:** To configure RADIUS CoA servers for use in Enhanced Subscriber Management, the server must be configured in the corresponding routing instance with the **accept-coa** command enabled.

### Legacy RADIUS Server Configuration

**Note:** It is recommended to migrate to the uniform RADIUS server configuration as described above to have additional functionality enabled.

To configure a RADIUS server in an authentication policy:

```
configure
subscriber-mgmt
  authentication-policy "auth-policy-1" create
    radius-authentication-server
      access-algorithm direct
      hold-down-time 30
      retry 3
      no source-address
      timeout 5
      router "Base"
      server 1 address 172.16.1.1 secret <shared secret> hash2 port 1812 pending-
requests-limit 4096
      server 2 address 172.16.1.2 secret <shared secret> hash2 port 1812 pending-
requests-limit 4096
    exit
    accept-authorization-change
    accept-script-policy "script-policy-2"
    coa-script-policy "script-policy-4"
    request-script-policy "script-policy-1"
  exit
exit
```

**Note:** In legacy RADIUS server configuration, to configure RADIUS CoA servers for use in Enhanced Subscriber Management, the server must be configured in the authentication policy with the **accept-authorization-change** command enabled. A CoA only server can be configured with the optional coa-only flag.

To configure a RADIUS server in a RADIUS accounting policy:

```
configure
subscriber-mgmt
  radius-accounting-policy "acct-policy-1" create
    radius-accounting-server
      access-algorithm direct
      retry 3
      timeout 5
      no source-address
      router "Base"
      server 1 address 172.16.1.1 secret <shared secret> hash2 port 1813
      server 2 address 172.16.1.2 secret <shared secret> hash2 port 1813
    exit
    acct-request-script-policy "script-policy-3"
  exit
exit
```

## RADIUS Authentication of Subscriber Sessions

This section describes the Alcatel-Lucent router acting as a Broadband Subscriber Aggregator (BSA).

Note that in the 7750 and 7710 TPSDA solutions, the Alcatel-Lucent 5750 Subscriber Services Controller (SSC) serves as the policy manager, DHCP and RADIUS server.

In this application, one of the required functions can be to authenticate users trying to gain access to the network. While sometimes the DHCP server (an SSC) can perform authentication, in most cases a RADIUS server (an SSC) is used to check the customer's credentials.

Note: Refer to section [DHCP Principles on page 314](#) for an explanation of DHCP and [DHCP Snooping on page 322](#) for an explanation of DHCP snooping.

For information about the RADIUS server selection algorithm, refer to the Security chapter in the OS System Management Guide.

If authentication is enabled, the router will temporarily hold any received DHCP discover message and will send a access-request message to a configured RADIUS server containing the client's MAC address and/or Circuit-ID (from the Option 82 field) as the user name. If and when access is granted by the RADIUS server, the router will then forward or relay the DHCP discover message to the DHCP server and thus allow an IP address to be assigned. If the RADIUS authentication request is denied, the DHCP message is dropped and an event is generated.

A typical initial DHCP scenario (after client bootup) is:

```

client          server
  ---discover---->
<----offer-----
  -----request---->
<-----ack-----

```

But, when the client already knows its IP address (when an existing lease is being renewed), it can skip straight to the request/ack phase:

```

client          server
  -----request---->
<-----ack-----

```

In the first scenario, the DHCP discover triggers an authentication message to RADIUS and the DHCP request also triggers RADIUS authentication. The previous reply is cached for 10 seconds, the second DHCP packet will not result in a RADIUS request.

In the second scenario, the DHCP request triggers an authentication message to RADIUS.

## RADIUS Authentication of Subscriber Sessions

If the optional subscriber management authentication policy **re-authentication** command is enabled, DHCP authentication is performed at every DHCP lease renew request. Only dynamic DHCP sessions are subject to remote authentication. Statically provisioned hosts are not authenticated.


## RADIUS Authentication Extensions

This section describes an extension to RADIUS functionality in the subscriber management context. As part of subscriber host authentication, RADIUS can respond with access-response message, which, in the case of an accept, can include several RADIUS attributes (standard and vendor-specific) that allow proper provisioning of a given subscriber-host.


Change-of-Authorization (CoA) messages as defined by RFC 3576, *Dynamic Authorization Extensions to Remote Authentication Dial In User Service (RADIUS)*, are supported. The goal of CoA messages is to provide a mechanism for “mid-session change” support through RADIUS.

## Triple Play Network with RADIUS Authentication

7450 ESS:



7750 SR



**Figure 42: Triple Play Aggregation Network with RADIUS-Based DHCP Host Authentication**

Figure 42 shows a flow of RADIUS authentication of DHCP hosts in the triple play aggregation environment. Besides granting the authentication of given DHCP host, the RADIUS server can include RADIUS attributes (standard and/or Vendor-Specific Attributes (VSAs)) which are then used by the network element to provision objects related to a given DHCP host.

RADIUS is a distributed client/server concept that is used to protect networks against unauthorized access. In the context of the router's subscriber management in TPSDA, the RADIUS client running on nodes sends authentication requests to the SSC.

RADIUS can be used to perform three distinct services:

- Authentication determines whether or not a given subscriber-host is allowed to access a specific service.
- Authorization associates connection attributes or characteristics with a specific subscriber host.
- Accounting tracks service use by individual subscribers.

The RADIUS protocol uses “attributes” to describe specific authentication, authorization and accounting elements in a user profile (which are stored on the RADIUS server). RADIUS messages contain RADIUS attributes to communicate information between network elements running a RADIUS client and a RADIUS server.

RADIUS divides attributes into two groups, standard attributes and Vendor-Specific Attributes (VSAs). VSA is a concept allowing conveying vendor specific configuration information in a RADIUS messages, discussed in RFC 2865, *Remote Authentication Dial In User Service (RADIUS)*. It is up to the vendor to specify the exact format of the VSAs. Alcatel-Lucent-specific VSAs are identified by vendor-id 6527.

## RADIUS Authorization Extensions

The following sections define different functional extensions and list relevant RADIUS attributes.

---

### Basic Provisioning of Authentication Extensions

In order to comply with RFC 4679, *DSL Forum Vendor-Specific RADIUS Attributes*, the software includes the following attributes in the authentication-request message:

- agent-circuit-id (as defined by DSL forum)
- agent-remote-id (as defined by DSL forum)

The following attributes can also be included if configured and provided by downstream equipment:

- actual-data-rate-upstream
- actual-data-rate-downstream
- minimum-data-rate-upstream
- minimum-data-rate-downstream
- access-loop-encapsulation

When the node is configured to insert (or replace) Option 82, the above mentioned attributes will have the content after this operation has been performed by the software.

In addition, the following standard RADIUS attributes will be included in authentication request messages (subject to configuration):

- NAS-identifier — string containing system-name
- NAS-port-id
- NAS-port-type — Values: 32 (null encap), 33 (dot1q), 34 (qinq), 15 (DHCP hosts), specified value (0 — 255)
- MAC-address (Alcatel-Lucent VSA 27)
- dhcp-vendor-class-id (Alcatel-Lucent VSA 36)
- calling-station-id

These will only be included in the access-request if they have been configured.

In order to provide the possibility to push new policies for currently active subscribers, the routers support commands to force re-authentication of the given subscriber-host. After issuing such a command, the router will send a DHCP force-renew packet, which causes the subscriber to renew its lease (provided it supports force-renew). The DHCP request and ACK are then authenticated and processed by the routers as they would be during a normal DHCP renew.



## Calling-Station-ID

A **calling-station-id** can be configured at SAP level and can be included in the RADIUS authentication and accounting messages. This attribute is used in legacy BRAS to identify the user (typically phone number used for RAS connection). In the broadband networks this was replaced by circuit-id in Option 82. However, the Option 82 format is highly dependent on access-node vendor, which makes interpretation in management servers (such as RADIUS) troublesome. Some operators use the calling-station-id attribute as an attribute indicating the way the circuit-id should be interpreted. The calling-station-id attribute can be configured as a string which will be configured on the SAP. It can also be configured to use the **sap-id**, **remote-id** or **mac-address**.

## Subscriber Session Timeout

To limit the lifetime of a PPP session or DHCPv4 host to a fixed time interval, a timeout can be specified from RADIUS. By default, a PPP session or DHCPv4 host has no session timeout (infinite).

For PPP sessions, a session-timeout can be configured in the ppp-policy. A RADIUS specified session-timeout overrides the CLI configured value.

```
subscriber-mgmt
  ppp-policy "ppp-policy-1" create
    session-timeout 86400
  exit
exit
```

When the session timeout expires a PPP session is terminated and a DHCPv4 host deleted. For a DHCPv4 host, a DHCP release message is also sent to the server .

The following two attributes can be used in RADIUS Access-Accept and CoA messages to limit the PPP session or DHCPv4 host session time ([Table 11](#)):

**Table 11: Subscriber Session Timeout**

Attribute ID	Attribute Name	Type	Limits	Purpose and Format
27	Session-Timeout	integer	2147483647 seconds	0 = infinite (no session-timeout) [1.. 2147483647] in seconds For example: Session-Timeout = 3600
26-6527-160	Alc-Relative-Ses-sion-Timeout	integer	[0..2147483647 ] seconds	0 = infinite (no session-timeout) [1..2147483647] in seconds For example: Alc-Relative-Session-Timeout = 3600

When specified in a RADIUS Access-Accept message, both attributes specify an absolute value for session timeout. When specified in a RADIUS CoA message, attribute [26-6527-160] Alc-Relative-Session-Timeout specifies a relative session timeout value in addition to the current session time while attribute [27] Session-Timeout specifies an absolute session timeout value. If the current session time is greater than the received Session-Timeout, a CoA NAK is sent with error cause “Invalid Attribute Value (407)”.

Only one of the above attributes to specify a session timeout can be present in a single RADIUS message. An event is raised when both are specified in a single message.

The output of the “show service id <service-id> ppp session detail” CLI command contains following fields related to session timeout for PPP sessions:

- Up Time: the PPP session uptime
- Session Time Left: the remaining time before the session is terminated
- RADIUS Session-TO: the RADIUS received session timeout value.

The output of the “show service id <service-id> dhcp lease-state detail” CLI command contains following fields related to session timeout for DHCPv4 hosts:

- Up Time: the DHCPv4 host uptime
- Remaining Lease Time: the remaining time before the lease expires in the DHCP server. The client should renew its lease before this time.
- Remaining SessionTime: the remaining time before the DHCPv4 host is deleted
- Session-Timeout: the DHCPv4 host is deleted when its uptime reaches the Session-Timeout value.
- Lease-Time: the lease time specified by the DHCPv4 server

**Note:**

In a radius-proxy scenario or when a DHCPv4 host is created with a RADIUS CoA message, the RADIUS attribute [26-6527-174] Alc-Lease-Time attribute must be used to specify the lease time. If the [26-6527-174] Alc-Lease-Time is not present in these scenarios, then the RADIUS attribute [27] Session-Timeout is interpreted as DHCPv4 lease time.

## Domain Name in Authentication

In many networks, the user name has specific meaning with respect to the domain (ISP) where the user should be authenticated. In order to identify the user correctly, the user name in an authentication-request message should contain a domain-name. The domain-name can be derived from different places. In PPPoE authentication the domain name is given by the PPPoE client with the user name used in PAP or CHAP authentication. For DHCP hosts similar functionality is implemented by a “pre-authentication” lookup in a local user database before performing the RADIUS request.

For example, it can be derived from option60 which contains the vendor-specific string identifying the ISP the set-box has been commissioned by.

To append a domain name to a DHCP host, the following configuration steps should be taken:

- Under the (group or IP) interface of the service, a local user database should be configured in the DHCP node and no authentication policy should be configured.
- In the local user database, there should be a host entry containing both the domain name to be appended and an authentication policy that should be used for RADIUS authentication of the host. The host entry should contain no other information needed for setting up the host (IP address, ESM string), otherwise the DHCP request will be dropped.
- In the authentication policy, the **user-name-format** command should contain the parameter **append** *domain-name*.

---

## RADIUS Reply Message for PPPoE PAP/CHAP

The string returned in a [18] Reply-Message attribute in a RADIUS Access-Accept is passed to the PPPoE client in the CHAP Success or PAP Authentication-Ack message.

The string returned in a [18] Reply-Message attribute in a RADIUS Access-Reject is passed to the PPPoE client in the CHAP Failure or PAP Authentication-Nak message.

When no [18] Reply-Message attribute is available, the SROS default messages are used instead: “CHAP authentication success” or ”CHAP authentication failure” for CHAP and “Login ok” or ”Login incorrect” for PAP.

## **radius-server-policy retry Attempt Overview**

This feature maximizes the use of the remaining healthy RADIUS servers for subscriber authentication and accounting. After the hold down time expires, a single radius message is used to determine the status of the RADIUS server. If the server remains unresponsive after waiting for a single timeout interval (without any retries) then it is placed back into the hold down state. If the RADIUS server responds then it is used for subscriber authentication and accounting with the rest of the healthy servers.

## Provisioning of Enhanced Subscriber Management (ESM) Objects

In the ESM concept on network elements, a subscriber host is described by the following aspects:

- subscriber-id-string
- subscriber-profile-string
- sla-profile-string
- ancp-string
- intermediate-destination-identifier-string
- application-profile-string

This information is typically extracted from DHCP-ACK message using a Python script, and is used to provision subscriber-specific resources such as queues and filter entries. As an alternative to extracting this information from DHCP-ACK packet, provisioning from RADIUS server is supported.

As a part of this feature, the following VSAs have been defined:

- alc-subscriber-id-string — Contains a string which is interpreted as a subscriber-id.
- alc-subscriber-profile-string — Contains a string which is interpreted as a subscriber profile
- alc-sla-profile-string — Contains string which is interpreted as an SLA profile.
- alc-ancp-string — Contains string which is interpreted as an ANCP string.
- alc-int-dest-id-string — Contains a string which is interpreted as an intermediate destination ID
- alc-app-profile-string — Contains a string which is interpreted as an application profile

Note that these strings can be changed in a CoA request.

When RADIUS authentication response messages contain the above VSAs, the information is used during processing of DHCP-ACK message as an input for the configuration of subscriber-host parameters, such as QoS and filter entries.

If ESM is not enabled on a given SAP, information in the VSAs is ignored.

If ESM is enabled and the RADIUS response does not include all ESM-related VSAs (an ANCP string is not considered as a part of ESM attributes), only the subscriber-id is mandatory (the other ESM-related VSAs are not included). The remaining ESM information (sub-profile, sla-profile) will be extracted from DHCP-ACK message according to “normal” flow (Python script, etc.).

If the profiles are missing from RADIUS, they are not extracted from the DHCP data with Python to prevent inconsistent information. Instead, the data will revert to the configured default values.

However, if the above case, a missing subscriber ID will cause the DHCP request to be dropped. The DHCP server will not be queried in that case.

When no DHCP server is configured, DHCP-discover/request messages are discarded.

### Provisioning IP Configuration of the Host

The other aspect of subscriber-host authorization is providing IP configuration (ip-address, subnet-mask, default gateway and dns) through RADIUS directory rather than using centralized DHCP server. In this case, the node receiving following RADIUS attributes will assume role of DHCP server in conversation with the client and provide the IP configuration received from RADIUS server.

These attributes will be accepted only if the system is explicitly configured to perform DHCP-server functionality on a given interface.

The following RADIUS attributes will be accepted from authentication-response messages:

- framed-ip-address — The IP address to be configured for the subscriber-host.
- framed-ip-netmask — The IP network to be configured for the subscriber host. If RADIUS does not return a netmask, the DHCP request is dropped.
- framed-pool — The pool on a local DHCP server from which a DHCP-provided IP address should be selected.
- alc-default-router — The address of the default gateway to be configured on the DHCP client.
- alc-primary-dns — The DNS address to be provided in DHCP configuration.
  - Juniper VSA for primary DNS.
  - Redback VSA for primary DNS.
- alc-secondary-dns
  - Juniper VSA for secondary DNS.
  - Redback VSA for secondary DNS.
- alc-lease-time — Defines the lease time.
- session-timeout — Defines the lease time in absence of the alc-lease-time attribute.
- NetBIOS
  - alc-primary-nbns
  - alc-secondary-nbns

---

### RADIUS Based Authentication in Wholesale Environment

In order to support VRF selection, the following attributes are supported:

- alc-retail-serv-id — Indicates the service-id of the required retail VPRN service configured on the system.

## Change of Authorization and Disconnect-Request

In a typical RADIUS environment, the network element serves as a RADIUS client, which means the messages are originated by a routers. In some cases, such as “mid-session” changes, it is desirable that the RADIUS server initiates a CoA request to impose a change in policies applicable to the subscriber, as defined by RFC 3576.

To configure a RADIUS server to accept CoA and Disconnect Messages is achieved in one of the following ways:

1. Configure up to 64 RADIUS CoA servers per routing instance:

```
config>router>radius-server#
config>service>vprn>radius-server#

server "coa-1" address 10.1.1.1 secret <shared-secret> hash2 create
accept-coa
exit
```

This is the preferred method.

2. Configure up to 16 RADIUS CoA servers per authentication policy.

```
config>subscr-mgmt>auth-plcy#

accept-authorization-change
```

The UDP port for CoA and Disconnect Messages is configurable per system with the command:

```
config>aaa#

radius-coa-port {1647|1700|1812|3799}
```

Note that there is a priority in the functions that can be performed by CoA. The first matching one will be performed:

- If the CoA packet contains a force-renew attribute, the subscriber gets a force-renew DHCP packet. This function is not supported for PPPoE or ARP hosts.
- If the CoA packet contains a create-host attribute, a new lease-state is created. Only DHCP lease-states can be created by a CoA message. PPPoE sessions and ARP hosts cannot be created.
- Otherwise, the ESM strings are updated.

There are several reasons for using RADIUS initiated CoA messages:

1. Changing ESM attributes (SLA or subscriber profiles) or queues/policers/schedulers rates of the given subscriber host — CoA messages containing the identification of the given subscriber-host along with new ESM attributes.

2. Changing (or triggering the change) of IP configuration of the given subscriber-host — CoA messages containing the identification of the given subscriber-host along with VSA indicating request of forcerenew generation.
3. Configuring new subscriber-host — CoA messages containing the full configuration for the given host.

If the changes to ESM attributes are required, the RADIUS server will send CoA messages to the network element requesting the change in attributes included in the CoA request:

- attribute(s) to identify a single or multiple subscriber host(s): “NAS-Port-Id + IP address/prefix” or “Acct-Session-Id” or “Alc-Subsc-ID-Str”
  - Nas-Port-Id attribute + single IP address/prefix attribute:
    - Framed-IP-Address
    - Alc-Ipv6-Address
    - Framed-Ipv6-Prefix
    - Delegated-Ipv6-Prefix
  - Acct-Session-Id (number format)
  - Alc-Subsc-ID-Str
- alc-subscriber-profile-string
- alc-sla-profile-string
- alc-ancp-string
- alc-app-profile-string
- alc-int-dest-id-string
- alc-subscriber-id-string
- alc-subscriber-qos-override

Note that if the subscriber-id-string is changed while the ANCP string is explicitly set, the ANCP-string **must** be changed simultaneously. When changing the alc-subscriber-id-string, the lease state is temporarily duplicated, causing two identical ANCP-strings to be in the system at the same time. This is not allowed.

As a reaction to such message, the router changes the ESM settings applicable to the given host.

If changes to the IP configuration (including the VRF-id in the case of wholesaling) of the given host are needed, the RADIUS server may send a CoA message containing VSA indicating request for forcerenew generation:

- attribute(s) to identify a single or multiple subscriber host(s): “NAS-Port-Id + IP address/prefix” or “Acct-Session-Id” or “Alc-Subsc-ID-Str”:
  - Nas-Port-Id attribute + single IP address/prefix attribute:
    - Framed-IP-Address
    - Alc-Ipv6-Address



- Framed-Ipv6-Prefix
- Delegated-Ipv6-Prefix
- Acct-Session-Id (number format)
- Alc-Subsc-ID-Str
- alc-force-renew
- alc-force-nak

As a reaction to such message, router will generate a DHCP forcerenew message for the given subscriber host. Consequently, during the re-authentication, new configuration parameters can be populated based on attributes included in Authentication-response message. The force-NAK attribute has the same function as the force-renew attribute, but will cause the ESR to reply with a NAK to the next DHCP renew. This will invalidate the lease state on the ESR and force the client to completely recreate its lease, making it possible to update parameters that cannot be updated through normal CoA messages, such as IP address or address pool.

If the configuration of the new subscriber-host is required, RADIUS server will send a CoA message containing VSA request new host generation along with VSAs specifying all required parameters.

- alc-create-host
- alc-subscriber-id-string — This attribute is mandatory in case ESM is enabled, and optional for new subscriber host creation otherwise.
- NAS-port-id — This attribute indicates the SAP where the host should be created.
- framed-ip-address —
- alc-client-hw-address — A string in the xx:xx:xx:xx:xx:xx format. This attribute is mandatory for new subscriber-host creation.
- alc-lease-time — Specifies the lease time. If both session-timeout and alc-lease-time are not present, then a default lease time of 7 days is used.
- session-timeout — Specifies the lease time in absence of the alc-lease-time attribute. If both session-timeout and alc-lease-time are not present, then a default lease time of 7 days is used.
- alc-retail-svc-id — This is only used in case of wholesaling for selection of the retail service. Indicates the service-id of the required retail VPRN service configured on the system.
- Optionally other VSAs describing given subscriber host. Obviously, if the ESM is enabled, but the CoA message does not contain ESM attributes the new host will not be created.

After executing the requested action, the router element responds with an ACK or NAK message depending on the success/failure of the operation. In case of failure (and hence NAK response), the element will include the error code in accordance with RFC 3576 definitions if an appropriate error code is available.

Supporting CoA messages has security risks as it essentially requires action to unsolicited messages from the RADIUS server. This can be primarily the case in an environment where RADIUS servers from multiple ISPs share the same aggregation network. To minimize the security risks, the following rules apply:

## Provisioning of Enhanced Subscriber Management (ESM) Objects

- Support of CoA messages is disabled by default. They can be enabled on a per RADIUS server or authentication-policy basis.
- When CoA is enabled, the node will listen and react only to CoA messages received from RADIUS servers. In addition, CoA messages must be protected with the key corresponding to the given RADIUS server. All other CoA messages will be silently discarded.

In all cases (creation, modification, forcere Renew) subscriber host identification attributes are mandatory in the CoA request: “NAS-Port-Id + IP” or “Acct-Session-Id” or “Alc-Subsc-ID-Str”

- Nas-Port-Id + single IP address/prefix:
  - Nas-Port-Id
  - Framed-IP-Address
  - Alc-Ipv6-Address
  - Framed-Ipv6-Prefix
  - Delegated-Ipv6-Prefix
- Acct-Session-Id (number format)
- Alc-Subsc-ID-Str

When there are no subscriber host identification attributes present in the CoA, the message will be NAK'd with corresponding error code.

- hosts, meaning only subscriber host to which the given authentication-policy is applicable.
- Receiving CoA message with the same attributes as currently applicable to the given host will be responded to with an ACK message.
- In case of dual homing (through SRRP), the RADIUS server should send CoA messages to both redundant nodes and this with all corresponding attributes (NAS-port-id with its local meaning to corresponding node).
- In the case of change requests, the node which has the given host active (active-sap for master-sap for SRRP) will process the RADIUS message and reply to RADIUS. The standby node will always reply with a NAK.
- In the case of create requests, the active node (the SAP described by NAS-port-id is “active” or “master”). Both nodes will reply, but the standby will NAK the request.

The properties of an existing RADIUS-authenticated PPPoE session can be changed by sending a Change of Authorization (CoA) message from the RADIUS server. Processing of a CoA is done in the same way as for DHCP hosts, with the exception that only the ESM settings can be changed for a PPPoE session (the force-renew attribute is not supported for PPPoE sessions and a Create-Host CoA will always generate a DHCP host.)

For terminating PPPoE sessions from the RADIUS server, the disconnect-request message can be sent from the RADIUS server. This message triggers a shutdown of the PPPoE session. The attributes needed to identify the PPPoE session are the same as for DHCP hosts.

## RADIUS-Based Accounting

When a router is configured to perform RADIUS-based accounting, at the creation of a subscriber-host, it will generate an accounting-start packet describing the subscriber-host and send it to the RADIUS accounting server. At the termination of the session, it will generate an accounting-stop packet including accounting statistics for a given host. The router can also be configured to send an interim-accounting message to provide updates for a subscriber-host.

The exact format of accounting messages, their types, and communication between client running on the routers and RADIUS accounting server is described in RFC 2866, *RADIUS Accounting*. The following describes a few specific configurations.

In order to identify a subscriber-host in accounting messages different RADIUS attributes can be included in the accounting-start, interim-accounting, and accounting-stop messages. The inclusion of the individual attributes is controlled by the following commands.

```
configure
  subscriber-mgmt
    radius-accounting-policy <name>
      include-radius-attribute
        [no] acct-authentic
        [no] acct-delay-time
        [no] called-station-id
        [no] calling-station-id
        [no] circuit-id
        [no] delegated-ipv6-prefix
        [no] dhcp-vendor-class-id
        [no] framed-interface-id
        [no] framed-ip-addr
        [no] framed-ip-netmask
        [no] framed-ipv6-prefix
        [no] framed-route
        [no] framed-ipv6-route
        [no] ipv6-address
        [no] mac-address
        [no] nas-identifier
        [no] nas-port
        [no] nas-port-id
        [no] nas-port-type
        [no] nat-port-range
        [no] remote-id
        [no] sla-profile
        [no] sub-profile
        [no] subscriber-id
        [no] tunnel-server-attrs
        [no] user-name
        [no] wifi-rssi
        [no] alc-acct-triggered-reason
        [no] access-loop-options
        [no] all-authorized-session-addresses
        [no] detailed-acct-attributes
        [no] std-acct-attributes
        [no] v6-aggregate-stats
```

## Provisioning of Enhanced Subscriber Management (ESM) Objects

RADIUS volume accounting attributes are depending on the type of volume reporting and can be controlled via an **include-radius-attribute** CLI command. Multiple volume reporting types can be enabled simultaneously:

```
configure
  subscriber-mgmt
    radius-accounting-policy <name>
      include-radius-attribute
        [no] detailed-acct-attributes
        [no] std-acct-attributes
        [no] v6-aggregate-stats
```

where:

**detailed-acct-attributes** — Report detailed per queue and per policer counters using RADIUS VSAs (enabled by default). Each VSA contains a queue or policer id followed by the stat-mode or 64 bit counter. The VSA's included in the Accounting messages is function of the context (policer or queue, stat-mode, MDA type, ...):

```
[26-6527-107] Alc-Acct-I-statmode
[26-6527-127] Alc-Acct-O-statmode
[26-6527-19] Alc-Acct-I-Inprof-Octets-64
[26-6527-20] Alc-Acct-I-Outprof-Octets-64
[26-6527-21] Alc-Acct-O-Inprof-Octets-64
[26-6527-22] Alc-Acct-O-Outprof-Octets-64
[26-6527-23] Alc-Acct-I-Inprof-Pkts-64
[26-6527-24] Alc-Acct-I-Outprof-Pkts-64
[26-6527-25] Alc-Acct-O-Inprof-Pkts-64
[26-6527-26] Alc-Acct-O-Outprof-Pkts-64
[26-6527-39] Alc-Acct-OC-O-Inprof-Octets-64
[26-6527-40] Alc-Acct-OC-O-Outprof-Octets-64
[26-6527-43] Alc-Acct-OC-O-Inprof-Pkts-64
[26-6527-44] Alc-Acct-OC-O-Outprof-Pkts-64
[26-6527-69] Alc-Acct-I-High-Octets-Drop_64
[26-6527-70] Alc-Acct-I-Low-Octets-Drop_64
[26-6527-71] Alc-Acct-I-High-Pack-Drop_64
[26-6527-72] Alc-Acct-I-Low-Pack-Drop_64
[26-6527-73] Alc-Acct-I-High-Octets-Offer_64
[26-6527-74] Alc-Acct-I-Low-Octets-Offer_64
[26-6527-75] Alc-Acct-I-High-Pack-Offer_64
[26-6527-76] Alc-Acct-I-Low-Pack-Offer_64
[26-6527-77] Alc-Acct-I-Unc-Octets-Offer_64
[26-6527-78] Alc-Acct-I-Unc-Pack-Offer_64
[26-6527-81] Alc-Acct-O-Inprof-Pack-Drop_64
[26-6527-82] Alc-Acct-O-Outprof-Pack-Drop_64
[26-6527-83] Alc-Acct-O-Inprof-Octs-Drop_64
[26-6527-84] Alc-Acct-O-Outprof-Octs-Drop_64
[26-6527-91] Alc-Acct-OC-O-Inpr-Pack-Drop_64
```

```

[26-6527-92] Alc-Acct-OC-O-Outpr-Pack-Drop_64
[26-6527-93] Alc-Acct-OC-O-Inpr-Octs-Drop_64
[26-6527-94] Alc-Acct-OC-O-Outpr-Octs-Drop_64
[26-6527-108] Alc-Acct-I-Hiprio-Octets_64
[26-6527-109] Alc-Acct-I-Lowprio-Octets_64
[26-6527-110] Alc-Acct-O-Hiprio-Octets_64
[26-6527-111] Alc-Acct-O-Lowprio-Octets_64
[26-6527-112] Alc-Acct-I-Hiprio-Packets_64
[26-6527-113] Alc-Acct-I-Lowprio-Packets_64
[26-6527-114] Alc-Acct-O-Hiprio-Packets_64
[26-6527-115] Alc-Acct-O-Lowprio-Packets_64
[26-6527-116] Alc-Acct-I-All-Octets_64
[26-6527-117] Alc-Acct-O-All-Octets_64
[26-6527-118] Alc-Acct-I-All-Packets_64
[26-6527-119] Alc-Acct-O-All-Packets_64

```

**std-acct-attributes** — Report IPv4 and IPv6 aggregated forwarded counters using standard RADIUS attributes (disabled by default):

```

[42] Acct-Input-Octets
[43] Acct-Output-Octets
[47] Acct-Input-Packets
[48] Acct-Output-Packets
[52] Acct-Input-Gigawords
[53] Acct-Output- Gigawords

```

**v6-aggregate-stats** — Report IPv6 aggregated forwarded counters of queues and policers in stat-mode v4-v6 using using RADIUS VSAs (disabled by default):

```

[26-6527-194] Alc-IPv6-Acct-Input-Packets
[26-6527-195] Alc-IPv6-Acct-Input-Octets
[26-6527-196] Alc-IPv6-Acct-Input-GigaWords
[26-6527-197] Alc-IPv6-Acct-Output-Packets
[26-6527-198] Alc-IPv6-Acct-Output-Octets
[26-6527-199] Alc-IPv6-Acct-Output-Gigawords

```

In addition to accounting-start, interim-accounting, and accounting-stop messages, a RADIUS client on a routers will send also accounting-on and accounting-off messages. An accounting-on message will be sent when a given RADIUS accounting-policy is applied to a given subscriber-profile, or the first server is defined in the context of an already applied policy. The following attributes will be included in such message:

- NAS-identifier
- alc-subscriber-profile-string
- Accounting-session-id
- Event-timestamp

## Provisioning of Enhanced Subscriber Management (ESM) Objects

Accounting-off messages will be sent at following events:

- An accounting policy has been removed from a sub-profile.
- The last RADIUS accounting server has been removed from an already applied accounting policy.

These messages contain following attributes:

- NAS-identifier
- alc-subscriber-profile-string
- Accounting-session-id
- Accounting-terminate-cause
- Event-timestamp

In case of dual homing, both nodes will send RADIUS accounting messages for the host, with all attributes as it is locally configured. The RADIUS log files on both boxes need to be parsed to get aggregate accounting data for the given subscriber host regardless the node used for forwarding.

For RADIUS-based accounting, a custom record can be defined to refine the data that is sent to the RADIUS server. Refer to the Configuring an Accounting Custom Record in the OS System Management Guide for further information.

## Accounting Modes Of Operation

This section is applicable to the 7750 SR or the 7450 ESS in mixed mode

There are three basic accounting models in 7750 SR or the 7450 ESS in mixed mode:

- Per queue-instance
- Per Host
- Per Session

Each of the basic models can optionally be enabled to send interim-updates. Inclusion/exclusion of interim-updates will depend on whether volume based (start/interim-updates/stop) or time based (start/stop) accounting is required.

The difference between the three basic accounting models is in its core related to the processing of the acct-session-id for each model. The differences are related to:

- acct-session-id generation within each model.
- outcome in response to the CoA action relative to the targeted acct-session-id.

The counters for volume-based accounting are collected from queues or policers that are instantiated per sla-profile instance (SPI) on non-HSMDA based hardware or per subscriber on HSMDA based hardware. This is true irrespective of which model of accounting (or combination of models) is deployed. Within accounting context, the SPI on non-HSMDA or subscriber on HSMDA equates to queue-instance.

Table 12 summarizes the key differences between various accounting modes of operation that are supported. Interim-updates for each individual mode can be enabled/disabled via configuration (interim-updates keyword as an extension to the commands that enable three basic modes of accounting). This is denoted by the IU-Config keyword under the 'I-U' column in the table. The table also shows that any two combinations of the three basic models (including their variants for volume/time based accounting) can be enabled simultaneously.

**Table 12: Accounting Modes of Operation**

Accounting Mode	Accounting Entity	START	I-U	STOP	Acct-session-id	Acct-multi-session-id
queue-instance-accounting	queue-instance	X	IU-config	X	X	
	session					
	host					

**Table 12: Accounting Modes of Operation (Continued)**

Accounting Mode	Accounting Entity	START	I-U	STOP	Acct-session-id	Acct-multi-session-id
session-accounting	queue-instance					
	session	X	IU-config	X	X	q-instance
	host					
host-accounting	queue-instance					
	session					
	host	X	IU-config	X	X	queue-instance
queue-instance-accounting + host-accounting	queue-instance	X	IU-config	X	X	queue-instance
	session					
	host	X	IU-config	X	X	queue-instance
queue-instance-accounting + session-accounting	queue-instance	X	IU-config	X	X	queue-instance
	session	X	IU-config	X	X	queue-instance
	host					
session-accounting + host-accounting	queue-instance					queue-instance
	session	X	IU-config	X	X	
	host	X		X	X	SESSION

Note that hosts within the targeted CoA entity will be affected as follows:

- If the CoA target is the session, then both constituting members (IPv4 and IPv6) of the dual-stack host will be affected.
- If the CoA target is the queuing-instance, then up to 32 hosts that are sharing that SPI will be affected.



The same principle applies to LI.

The accounting behavior (accounting messages and accounting attributes) in case that the SPI is changed via CoA depends on the accounting mode of operation. On non-HSMDA hardware, the behavior is the following:

- SPI change in conjunction with per queuing instance accounting will trigger a STOP for the old SPI and a START for the new SPI with corresponding counters. Acct-session-id/Acct-Multi-Session-Id will be unique per SPI. Note that Acct-Multi-Session-Id is only generated if per queuing-instance accounting mode of operation is combined with some other mode of operation (host or session).
- SPI change in conjunction with per host or per session accounting (no interim updates for either method) will NOT trigger any new accounting messages. In other words, SPI change will go unnoticed from the perspective of the accounting server until the host/session is terminated. When the host/session is terminated a STOP will be sent with the VSA carrying the latest SPI name and the acct-multi-session-id attribute of the latest SPI. Acct-session-id will stay the same during the lifetime of the host. Counters are not included in STOP (interim-update not enabled).

SPI change in conjunction with per host accounting with interim-updates or per session accounting with interim-updates will trigger two interim-update messages:

- One with the old counters (terminated queues) and the old SPI name VSA. This behavior is similar to the triggered STOP message in per queuing-instance accounting upon SPI change.
- One with the new counters (new queues instantiated), the VSA carrying the new SPI name and the new acct-multi-session-id referencing the new SPI. This behavior is similar to the triggered START message in per queuing-instance when SPI is changed in per queuing-instance accounting.

On HSMDA, no START/STOPS are sent since queues are not re-instantiated on ingress or egress.

## Per Session Accounting

In the per session accounting mode of operation the accounting message stream<sup>1</sup> (START/INTERIM-UPDATE/STOP) is generated per session.

- A session is defined for PPPoE hosts for which a state is maintained. The state of the host (single stack or dual-stack) is normally refreshed via PPP keepalives. Each PPPoE host of the same address family (v4 or v6) corresponds to a unique session which is identified by the <session-ID, mac> combination.

In dual-stack PPPoE case, IPv4 and IPv6 hosts are tied to the same (LCP) session. A single authentication request is initiated for such session (triggered by the first host that initiates the session).

For a single stack PPPoE host, the behavior defined in the per session accounting model is indistinguishable from the per host accounting model. The per session accounting model makes difference in behavior only for dual stack PPPoE hosts.

The following are the properties of the Per Session Accounting model:

- A single accounting session ID (acct-session-id) is generated per (PPPoE) session and it can optionally be sent in RADIUS Access-Request message.
- This acct-session-id is synchronized via MCS in dual homing environment.
- The accounting messages (START, INTERIM-UPDATE, STOP) carry the acct-multi-session-id attribute denoting the sla-profile instance with which the session is associated.
- The counters are collected from the queues instantiated through the sla-profile instance. If multiple sessions are sharing the same sla-profile instance, the counters are aggregated. In other words, counters per individual session cannot be extracted from the aggregated count.
- RADIUS triggered changes and LI are applicable per session:
  - Queue/policer RADIUS overrides — Parameters for the referenced queue/policer within the session are changed accordingly.
  - Subscriber aggregate rate limits, scheduler rates and arbiter rates are changed accordingly.
  - CoA DISCONNECT brings down the entire session.
  - LI — Activation based on the session acct-session-id affects the hosts within the session (dual-stack).
  - SLA profile instance change affects all hosts (or sessions) sharing the same sla-profile instance (SPI). If the SPI is changed on a non-HSMDA based MDA, then queues are re-instantiated and counters are reset.
- All applicable IP addresses (v4 and v6 – including all v6 attributes – alc-ipv6-address, framed-ipv6-prefix, delegated-ipv6-prefix) are present in accounting messages for the session.

---

1. The accounting message stream refers to a collection of accounting messages (START/INTERIM-UPDATE/STOP) sharing the same acct-session-id.

## **Caveats**

Per session accounting is supported for entities that have concept of a session. Currently only PPPoE hosts (single or dual-stack) fall into this category.

## RADIUS Per Host Accounting

In SR-OS, the accounting paradigm is based on SLA profile instances yet this is at odds with traditional RADIUS authentication and accounting which is host-centric. In previous SR-OS releases, it was possible to have many hosts sharing a common SLA profile instance, and thus accounting and QoS parameters. Complications would arise with RADIUS accounting because Accounting-Start and Accounting-Stop are a function of sla-profile instance and not the hosts — this meant that some host-specific parameters (like framed-ip-address) would not be consistently included in RADIUS accounting.

Currently, dual-stack subscribers are really two different hosts sharing a single sla-profile instance. A new RADIUS accounting mode has been introduced to support multiple-host environments.

Under accounting-policy, a host-accounting command allows configurable behavior.

---

## No Host-Accounting

In prior releases and when no host-accounting is configured, the accounting behavior is as follows:

- A RADIUS accounting start message is sent when the SLA-profile instance is created. It contains accounting (octets/packets) and the framed-ip-address of the host which caused the sla-profile instance to be created.
  - Additional hosts may bind to the sla-profile instance at any time, but no additional Accounting messages are sent during these events.
  - If the original host disconnects then future Accounting messages will use an IP address of one of the remaining hosts.
  - When the final host associated with an sla-profile instance disconnects an Accounting Stop message will be sent.
- 

## Host-Accounting Enabled

When host-accounting is configured, additional RADIUS accounting messages are created for host activity in addition to messages for common queue accounting. The behavior is as follows:

- A RADIUS accounting start message is sent each time a host is authenticated. It contains the framed-ip-address among other things. It does not contain any octet or packet counts.
- A RADIUS accounting start message is sent each time a sla-profile instance is created.
- Whenever a host disconnects a RADIUS accounting stop message is sent for that host.
- If all host associated with an sla-profile instance disconnect, a RADIUS Accounting Stop message is sent for that instance.

This new behavior means certain AVP may be in either host; sla-profile instance or both accounting records.

Note that interim-acct records are not sent for hosts, only the start- and stop-accting messages.

<b>RADIUS Accounting AVP</b>	<b>Include-radius-attrs Acct/ Auth</b>	<b>Host Accounting</b>	<b>SLA-Profile Accounting</b>
User-Name	Yes/No	Yes	No
NAS-Identifier	Yes/No	Yes	Yes
NAS-Ip-Address	No/No	Yes	Yes
NAS-Port-Id	Yes/Yes	Yes	No
Nas-Port	Yes/No	Yes	No
NAS-Port-Type	Yes/Yes	Yes	No
Service-Type	No/No	Yes	No
Framed-Protocol	No/No	Yes	No
Framed-Ip-Address	Yes/No	Yes	No
Framed-Ip-Netmask	Yes/No	Yes	No
Framed-Route	No/No	Yes	No
Class	No/No	Yes	No
Session-Timeout	No/No	Yes	Yes
Circuit-Id VSA	Yes/Yes	Yes	No
Called-Station-Id	Yes/Yes	Yes	No
Calling-Station-Id	Yes/Yes	Yes	No
MAC-Addr VSA	Yes/Yes	Yes	No
Remote-Id VSA	Yes/Yes	Yes	No
Acct-Input-Octets	No/No	No	Yes
Acct-Output-Octets	No/No	No	Yes
Acct-Input-Gigawords	No/No	No	Yes
Acct-Output-Gigawords	No/No	No	Yes
Acct-Session-Id	No/No	Yes	Yes
Acct-Session-Time	No/No	Yes	Yes

## Provisioning of Enhanced Subscriber Management (ESM) Objects

<b>RADIUS Accounting AVP</b>	<b>Include-radius-attrs Acct/ Auth</b>	<b>Host Accounting</b>	<b>SLA-Profile Accounting</b>
Acct-Input-Packets	No/No	No	Yes
Acct-Output-Packets	No/No	No	Yes
Acct-Multi-Session-Id	No/No	Yes	Yes
Actual-Data-Rate-Upstream	No/No	Yes	No
Actual-Data-Rate-Downstream	No/No	Yes	No
Access-Loop-Encapsulation	No/Yes	Yes	No
Alc-Accounting	No/No	No	Yes
Alc-Subscriber-Id	Yes/No	Yes	Yes
Alc-Subscriber-Profile-String	Yes/No	Yes	Yes
Alc-Sla-Profile-String	Yes/No	Yes	Yes

## Accounting Interim Update Message Interval

The interval between two RADIUS Accounting Interim Update messages can be configured in the RADIUS accounting policy with the **update-interval** command, for example:

```
configure
  subscriber-mgmt
    radius-accounting-policy "acct-policy-1" create
      update-interval 60
      update-interval-jitter absolute 600
```

A RADIUS specified interim interval (attribute [85] Acct-Interim-Interval) overrides the CLI configured value.

By default, a random delay of 10% of the configured **update-interval** is added to the update-interval between two Accounting Interim Update messages. This jitter value can be configured with the **update-interval-jitter** to an absolute value in seconds between zero and 3600. The effective maximum random delay value is the minimum value of the configured absolute jitter value and 10% of the configured **update-interval**.

A value of zero will send the Accounting Interim Update message without introducing an additional random delay.

---

## Class Attribute

The RADIUS class-attribute helps to aid in user identification

User identification is used to correlate RADIUS accounting messages with the given user. During authentication process, the RADIUS authentication server inserts a class-attribute into the RADIUS authenticate response message and then the router echoes this class attribute in all RADIUS accounting messages.

---

## User Name

The user-name, which is used for user authentication (user-name attribute in RADIUS authentication request), can be included in RADIUS accounting messages. Per RFC 2865, when a RADIUS server returns a (different) user-name attribute, the changed user name will be used in accounting and not the originally sent user name.

## Accounting-On and Accounting Off

For RADIUS servers configured in a RADIUS server policy, the accounting on/off behavior is controlled via the **acct-on-off** command in the radius-server-policy.

By default, no Accounting-On or Accounting-Off messages are sent (**no acct-on-off**).

With the **acct-on-off** command configured in the radius-server-policy:

- An Accounting-On is sent for the following:
  - When the system is powered on.
  - After a system reboots.
  - When the **acct-on-off** command is added to the **radius-server-policy** configuration.
  - User triggered via CLI: tools perform `aaa acct-on`
- An Accounting-Off is sent for the following:
  - Before a user initiated system reboot.
  - When the **acct-on-off** command is removed from the **radius-server-policy** configuration.
  - User triggered via CLI: tools perform `aaa acct-off`.

The Accounting-On or Accounting-Off message is sent to the servers configured in the radius-server-policy, following the configured access-algorithm until an Accounting Response is received. If the first server responds, no message is sent to the other servers.

The Accounting-On message is repeated until an Accounting Response message is received from a RADIUS server: If after the configured retry/timeout timers for each RADIUS server in the radius-server-policy no response is received then the process starts again after a fixed one minute wait interval.

The Accounting-Off message is attempted once: If after the configured retry/timeout timers for each RADIUS server in the radius-server-policy no response is received then no new attempt is made.

It is possible to block a radius-server-policy until an Accounting Response is received from one of the RADIUS servers in the radius-server-policy that acknowledges the reception of an Accounting-On. The radius-server-policy cannot be used by applications for sending RADIUS messages until the state becomes “Not Blocked”. This is achieved with the optional “oper-state-change” flag, for example:

```
configure
aaa
    radius-server-policy "aaa-server-policy-1" create
        acct-on-off oper-state-change
        servers
            router "Base"
            server 1 name "server-1"
        exit
    exit
exit
```



If multiple radius-server-policies are in use for different applications (for example, authentication and accounting) and an Accounting-On must be sent for only one radius-server-policy, it is possible to tie the acct-on-off states of both policies together using an acct-on-off-group. With this configuration, it is possible to block the authentication servers until the accounting servers are available. An acct-on-off-group can be referenced by:

- a single radius-server-policy as controller: the acct-on-off oper-state of the acct-on-off-group is set to the acct-on-off oper-state of the radius-server-policy (acts as master)
- multiple radius-server-policies as monitor: the acct-on-off oper-state of the radius-server-policy is inherited from the acct-on-off oper-state of the acct-on-off group. (acts as a slave)

```
configure
aaa
    acct-on-off-group "group-1" create
        description "Grouping of radius-server-policies acct-on-off"
    exit
    radius-server-policy "aaa-server-policy-1" create
        acct-on-off oper-state-change group "group-1"
    servers
        router "Base"
        server 1 name "server-1"
    exit
    radius-server-policy "aaa-server-policy-2" create
        acct-on-off monitor-group "group-1"
    servers
        router "Base"
        server 1 name "server-2"
    exit
exit
```

It is possible to force an Accounting-On or Accounting-Off message for a radius-server-policy with acct-on-off enabled using following CLI commands:

**tools perform aaa acct-on [radius-server-policy *policy-name*] [force]**

**tools perform aaa acct-off [radius-server-policy *policy-name*] [force] [acct-terminate-cause *number*]**

If an Accounting-On was sent to the radius-server-policy and it was acknowledged with an Accounting Response then a new Accounting-On can only be sent with the “force” flag.

If an Accounting-Off was sent to the radius-server-policy and it was acknowledged with an Accounting Response then a new Accounting-Off can only be sent with the “force” flag. The Acct-Terminate-Cause value in the Accounting-Off can be overwritten.

Use the following CLI command to display the Accounting On/Off information for a radius-server-policy:

```
# show aaa radius-server-policy "aaa-server-policy-3" acct-on-off
=====
RADIUS server policy "aaa-server-policy-3" AcctOnOff info
=====
Oper state           : on
Session Id          : 242FFF0000008F512A3985
Last state change   : 02/24/2013 16:06:41
Trigger             : startUp
Server              : "server-1"
=====
```

The operational state provides following state information: The sending of the Accounting-On or Accounting-Off message is ongoing (sendAcctOn, SendAcctOff), is successfully responded (on, off) or no response received (OffNoResp).

The Session-Id is a unique identifier for each RADIUS server policy accounting Accounting-On/ Accounting-Off sequence.

The Trigger field shows what triggered the Accounting On or Accounting Off message. If the radius-server-policy is part of an acct-on-off group then the group name is shown in brackets.

The Server field shows which server in the RADIUS server policy responded to the Accounting-On or Accounting-Off message.

To display the acct-on-off state of a radius-server-policy, use the command, for example:

```
# show aaa radius-server-policy "aaa-server-policy-3"
=====
RADIUS server policy "aaa-server-policy-3"
=====
Description           : (Not Specified)
Acct Request script policy : script-policy-1
Auth Request script policy : script-policy-1
Accept script policy    : script-policy-1
Acct-On-Off            : Enabled (state Blocked)
-----
RADIUS server settings
-----
Router                 : "Base"
Source address          : (Not Specified)
Access algorithm        : direct
Retry                  : 3
Timeout (s)            : 5
Hold down time (s)     : 30
Last management change  : 02/20/2013 13:32:05
=====
Servers for "aaa-server-policy-3"
=====
```

Idx	Name	Address	Port Auth/Acct	Oper State
1	server-3	172.16.1.10	1812/1813	unknown

```
=====
```

The Acct-On-Off field indicates if the sending of Accounting-On and Accounting-Off messages is enabled or disabled. If enabled, the oper-state is displayed: state Blocked or state Not Blocked. When Blocked, the radius-server-policy cannot be used to send RADIUS messages.

To display acct-on-off-group information, use following command, for example:

```
# show aaa acct-on-off-group "group-1"
=====
Acct-On-Off-Group Information
=====
acct on off group name          : group-1
- controlling Radius-Server-policy :
    aaa-server-policy-1
- monitored by Radius-Serer-policy :
    aaa-server-policy-2

-----
Nbr of Acct-on-off-groups displayed : 1
-----
=====
```

## RADIUS Accounting Message Buffering


When all servers in a radius-server-policy are unreachable, it is possible to buffer the Accounting Stop and Accounting Interim-Update messages for up to 25 hours. When a RADIUS server becomes reachable again then the messages in the buffer are retransmitted.

RADIUS Accounting message buffering parameters can be configured per message type, for example:

```
configure
aaa
    radius-server-policy "aaa-server-policy-1" create
        servers
            router "Base"
            buffering
                acct-interim min 60 max 3600 lifetime 12
                acct-stop min 60 max 3600 lifetime 12
            exit
        server 1 name "server-1"
    exit
exit
exit
```

When RADIUS accounting message buffering is enabled:

1. The message is stored in the buffer, a lifetime timer is started and the message is sent to the RADIUS server.
2. If after  $\text{retry} \times \text{timeout}$  seconds no RADIUS accounting response is received for the Accounting Interim Update or Accounting Stop then a new attempt to send the message is started after minimum  $[(\text{min-val} \times 2n), \text{max-val}]$  seconds.
3. Repeat step 2 until:
  - a. RADIUS accounting response is received, or
  - b. the lifetime of the buffered message expires, or
  - c. (if the buffered message is an Accounting Interim-Update only) A new Accounting Interim-Update or an Accounting Stop or for the same accounting session-id and radius-server-policy is stored in the buffer, or
  - d. the message is manually purged from the message buffer via a clear command
4. The message is purged from the buffer as shown in [Figure 43](#).



**Figure 43: Purging Message from Buffer**

When Accounting Interim-Update message buffering is enabled, it is recommended to also enable Accounting Stop message buffering. This will guarantee the message ordering per accounting session.

Use following clear command to manually delete messages from the RADIUS accounting message buffer:

**# clear aaa radius-server-policy *policy-name* msg-buffer [*acct-session-id acct-session-id*]**

When specifying the Acct-Session-Id, only that specific message will be deleted from the message buffer. If no Acct-Session-Id is specified, all messages for that radius-server-policy are deleted from the message buffer.

Use the following show commands to display the RADIUS accounting message buffer statistics:

```
# show aaa radius-server-policy "aaa-server-policy-1" msg-buffer-stats
=====
RADIUS server policy "aaa-server-policy-1" message buffering stats
=====
buffering acct-interim      : enabled
  min interval (s)         : 60
  max interval (s)         : 3600
  lifetime (hrs)           : 12
buffering acct-stop        : enabled
  min interval (s)         : 60
  max interval (s)         : 3600
  lifetime (hrs)           : 12

Statistics
-----
```

## Provisioning of Enhanced Subscriber Management (ESM) Objects

```
Total acct-stop messages in buffer          : 0
Total acct-interim messages in buffer        : 5
Total acct-stop messages dropped (lifetime expired) : 0
Total acct-interim messages dropped (lifetime expired) : 0
Last buffer clear time                      : N/A
Last buffer statistics clear time           : N/A
-----
=====
```

Use following clear command to reset the RADIUS accounting message buffer statistics:

**# clear aaa radius-server-policy *policy-name* statistics msg-buffer-only**

Use following tools commands to display the RADIUS accounting message buffer content:

**# tools dump aaa radius-server-policy *policy-name* msg-buffer [session-id *acct-session-id*]**

For example:

```
# tools dump aaa radius-server-policy "aaa-server-policy-1" msg-buffer
=====
RADIUS server policy "aaa-server-policy-1" message buffering
=====
message type Acct-Session-Id                      remaining lifetime
-----
acct-interim 242FFF0000009A512B36FC                0d 11:58:54
acct-interim 242FFF0000009B512B36FC                0d 11:58:48
acct-interim 242FFF0000009C512B36FC                0d 11:58:30
acct-interim 242FFF0000009D512B36FC                0d 11:58:29
acct-interim 242FFF0000009E512B36FC                0d 11:59:05
-----
No. of messages in buffer: 5
=====
```

When specifying the Acct-Session-Id, the message details are displayed.

## Sending an Accounting Stop Message upon a RADIUS Authentication Failure of a PPPoE Session

In scenarios where RADIUS authentication is used for PPPoE sessions, an accounting stop message can be generated to notify the RADIUS servers in case of an authentication failure.

The failure events are categorized in three categories:

- **“on-request-failure”** — All failure conditions between the sending of an Access-Request and the reception of an Access-Accept or Access-Reject.
- **“on-reject”** — When an Access-Reject is received.
- **“on-accept-failure”** — All failure conditions that appear after receiving an Access-Accept and before successful instantiation of the host or session.

Each of the categories can be enabled separately in the RADIUS authentication policy.

In the Enhanced Subscriber Management (ESM) model, the RADIUS accounting server is found after authentication and host identification as part of the subscriber profile configuration. To report authentication failures to accounting servers, an alternative RADIUS accounting policy configuration is required: local user database pre-authentication can provide the RADIUS authentication policy to be used for authentication and the RADIUS accounting policy to be used for authentication failure reporting. A duplicate RADIUS accounting policy can be specified if the accounting stop resulting from a RADIUS authentication failure must also be sent to a second RADIUS destination.

```
configure
subscriber-mgmt
  local-user-db "ludb-1" create
  ppp
    match-list username
    host "default" create
      auth-policy "auth-policy-1"
      password ignore
      acct-policy "acct-policy-1" duplicate "acct-policy-2"
      no shutdown
    exit
  exit
  no shutdown
exit
authentication-policy "auth-policy-1" create
  pppoe-access-method pap-chap
  include-radius-attribute
    - - - snip - - -
  exit
  send-acct-stop-on-fail on-request-failure on-reject on-accept-failure
  radius-server-policy "aaa-server-policy-1"
exit
radius-accounting-policy "acct-policy-1" create
  - - - snip - - -
  radius-server-policy "aaa-server-policy-1"
exit
radius-accounting-policy "acct-policy-2" create
  - - - snip - - -
```

## Provisioning of Enhanced Subscriber Management (ESM) Objects

```
radius-server-policy "aaa-server-policy-2"  
exit
```

To enable local user database pre-authentication, use the user-db configuration in the capture SAP and in the group-interface. For example:

```
configure  
service  
  vpls 10 customer 1 create  
    sap 1/1/1:1.* capture-sap create  
      trigger-packet pppoe  
      pppoe-policy "ppp-policy-1"  
      pppoe-user-db "ludb-1"  
    exit  
  no shutdown  
exit  
ies 1000 customer 1 create  
  subscriber-interface "sub-int-1" create  
    - - - snip - - -  
    group-interface "group-int-1-1" create  
      - - - snip - - -  
      pppoe  
        policy "ppp-policy-1"  
        user-db "ludb-1"  
        no shutdown  
      exit  
    exit  
  exit  
no shutdown  
exit
```



# Enhanced Subscriber Management Overview

---

## Enhanced Subscriber Management Basics

In residential broadband networks numerous subscribers can be provisioned that can require significant changes on a daily basis. Manually configuring the applicable parameters for each subscriber would be prohibitive. The Alcatel-Lucent 7450 ESS has been designed to support fully dynamic provisioning of access, QoS and security aspects for residential subscribers using DHCP to obtain an IP address. Enabling Enhanced Subscriber Management drastically reduces the configuration burden.

Enhanced Subscriber Management in the 7450 ESS supports many vendor's access nodes and network aggregation models, including VLAN per customer, per service or per access node.

---

## Standard and Enhanced Subscriber Management

The system can switch between standard and enhanced subscriber management modes on a per SAP basis. The Enhanced Subscriber Management mode is supported on the SR-7 and SR-12 chassis and on the ESS-7 chassis.

Some functions are common between the standard and enhanced modes. These include DHCP lease management, static subscriber host definitions and anti-spoofing. While the functions of these features may be similar between the two modes, the behavior is considerably different.

- Standard mode — The system performs SLA enforcement functions on a per SAP basis, that is, the attachment to a SAP with DHCP lease management capabilities. The node can authenticate a subscriber session with RADIUS based on the MAC address, the circuit-id (from Option 82) or both. It will then maintain the lease state in a persistent manner. It can install anti-spoofing filters and ARP entries based on the DHCP lease state. Static subscriber hosts are not required to have any SLA or subscriber profile associations and are not required to have a subscriber identification string defined.
- Enhanced mode — When enabled on a SAP, the system expands the information it stores per subscriber host, allowing SLA enforcement and accounting features on a per subscriber basis. The operator can create a subscriber identification policy that will include a URL to a user-space script that assists with the subscriber host identification process.
  - A subscriber host is identified by a subscriber identification string instead of the limited Option 82 values (although, the identification string is normally derived from string manipulation of the Option 82 fields). A subscriber identification policy is used to process the dynamic host DHCP events to manage the lease state information

stored per subscriber host. The static subscriber hosts also must have subscriber identification strings associations to allow static and dynamic hosts to be grouped into subscriber contexts.

- Further processing by the subscriber identification policy derives the appropriate subscriber and SLA profiles used to define the hierarchical virtual schedulers for each subscriber and the unique queuing and filtering required for the hosts associated with each subscriber
- The SLA profile information is used to identify which QoS policies and which queues/policers, and also which egress hierarchical virtual schedulers, will be used for each subscriber host (dynamic or static).
- The system performs SLA enforcement functions on a per subscriber SLA profile instance basis. SLA enforcement functions include QoS (classification, filtering and queuing), security (filtering), and accounting.

When the enhanced mode is enabled on a SAP (see [Subscriber SAPs on page 775](#)), first, the router ensures that existing configurations on the SAP do not prevent proper enhanced mode operation. If any one of the following requirements is not met, enhanced mode operation is not allowed on the SAP:

- Anti-spoofing filters must be enabled and configured as IP+MAC matching.
- Any existing static subscriber hosts must have:
  - An assigned subscriber identification string.
  - An assigned subscriber profile name.
  - An assigned SLA profile name.
- The system must have sufficient resources to create the required SLA profile instances and schedulers.

When the router successfully enables the enhanced mode, the current dynamic subscriber hosts are not touched until a DHCP message event occurs that allows re-population of the dynamic host information. Thus, over time, the dynamic subscriber host entries are moved from SAP-based queuing and SAP-based filtering to subscriber-based queuing and filtering. In the event that a dynamic host event cannot be processed due to insufficient resources, the DHCP ACK message is discarded and the previous host lease information is retained in the system.

## Subscriber Management Definitions

---

### Subscriber

A subscriber is typically defined by a unique subscriber identifier to which an assortment of policies (or subscriber profile) can be applied. A subscriber typically (but not always) maps into a VLAN, a VPI/VCI pair, an “ifentry” (a logical interface such as a SAP), a (source) MAC or IP address or a physical port, which uniquely identify a billable entity for the service provider.

---

### Subscriber Management

The management of all services, policies, AAA functions and configurations that relate to the concept of a subscriber. Subscriber management can be configured in a variety of ways, but it is critical that subscriber management integrates seamlessly with element and service management across the broadband infrastructure, via for instance, the Alcatel-Lucent 5750 Subscriber Services Controller (SSC). Subscriber management can also be implemented through CLI or scripted commands at the platform level, whereby a network administrator would manually configure the set of QoS, security, AAA or anti-spoofing functions that relate to a particular billable entity or subscriber. Subscriber management is typically centralized and highly integrated with the element, services and middleware management functions for streamlined management, flowthrough provisioning, and accelerated service activation, with minimized operating expenditures.

---

### Subscriber Policy Enforcement

Is the set of actual enforcement functions that are implemented relative to a given subscriber, possibly at multiple enforcement points in the infrastructure and as a result of a match between the subscriber profile which was defined by the subscriber management suite (Alcatel-Lucent’s 5750 SSC) and actual traffic patterns. Examples include for instance, the shaping, policing or rate limiting of traffic or the traffic of a given subscriber being dropped because it matched or violated any specific rule (packet with a mismatch between MAC and IP address suggesting an address spoof for instance)

---

### Subscriber SAPs

A subscriber SAP is a service access point (SAP) where enhanced subscriber management is active. Enhanced subscriber management must be explicitly enabled on a per-SAP basis with the CLI **sub-sla-mgmt** command.

A subscriber SAP can be used by a single subscriber or support multiple subscribers simultaneously. Each subscriber can be represented by one or multiple subscriber hosts on the subscriber SAP. If enhanced subscriber management is enabled on a SAP, any configured QoS and

IP filter policies defined on the SAP are ignored. A subscriber SAP must refer to an existing subscriber identification policy.

---

### Hosts and Subscribers

A host is a device identified by a unique combination of IP address and MAC address. Typically, the term “subscriber host” is used instead of the “host”.

A host can be an end-user device, such as a PC, VoIP phone or a set top box, or it can be the user’s Residential Gateway (RGW) if the RGW is using Network Address Translation (NAT).

Each subscriber host must be either statically provisioned or dynamically learned by the system. The host’s IP address + MAC address are populated in the subscriber host table on the appropriate SAP to allow packets matching the IP address and MAC address access to the provider’s network.

- A dynamic subscriber host is dynamically learned by the system through the DHCP snooping or relay process. Each subscriber SAP created on the system is configured (using the lease-populate command) to monitor DHCP activity between DHCP clients reached through the SAP and DHCP servers. DHCP ACKs from the DHCP server are used to determine that a certain IP address is in use by a specific DHCP client. This client IP address association is treated by the system as a dynamic subscriber host.
- When it is not possible to dynamically learn a subscriber host through DHCP, a static subscriber host can be created directly on a subscriber SAP. Since a subscriber identification policy is not applicable to static subscriber hosts, the subscriber identification string, subscriber profile and SLA profile must be explicitly defined with the hosts IP address and MAC address.

A subscriber (in the context of the router) is a collection of hosts getting common (overall) treatment. It is expected that this group of hosts originate from the same site and all hosts of a subscriber are reached by the same physical path (such as a DSL port).

Once a subscriber host is known by the system, it is associated with a subscriber identifier and an SLA profile instance. Subscriber hosts with a common subscriber identifier are considered to be owned by the same subscriber.

Depending on the network model, hosts associated with a single subscriber can be associated with a single subscriber SAP or spread across multiple subscriber SAPs on the same port.

## Subscriber Identification Policy

The subscriber identification policy contains the URL definitions for the Programmable Subscriber Configuration Policy (PSCP) scripts used for DHCP ACK message processing. Up to three URLs can be defined per subscriber identification policy. These are designated as primary, secondary and tertiary. Each URL can be individually enabled or disabled. Only one script (the URL with the highest priority active script) is used at any one time to process DHCP ACK messages. If the system detects an error with a specified script, the URL is placed in an operationally down state. If the script is shutdown, it is placed in an administratively down state. A script that is operationally or administratively down is considered inactive. The system automatically reverts to the highest priority active script. If a script becomes operationally down, it must be cycled through the administratively down then administratively up states for the system to attempt to reactivate the script.

Multiple subscriber identification policies are provided for the event that access nodes (such as DSLAMs) from different vendors are attached to the same router. Each policy's active script can be explicitly defined to process the various DHCP message formats or idiosyncrasies of each vendor.

If a script is changed, it must be reloaded by disabling and re-enabling any URL which refers to the changed script (a **shutdown** command followed by a **no shutdown** command).

Each subscriber identification policy can also contain a subscriber profile map and/or an SLA profile map. The subscriber profile map creates a mapping between the sub-profile-strings returned from the active script with an existing subscriber profile name. The SLA profile map is used to create a mapping between the sla-profile-strings returned from the active script with an existing SLA profile name.

The subscriber identification policy is designed to accept a DHCP ACK message destined for a subscriber host and return up to three string values to the system;

- The subscriber identification string (mandatory)
- The subscriber profile string (optional)
- The SLA profile string (optional).

These strings are used to derive the subscriber profile and the SLA profile to be used for this host. See [Using Scripts for Dynamic Recognition of Subscribers on page 808](#).

### Subscriber Identification String

Subscribers are managed by the router through the use of subscriber identification strings. A subscriber identification string uniquely identifies a subscriber.

The subscriber identification string is the index key to any entry in the active subscriber table, and thus must always be available. It is derived as follows:

- For dynamic hosts, the subscriber identification string is derived from the DHCP ACK message sent to the subscriber host.
  - The DHCP ACK message is processed by a subscriber identification script which has the capability to parse the message into an alternative ASCII string value.
  - If enhanced subscriber management is disabled, the default value for the string is the content of the Option 82 circuit-id and remote-id fields interpreted as an octet string.
- For static hosts, the subscriber identification string must be explicitly defined with each static subscriber host.

When multiple hosts are associated with the same subscriber identification string, they are considered to be host members of the same subscriber. Hosts from multiple SAPs can be members of the same subscriber, but for proper virtual scheduling to be performed all hosts of a subscriber must be active on the same IOM.

When the first host (either dynamic or static) is created with a certain subscriber identification string, an entry is created in the active subscriber table. The entries are grouped by their subscriber identification string.

---

### Subscriber Profile

The subscriber profile is a template which contains those hierarchical QoS (HQoS) and accounting settings which are applicable to all hosts belonging to the same subscriber. These include:

- Ingress and egress scheduler policy HQoS
- Accounting policy
- RADIUS accounting policy

Subscribers are either explicitly mapped to a subscriber profile template or are dynamically associated with a subscriber profile.

Attempting to delete any subscriber profile (including the profile named 'default') while in use by an existing active subscriber will fail.

## SLA Profile

For the purpose of supporting multiple service types (such as high speed Internet (HSI), voice over IP (VoIP), video on demand (VoD) and Broadcast TV) for a single subscriber, the hosts associated with a subscriber can be subdivided into multiple SLA profiles.

The SLA profile contains those QoS and security settings which are applicable to individual hosts. An SLA profile acts like a template and can be used by many subscribers at one time. Settings in the SLA profile include:

- Egress and ingress QoS settings
- Egress scheduler policy HQoS
- Egress and ingress IP filters
- Host limit

If the SLA profile does not explicitly define an ingress or egress QoS policy, the default SAP ingress or default SAP egress QoS policy is used.

Refer to [Determining the SLA Profile on page 812](#) for information on how the SLA profile is determined for dynamic hosts.

---

## Explicit Subscriber Profile Mapping

An explicit mapping of a subscriber identification string to a specific subscriber profile can be configured.

An explicit mapping overrides all default subscriber profile definitions while processing a DHCP ACK. In an environment where dynamic and static hosts coexist in the context of a single subscriber, care will be taken to not define a subscriber profile in the explicit subscriber map that conflicts with the subscriber profile provisioned for the static host(s). If such a conflict occurs, the DHCP ACKs will be dropped.

An explicit mapping of a subscriber identification string to the subscriber profile name 'default' is not allowed. However, it is possible for the subscriber identification string to be entered in the mapping table without a defined subscriber profile which can result in the explicitly defined subscriber to be associated with the subscriber profile named 'default'.

Attempting to delete a subscriber profile that is currently defined in an explicit subscriber identification string mapping will fail.

The explicit mapping entries can be removed at any time.

## ESM for IPv6

ESM for IPv6 is supported on 7750 chassis with at least IOM3-XP cards or equivalent or in 7450 chassis operating in Mixed Mode (containing one or more IOM3-XP cards that have the 7750 SR feature set enabled.) ESM for IPv6 is supported with RADIUS as the backend authentication and authorization mechanism.

### Models

- [PPPoE Host](#)
- [PPPoE RG](#)
- [IPoE Host/RG](#)

### PPPoE Host

For PPPoE, the ESR suggests the 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 indicates the setup of a PPPoE host is Framed-IPv6-Prefix, which should contain a /64 prefix for the client.

When a PPPoE host has successfully completed the IPv6CP negotiation, the ESR will transmit a Router Advertisement to the PPPoE host containing the suggested prefix and any other options that are configured. The client may use this information to pick one or more addresses from the suggested prefix; all addresses within the prefix are forwarded towards the client.

Alternatively, the Recursive DNS Server (RDNSS) Option as defined in RFC 6106, *IPv6 Router Advertisement Options for DNS Configuration*, can be included in IPv6 Router Advertisements for DNS name resolution of IPv6 SLAAC hosts. The following CLI command includes the DNS info in IPv6 Router Advertisements for SLAAC hosts and sets the RDNSS lifetime:

```
config>service>ies>sub-if>grp-if>ipv6>rtr-adv
config>service>vprn>sub-if>grp-if>ipv6>rtr-adv

[no] dns-options

      [no] include-dns      - Set/reset inclusion of the RDNSS server
                           option 25 on this group-interface
      [no] rdNSS-lifetime  - Maximum time the RDNSS address is valid
                           in this group-interface
```

The source for DNS information to be included in Router Advertisements for IPv6 SLAAC hosts, can be (listed in priority order):

1. Local User Database IPv6 options:



```
configure subscriber-mgmt local-user-db <ludb-name> dhcp|ppp host <host-name>
options6 dns-server <ip-address> [<ip-address>... (up to 4 max)]
```

2. RADIUS attributes [26-6527-105] Alc-Ipv6-Primary-Dns and [26-6527-106] Alc-Ipv6-Secondary-Dns
3. Default IPv6 DNS Server configured at the group interface:

```
configure service ies|vprn <svc-id> subscriber-interface <sub-int-name> ipv6 default-
dns <ipv6-address> [secondary <secondary-ipv6-address>]
```

**Note:** A default IPv6 server configuration at the group interface is a last resort IPv6 DNS info that can be used for IPoEv6 hosts (IA\_NA, IA\_PD and SLAAC) and PPPoEv6 hosts (IA\_NA, IA\_PD and SLAAC).

## PPPoE RG

Initially, a PPPoE RG follows the same procedure as a PPPoE host: the ESR receives a prefix from RADIUS (in this case through a Delegated-IPv6-Prefix attribute), which is used as a trigger to suggest the IPv6CP protocol to the client. The prefix that is suggested to the client should have the same prefix length as configured under the subscriber interface ipv6 node (delegated-prefix-length). This length should be between 48 and 64 bits, inclusive.

After the IPv6CP protocol has completed, however, the client should 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 downstream interfaces. During DHCPv6, no extra RADIUS request will be made; the information is stored during the initial (PPPoE or PPP) authentication until the client starts DHCPv6.

Only after DHCPv6 has completed, the IPv6 subscriber host will be instantiated and the ESR will start sending Router Advertisements (if configured.) The router advertisements will 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 ESR.

## IPoE Host/RG

Similar to an IPv4 DHCP client, a DHCPv6 client is authenticated at its Solicit message, where it can request one or more addresses or prefixes. The address and prefix types supported are IA\_NA (Non-Temporary Address) through the Alc-IPv6-Address RADIUS attribute and IA\_PD (Delegated Prefix) through the Delegated-IPv6-Prefix attribute. Contrary to the IPv4 case, the ESR will always reply to a DHCPv6 request because the client may request more than one address or prefix simultaneously and not all of the requests may be honored.

The DHCPv6 protocol handling and Router Advertisement behavior are similar to the PPPoE RG case above, with the exception that for an IA\_NA address, the entire /64 prefix containing the address is allocated to the client.

For SLAAC prefix assignment, authentication is triggered on router-solicit message. The SLAAC prefix can be assigned statically or dynamically. For a static SLAAC prefix, frame-ipv6-prefix, RADIUS attribute is used. For dynamic SLAAC prefix assignment from a local pool, Alc-slaac-ipv6-pool, RADIUS attribute is used.

## Setup

IPv6 ESM hosts are only supported in the Routed CO model (both VPRN and IES).

At the ipv6 node under the subscriber interface level, the length of the prefixes that are offered is defined through the delegated-prefix-length option. This setting is fixed for the subscriber interface and can not be changed once subscriber prefixes are defined.

Subscriber prefixes define the ranges of addresses that are offered on this subscriber interface. By default only these subscriber prefixes are exported to the routing protocols to keep the routing tables small. There are three types of subscriber interfaces:

- wan-host — A range of prefixes that are assigned to PPPoE hosts and as DHCPv6 IA\_NA addresses. These prefixes are always /64.
- pd — A range of prefixes that are assigned as DHCPv6 IA\_PD prefixes for DHCPv6 IPoE clients and for PPPoE RGs. The length of these prefixes is defined by the delegated-prefix-length.
- both — When both 'wan-host' and 'pd' are defined, the subscriber prefix is a range that can be used for both previous types. However, the delegated-prefix-length is restricted to /64 in this case.

The IPv6 node under the group interface contains the DHCPv6 proxy configuration and the router advertisement configuration.

## Behavior

- [Dual Stack](#)
  - [Router Advertisements \(RA\)](#)
  - [CoA and Disconnect-Request](#)
- 

### Dual Stack

Clients may support both IPv4 and IPv6 simultaneously (dual stack hosts.) In this case one subscriber host entry will be created for the IPv4 address family and one for the IPv6 instance. The scaling limits apply for all entries, regardless of address type.

For DHCP, these subscriber hosts are fully independent (as they are set up through different protocols), but for PPPoE hosts or RGs, the ESM information in both subscriber host entries is linked together through the PPPoE session.

---

### Router Advertisements (RA)

RA messages are started immediately after the subscriber host is instantiated and unsolicited messages are sent in the interval defined in the configuration. Apart from unsolicited RAs, the client may also send a router solicitation (RS) to explicitly request the information. RAs are throttled so that they are not sent more often than once every three seconds.

---

### CoA and Disconnect-Request

For IPv6 subscriber hosts, RADIUS-triggered mid-session changes and session terminations may identify the subscriber host to be changed by the same address or prefix that was originally returned from RADIUS. Only one address attribute (framed-IP-address, framed-IPv6-prefix, delegated-IPv6-prefix or Alc-IPv6-address) may be given in a single request.

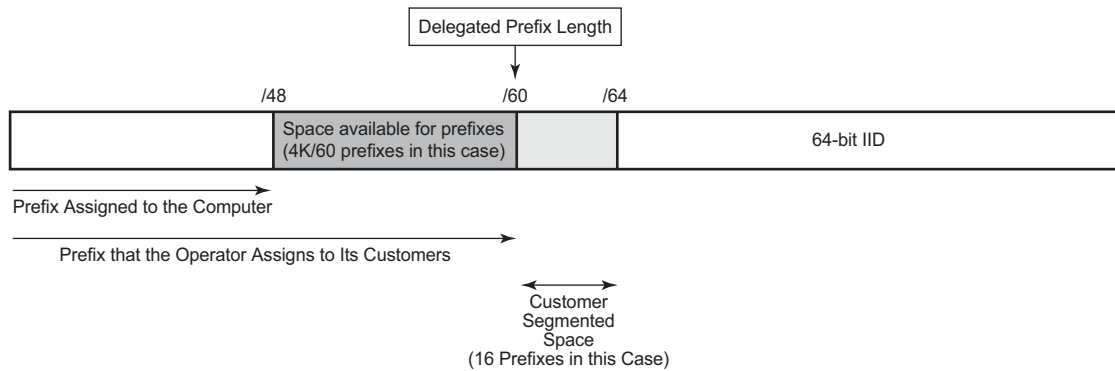
For PPPoE clients, changing either the IPv4 or IPv6 information will result in both the v4 and v6 subscriber host being modified (if they are contained within the same PPPoE session.)

The only CoA action that is allowed for IPv6 hosts is a change of ESM strings; creation of new hosts and forcing a DHCPv6 RENEW is not supported.

## Delegated-Prefix-Length

The delegated prefix length (DPL) is applicable to subscriber-hosts with IPv6 Prefix (IA-PD) assigned via DHCPv6 Server. IPv6 Prefix is more akin to a route than it is to an IP address. The length of the prefix plays crucial role in forwarding decisions, antispoofing, and prefix assignment via DHCPv6 pools in the local DHCPv6 Server.

The structure of an IPv6 prefix is shown in [Figure 44](#).



al\_0124

**Figure 44: IPv6 Prefix**

For example, a DHCPv6 server prefix pool contains an aggregated (configured) IPv6 prefix from which the delegated prefixes will be carved out. In [Figure 44](#) this aggregated IPv6 prefix has length of /48. In addition, the DHCPv6 server needs to know the length of the delegated prefix (in the above case /60). These two values are marking the boundary within which a unique delegated prefix will be selected. This is represented by the purple area in [Figure 44](#).

The delegated prefix length can be obtained via:

- RADIUS
  - Delegated-IPv6-Prefix attribute that contains the prefix and the length (Delegated-IPv6-Prefix = AAAA:BBBB::/56). The DPL in this case is /56.
  - Alc-Delegated-IPv6-Prefix-Length VSA (to be used in conjunction with the DHCPv6 pool name - Alc-Delegated-IPv6-Pool VSA)
- LUDB – configured via LUDB per IPoEv6/PPPoEv6 host:

```
configure
subscriber-mgmt
  local-user-db <name>
    dhcp | ppp
      host <name>
        ipv6-delegated-prefix-length [48..64]
```

This is to be used along with the DHCPv6 pool name (ipv6-delegated-prefix-pool) defined under the same CLI hierarchy.

Alternatively, the entire prefix, including the DPL can be returned via LUDB.

```
configure
  subscriber-mgmt
    local-user-db <name>
      dhcp | ppp
        host <name>
          ipv6-delegated-prefix <ipv6-prefix/prefix-length>
```

- DHCPv6 server – each DHCPv6 pool can optionally be configured with a DPL:

```
configure
  service/router
    dhcp6
      local-dhcp-server <name>
        pool <pool-name>
          delegated-prefix-length [48..64]
```

- Configured statically under the ipv6 CLI node of subscriber-interface. In this case the DPL is fixed for all subscriber-hosts under the subscriber-interface.

```
configure
  service ies/vprn
    subscriber-interface <name>
      ipv6
        delegated-prefix-length [48..64] | variable
```

## Order of Preference for DPL

In case that the DPL is statically provisioned under the subscriber-interface>ipv6> hierarchy, all hosts under this subscriber-interface will inherit this fixed DPL. In case that the DPL is provided via LUDB or RADIUS in addition to static configuration under the subscriber-interface then the LUDB or the RADIUS one MUST match the DPL that is statically provisioned under the subscriber-interface. Otherwise, the prefix instantiation in 7x50 will fail.

Note that the “no delegated-prefix-length” command under the **subscriber-interface>ipv6>** hierarchy means that the DPL is set to a default-value of 64.

When the delegated-prefix-length commands under the **subscriber-interface>ipv6>** hierarchy is set to variable, prefixes under such subscriber-interface can have different lengths and the DPL can be configured via one of the following means:

- LUDB
- RADIUS
- DHCP Server

### **DHCP Server Address Utilization and Delegated Prefix Length**

In case that the delegated prefix length is variable, for each consecutive address allocation request for the given delegated prefix, the DHCPv6 server will allocate the prefix at the end of the last delegated lease with the same delegated prefix length. This will minimize the address space fragmentation within the configured prefix.

## DHCPv6 Relay Agent

A DHCPv6 Relay Agent can support a 7x50 DHCPv6 local server (same or remote chassis) and a third party DHCPv6 external server.

An incoming DHCPv6 client message is relayed within the Relay-Forward message specified in RFC 3315. If the server responds with a valid address/prefix, the ESM process attempts to install it. If it fails, the DHCPv6 Relay Agent sends an explicit RELEASE to the server. There is no retransmission of DHCPv6 Relay-Forwards in the case of failure – it requires the client to re-start or re-send the original DHCPv6 message.

A Lightweight DHCPv6 Relay Agent may insert Relay Agent Information including the Interface ID option between the DHCPv6 client and the DHCPv6 Relay Agent.

Additional Relay Agents (non-LDRA) between the DHCPv6 client and the DHCPv6 Relay Agent are not supported.

## Configuring a DHCPv6 Relay Agent

A DHCPv6 Relay Agent is configured in the IPv6 DHCP6 context of a group-interface:

```
config>service>vprn>sub-if>grp-if>ipv6>dhcp6# relay ?
config>service>ies>sub-if>grp-if>ipv6>dhcp6# relay ?
  - no relay
  - relay

[no] client-applications - Configure the set of DHCP6 relay server client
                        applications
[no] description        - Description for DHCPv6 relay
[no] link-address       - Configure the link address of the DHCPv6 relay messages
[no] option             + Configure the DHCPv6 Relay information options
[no] server             - Configure the DHCPv6 server IPv6 address
[no] shutdown          - Administratively enable/disable DHCPv6 relay on this interface
[no] source-address     - Configure the source IPv6 address of the DHCPv6 relay messages
```

Up to eight DHCPv6 servers can be provisioned to be served by a DHCPv6 Relay Agent. A Relay-Forward is sent to all servers and the Relay-Replies from all servers are sent to the client.

The “client-applications” parameter specifies if the Relay Agent can be used for IPoE (dhcp) or PPP (ppp) hosts.

Optional configuration parameters:

- description — A free configurable description string.
- link-address — The link address field in the DHCPv6 Relay-Forward message header.

The link address can be configured to enable link-address based pool selection in a 7x50 DHCPv6 local server. The address must be one of the IPv6 prefixes configured at the ipv6 subscriber-prefixes context for a subscriber interface. If not configured, the system selects one of the prefixes.



- option: allows to configure following options to be inserted in the Relay-Forward message:
  - Interface-Id [18] — The interface ID option identifies the interface on which the DHCPv6 client message is received. The format options are the following:
    - ascii-tuple: — *host-name|service-id|group-interface-name|sap-id*
    - ifindex — Interface index for the group-interface
    - sap-id — SAP identifier (port and vlans)
    - string <string>: — A free configurable string (max. 80 chars)
  - Remote-Id [37] — Relay Agent Remote Id option contains the DHCPv6 client DHCP Unique Identifier (DUID).
- source-address: the source-address of the Relay-Forward messages.  
 If not configured, the outgoing interface IPv6 address is used. The source-address configuration is mandatory for a DHCP Relay Agent in a VPRN service when the DHCPv6 server is reachable via a tunnelled next-hop (MPLS).

---

## DHCPv6 Relay to Third Party DHCPv6 External Server

When the DHCPv6 Relay Agent is relaying to a third party DHCPv6 external server, following conditions should be met:

- The third party DHCPv6 server must return a unique IA\_PD IPv6 delegated prefix (/64 or lower) for each allocation. The length of the IA\_PD IPv6 delegated prefix must match the delegated-prefix-len configured on the subscriber interface on the 7750 DHCP L3 relay. This length is also included in the Relay-Forward message as PFX\_LEN option (3) in a Vendor-Specific-Information-Option (17).
- For IPv6oE routed CPE's, the 3rd party DHCPv6 server must return a unique IA\_NA IPv6 address (/128) from a different /64 subnet for each allocation.
- For IPv6oE hosts behind bridged CPE's,
  - the third party DHCPv6 server must return a unique IA\_NA IPv6 address (/128) from a different /64 subnet for each allocation (host) that belongs to a different CPE.
  - the third party DHCPv6 server may return a unique IA\_NA IPv6 address (/128) from the same /64 subnet for allocations (hosts) that belong to the same CPE and that are attached to the same vlan (SAP) on the BNG.

Following information is available to the third party DHCPv6 server in a Vendor-Specific-Information-Option (17) included in the Relay-Forward message:

- WAN\_POOL option (1) contains the pool name from which the IA\_NA IPv6 address should be allocated.
- PFX\_POOL option (2) contains the pool name from which the IA\_PD IPv6 delegated prefix should be allocated.

- PFX\_LEN option (3): contains the IA\_PD IPv6 delegated prefix length that should be allocated.

## DHCPv6 Local Server

A local DHCPv6 pool server for both addresses (IA\_NA) and prefixed (IA\_PD) manages the address and prefixes sent to either routing gateways or hosts.

Because IPv6 home networks lack NAT, the IPv6 addresses delegated to a routing gateway are in turn assigned to hosts in the home. These addresses are assigned with reasonably long (but configurable) lifetimes such that the loss of the WAN connection will not result in the IPv6 hosts in the LAN losing their IPv6 addresses. One consequence of these long lifetimes is that the IPv6 hosts will retain any IPv6 address provided the valid-lifetime is greater than zero. Should an operator delegate a prefix and then at a later time delegate a second IPv6 prefix, a host may end up with two or more valid prefixes. This situation plays havoc with IPv6 source address selection and may result in impaired service.

To overcome the problems of multiple IPv6 prefixes in the home, the operator must ensure that the individual subscriber has the same IPv6 prefix even across modem reboots (that is, if a subscriber session is destroyed and later re-created, an attempt should be made to use the previously delegated prefix). In release 8.0, the operator used RADIUS for all address and prefix assignment, but in release 9.0, with the introduction of the local DHCPv6 server, it requires the 7750 to process and maintain some state even after a session disconnects.

For the DHCPv6 local server to function, a DHCPv6 relay or proxy function must also operate alongside ESM. For the purposes of this document, to relay means to implement a DHCPv6 Relay as indicated in RFC 3315 : a relay encapsulates the client DHCP message within a DHCP Relay-Forward message and unicasts it to a specified destination.

A proxy is an internal concept. Unlike a DHCPv6 relay, the DHCPv6 proxy does NOT encapsulate the client message in a Relay-Forward, nor does it send packets towards the Local DHCPv6 Server. The DHCPv6 proxy is exclusively used as an interface between the RADIUS Access-Accept or local user database lookup and the DHCPv6 client in the consumer device.

The use of the DHCPv6 relay or proxy function depends on the attributes returned from authentication phase (RADIUS or LUDB).

1. DHCPv6 Proxy:

- If only IPv6 address/prefix information provided (Framed-IPv6-Prefix, Alc-IPv6-Address or Delegated-IPv6-Prefix).

2. DHCPv6 Relay:

- If no IPv6 address/prefix (Framed-IPv6-Prefix, Alc-IPv6-Address or Delegated-IPv6-Prefix) and no IPv6 pool (Framed-Pool, Delegated-Pool) information provided.

- If no IPv6 address/prefix (Framed-IPv6-Prefix, Alc-IPv6-Address or Delegated-IPv6-Prefix) and IPv6 pool (Framed-Pool, Delegated-Pool) information provided.

3. If both IPv6 address/prefix (Framed-IPv6-Prefix, Alc-IPv6-Address or Delegated-IPv6-Prefix) and IPv6 pool (Framed-Pool, Delegated-Pool) information are present, the DHCP packet is DROPPED.

**Note:** If IPv6 DNS parameters are returned in RADIUS AND a pool is specified then the DNS parameters are ignored. It is the DHCPv6 server that will need to reply with appropriate DNS servers.

## Dynamic Subscriber Host Processing

### Dynamic Tables

To support all processing for Enhanced Subscriber Management, several tables are maintained in the router (Figure 45).

- [Active Subscriber Table on page 794](#)
- [SLA Profile Instance Table on page 794](#)
- [Subscriber Host Table on page 794](#)
- [DHCP Lease State Table on page 796](#)

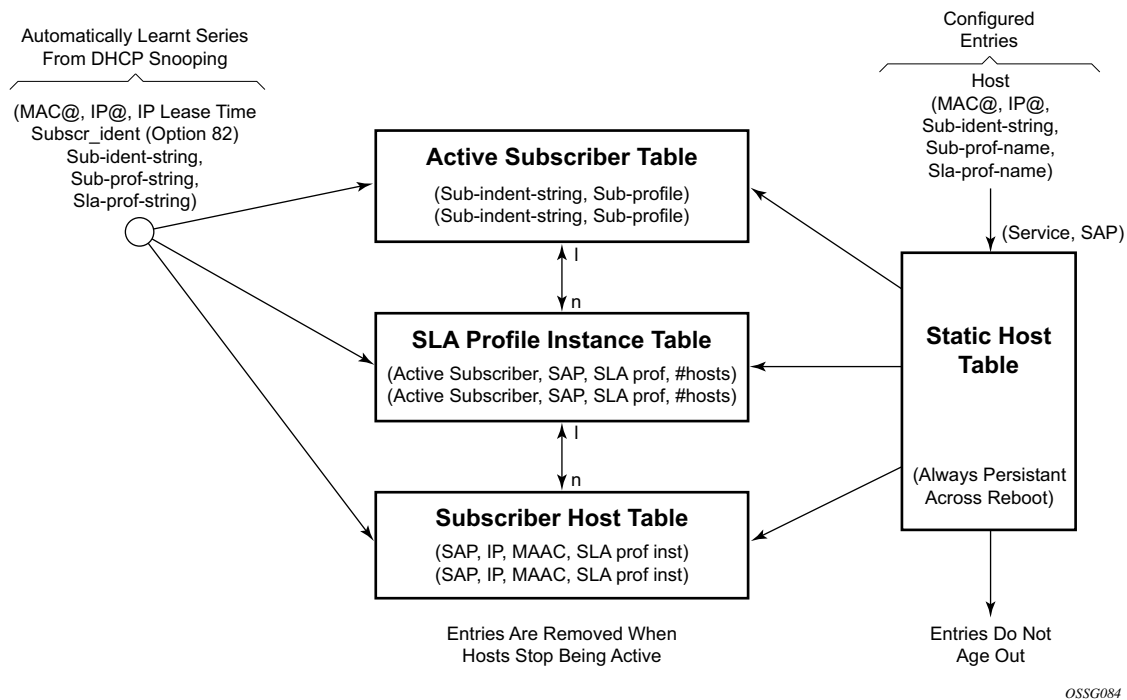


Figure 45: Enhanced Subscriber Management Dynamic Tables

### Active Subscriber Table

An entry is created in the active subscriber table when the first host (either dynamic or static) is created with a certain subscriber identification string. The entries are grouped by their subscriber identification string.

Fields for each entry in the active subscriber table include:

- The subscriber identification string (see [Subscriber Identification String on page 778](#)).
  - In use subscriber profiles (see [Subscriber Profile on page 778](#)).
- 

### SLA Profile Instance Table

An entry is created in the SLA profile instance table when the first subscriber host on a certain SAP is created that uses a certain SLA profile. All subsequent hosts of the same subscriber on the same SAP that use the same SLA profile will be associated with this entry. When the last host on this SAP, using this SLA profile disappears, the SLA profile instance is deleted from the table and the associated queues are removed.

SLA profile instances can not span multiple subscriber SAPs. If subscriber hosts from the same subscriber exist on multiple SAPs and are associated with the same SLA profile template, a separate SLA profile instance is created for each SAP.

Fields for each entry in the SLA profile instance table include:

- Active subscriber
  - SAP
  - SLA profile
  - Number of active subscriber hosts that share this instance
- 

### Subscriber Host Table

An entry is created in the subscriber host table if anti-spoofing is enabled as well as:

- The first host (dynamic or static) with a specific IP and MAC combination is created. If the anti-spoof is IP only, the MAC address is masked to all 0's. If anti-spoof is MAC, only the IP address is 0.0.0.0. All dynamic hosts and static hosts with the same IP and MAC combination will be associated with the same subscriber host entry. If the anti-spoof type includes IP (IP-only or IP/MAC), there can be at most two hosts associated with the entry: one dynamic and one static. If the anti-spoof type is MAC-only, there can be a combination of several dynamic and static hosts associated with the entry.
- The non-prof-traffic is provisioned. Both IP and MAC address are all 0's.

Fields for each entry in the subscriber host table include:

- SAP

- IP address
- MAC address
- SLA profile instance (enhanced mode only)

### **DHCP Lease State Table**

An entry in the DHCP lease state table is created for each dynamic host. Fields for each entry in the lease state table include:

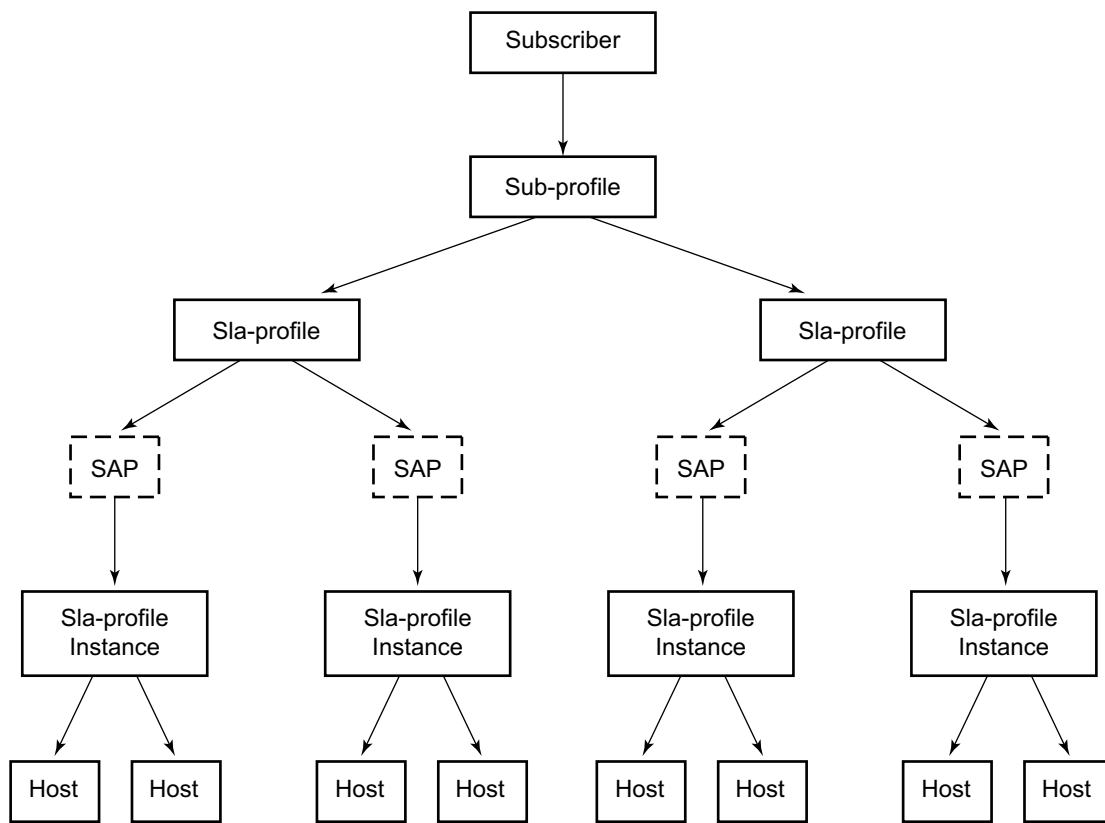
- Assigned IP address
- Assigned MAC address
- Persistence key



## Enhanced Subscriber Management Entities

Figure 46 illustrates the relationship between the main entities in Enhanced Subscriber Management:

- A subscriber is associated with only one subscriber profile.
- A subscriber can be associated with one or more SLA profile (a VPLS service with 2 different SAPs can have different SLA profiles for the same subscriber).
- A maximum of one SLA profile instance is generated (including ingress and egress queues) per SAP per SLA profile.
- One or more hosts can be assigned to each SLA profile instance (these will share the same queues).



OSSG085

**Figure 46: Relationship Between Enhanced Subscriber Management Entities**

## Instantiating a New Host

When a DHCP ACK is received for a new subscriber host on a particular SAP:

- The ACK message is parsed using the appropriate script.
- An entry is generated in the subscriber host table with indices:
  - The SAP on which the host resides
  - The assigned IP address
  - The assigned MAC address and as lookup parameters:
    - The subscriber profile and
    - The SLA profile to be used (derived from using the script).

If this is the first host of a subscriber, an HQoS scheduler is instantiated using the ingress and egress scheduler policies referred to in the subscriber profile. Otherwise, if the subscriber profile of the new host equals the subscriber profile of the existing subscriber, the new host is linked to the existing scheduler. If the subscriber profile is different from the subscriber profile of the existing subscriber, a new scheduler is created and all the hosts belonging to that subscriber are linked to this new scheduler. Notice that the new subscriber profile will not conflict with the subscriber profile provisioned for a static host or non-sub-traffic under the same SAP.

If this is the first host of a subscriber on a particular SAP using a particular SLA profile, an SLA profile instance is generated and added to the SLA profile instance table. This includes instantiating a number of queues, according to the ingress and egress QoS profiles referred to in SLA profile, optionally with some specific overrides defined in the SLA profile. Otherwise the host is linked to the existing SLA profile instance for this subscriber on this SAP.

Notes:

- Any QoS and IP filter policies defined on the SAP are still processed even if Enhanced Subscriber Management is enabled on the SAP. For IPv4 traffic that is dropped due to anti-spoofing, counters, logging, and mirroring can be used. All other Layer 2 traffic that is never blocked by anti-spoofing can be processed by applying a QoS policy on the SAP and can still be classified differently, by the dot1p value.
- If insufficient hardware resources (queues) or software resources (profile instances) are available to support the new host, the DHCP ACK is dropped and an event is generated.

## Packet Processing for an Existing Host

Whenever an IP packet arrives on a subscriber-facing SAP on which Enhanced Subscriber Management is enabled, a lookup is done in the subscriber host table using as the index the SAP, source IP address, and source MAC address.

- If there is no entry, this means that the host is not using his assigned IP address, so the packet is dropped;
- If there is an entry, this will refer to the subscriber profile and SLA profile to be used.

## ESM Host Lockout

This feature is applicable to the 7750 SR and the 7450 ESS.

This feature increasingly penalizes hosts that fail repeated login attempts within a configurable time interval. This is done by holding off on creation attempts for these hosts for a configured but adaptable time period. A transient failure, due to a mis-configuration, is quickly corrected and does not prevent the host from logging in within a reasonable amount of time. At the same time, a malicious client or a constantly mis-configured client is locked-out and will not take up resources impacting other clients.

A lockout time per host supports exponential back-off with each retry and failure cycle, starting with a configured minimum value and increasing up to a configured maximum. The lockout time can be reset to the configured minimum value if there is no failed retry within a configured time threshold. The configurable values include:

```
lockout-reset-time seconds
lockout-time [min seconds] [max seconds]
max-lockout-hosts hosts
```

If multiple retries/failure cycles occur within the lockout time, then lockout period is exponentially increased starting from configured minimum value up to the configured maximum value. The lockout is reset to the minimum value if there is no failed retry till this lockout time.

This mechanism is supported for both single and dual-stack PPPoE and IPoE (DHCP) hosts over 1:1 or N:1 static or managed SAPs. The hold-off timer maintenance is on a per host basis (as follows):

- For 1:1 VLAN (PPPoE or IPoE hosts) per <VLAN, MAC address>
- For N:1 VLAN (PPPoE or IPv4oE hosts) per <VLAN, agent-circuit-id, agent-remote-id, MAC@>
- For 1:1 VLAN (IPv6oE hosts) per <VLAN, DUID>

A show lockout state for hosts is supported, given one or more of <SAP, MAC@, agent-circuit-id, agent-remote-id>.

A clear lockout state is supported for hosts given one or more of <SAP, MAC@, agent-circuit-id, agent-remote-id>.

Any changes in configured lockout values will not apply to hosts currently under lockout and will only apply once these hosts are out of lockout.

## Functionality

ESM lockout is supported for dual-stack PPPoE hosts, dual-stack IPoE hosts and ARP hosts. ESM Lockout will track the following:

- PPPoE PADI and PADR
- DHCPv4 discover, DHCPv4 request, DHCPv6 solicit, DHCPv6 request
- ARP Request

During lockout, authentication and ESM host creation is suppressed. A lockout context will be created when a client first enters lockout. The context maintains state and timeout parameters for the lockout. If a lockout policy is configured for the underlying SAP for a host that has failed authentication or host creation, the host enters lockout for the configured minimum time (1 — 86400 seconds). When the lockout time expires, normal authentication and ESM host creation will be resumed on relevant PPP or DHCP messages. In case of another failure, the host will again enter the lockout state. The lockout time for the host on each failure will be exponentially increased up to the configured maximum time (1 — 86400 seconds). The lockout time for a client will be reset to the configured minimum value, and the corresponding lockout context will be deleted, if there is no authentication (and host creation) failure within a configured amount of time that needs to elapse after the client initially enters lockout. This time is called the **lockout-reset-time**.

The host identification for lockout includes <SAP, MAC@, circuit ID, remote ID>.

## ANCP and GSMP

- [ANCP on page 802](#)
  - [General Switch Management Protocol Version 3 \(GSMPv3\) on page 806](#)
- 

### ANCP

Access Node Control Protocol Management (ANCP) can provide the following information to the router:

- ANCP can communicate the current access line rate to the router. This allows the router to adjust the H-QoS subscriber scheduler with the correct rate or potentially change alarm when the rate goes below a set threshold. This allows a policy manager to change the entire policy when the rate drops below a minimal threshold value. The ANCP actual upstream synchronization rate is mapped to the ingress while ANCP actual downstream synchronization rate is mapped to the egress.
- The router can send DSL line OAM commands to complete an OAM test from a centralized point or when operational boundaries prevent direct access to the DSLAM.

When ANCP is used with Enhanced Subscriber Management (ESM), a new string ancp-string can be returned from the Python script or from RADIUS. If not returned it defaults to the subscriber ID.

ANCP version 0x31 and 0x32 are both supported and will be auto detected at the start of each ANCP session. Within version 0x32, partitioning is also supported.


Multiple partitions from the same Access Node are also supported. If partitions are used, they are automatically detected during the start of an ANCP session.

---

### Static ANCP Management

As depicted in [Figure 47](#), a DSLAM is connected to an aggregation network that is connecting the DSLAM to a BRAS. ANCP is used to provide SAP level rate management. The DSLAM in this application maintains multiple ANCP connections. The primary connection is to the BRAS, providing rate and OAM capabilities while the secondary is to the router to provide rate management.

7750 SR and 7450 ESS:




**Figure 47: Static ANCP Management Example**

## Enhanced Subscriber Management (ESM) Dynamic ANCP

In this application ANCP is used between the DSLAM and the router to provide line control. There are multiple attributes defined as described below. [Figure 48](#) depicts the connectivity model.

This application is used to communicate the following from the DSLAM to the router (the policy control point):

- Subscriber rate
- OAM



**Figure 48: ESM Dynamic ANCP Example**

---

## ANCP String

To support node communication with the access device the line rate, OAM commands, etc. the node can use an “ANCP string” that serves as a key in the out-of-band channel with the access node. The string can be either provisioned in the static case, retrieved from RADIUS or from the Python script.



## **ANCP Persistency Support**

Persistency is available for subscriber's ANCP attributes and is stored on the on-board compact flash card. ANCP data will stay persistence during an ISSU as well as nodal reboots. During recovery, ANCP attributes are first restored fully from the persistence file and incoming ANCP sessions are temporarily on hold. Afterwards new ANCP data can overwrite any existing values. This new data is then stored into the compact flash in preparation for the next event.

## General Switch Management Protocol Version 3 (GSMPv3)

General Switch Management Protocol version 3 (GSMPv3) is a generic protocol that allows a switch controller node to establish and maintain connections with one or more nodes to exchange operational information. Several extensions to GSMPv3 exist in the context of broadband aggregation. These extensions were proposed to allow GSMPv3 to be used in a broadband environment as additional information is needed to synchronize the control plane between access nodes (such as DSLAMs) and broadband network gateways (such as BRAS).

In the TPSDA framework, nodes fulfill some BRAS functionality, where per subscriber QoS enforcement is one of the most important aspects. To provide accurate per-subscriber QoS enforcement, the network element not only knows about the subscriber profile and its service level agreement but it is aware of the dynamic characteristics of the subscriber access circuit.

The most important parameters in this context are the subscriber-line capacity (DSL sync-rate) and the subscriber's channel viewership status (the actual number of BTV channels received by the given subscriber in any point in time). This information can be then used to adjust parameters of aggregate scheduling policy.

Besides, the above-mentioned information, GSMPv3 can convey OAM information between a switch controller and access switch. The node can operate in two roles:

- As the intermediate controller — The router terminates a connection from the DSLAM.
- As the terminating controller— The router fulfils full the roll of BRAS.

The DSL forum working documents recommends that a dedicated Layer 2 path (such as, a VLAN in an Ethernet aggregation network) is used for this communication to provide a certain level of security. The actual connection between DSLAM and BRAS is established at TCP level, and then individual messages are transported.

## DHCP Release Messages

The node supports DHCP release messages. A DHCP release message removes state from the DHCP server when the node rejects ACKs or removes hosts.

---

### DHCP Release

DHCP release messages will be controlled by the node and sent to the DHCP server to clear stale state. There are two examples:

1. If the node drops a DHCP ACK (because of resources, duplicate host or other reasons) the servers state must be cleared and the node will send a DHCP release.
  2. When a host state is removed, based on SHCV, ANCP, user clear, etc., the node will send a DHCP release to the server and the MAC will be flushed from SDPs. A new flag will allow the user to elect not to send the release message. If when using a clear lease command the host was removed by the user (using a clear command) a new flag will allow the user to elect not to send the release message.
- 

### DHCP Client Mobility

Client mobility allows the node to use host monitoring (SHCV, ANCP, split DHCP) to remove network and server state when a host is removed locally. This allows for MAC addressed learned and pinned to move based on policy parameters.

Subscriber Host Connectivity Verification (SHCV) configuration is mandatory. This allows clients to move from one SAP to another SAP in the same service. This is only applicable in a VPLS service and group interfaces.

The first DHCP message on the new SAP with same MAC address (and IP address for group-interfaces) will trigger SHCV and will always be discarded.

SHCV will check that the host is no longer present on the SAP where the lease is currently populated to prevent spoofing. When SHCV detects that the host is not present on the original SAP, the lease-state will be removed. The next DHCP message on the new SAP can initiate the host.

---

### DHCP Lease Control

DHCP lease control allows the node to be configured to present a different lease to the client. This can be used to monitor the health of the client.

## Using Scripts for Dynamic Recognition of Subscribers

Whenever a host belonging to a subscriber is activated (when a PC or set-top box (STB) is turned on), the host will typically request an IP address from the network using DHCP. Refer to [DHCP Management on page 313](#) for an explanation of DHCP and DHCP snooping in the router.

The DHCP ACK response from the DHCP server can be parsed and the contents of the message can be used to identify the “class” to which this host belongs, and thus, the QoS and security settings to apply.

The information necessary to select these settings can be codified in, the IP address by the DHCP server and/or the Option 82 string inserted by the DSLAM or other access node.

---

## Python Language and Programmable Subscriber Configuration Policy (PSCP)

PSCP is an identification mechanism using the Python scripting language. The PSCP references a Python script that can use regular expressions to derive the sub-ident-string, sub-profile-string and sla-profile-string from the DHCP response. A tutorial of regular expressions is beyond the scope of this guide, and can be found on the Internet (refer to <http://www.amk.ca/python/howto/regex/>).

A tutorial of Python is beyond the scope of this guide but can be found on the Internet (refer to <http://www.python.org/>).

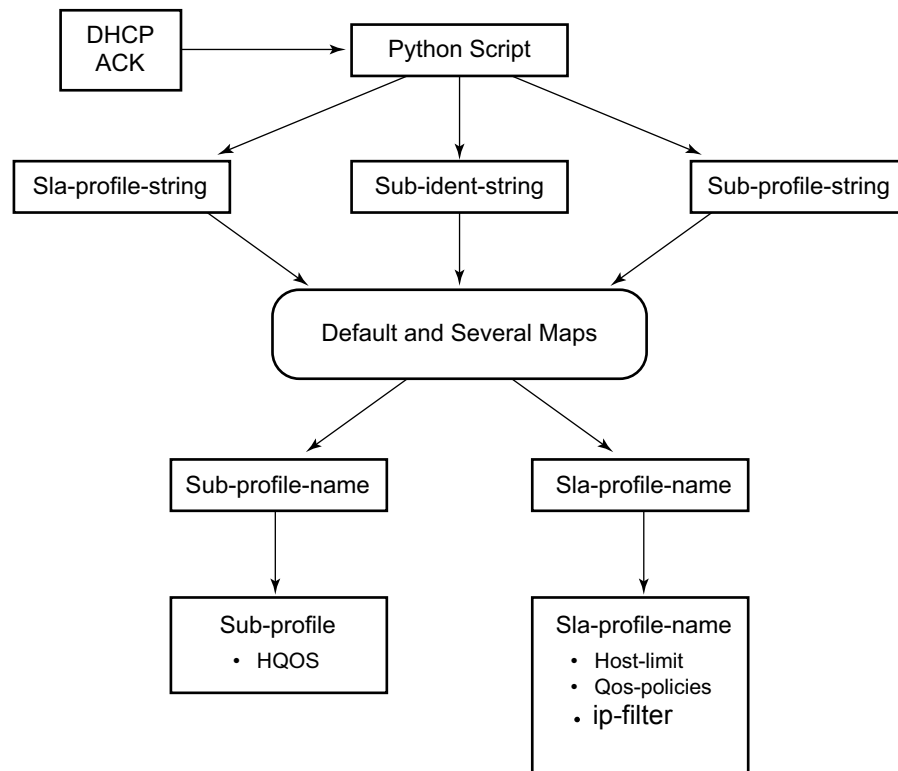
Example scripts, using some regular expressions, can be found in [Sample Python Scripts on page 1474](#). Additional information about the service manager scripting language, see [Python Script Support for ESM on page 1465](#).

One or more scripts can be written by the operator and stored centrally on a server (in a location accessible by the router). They are loaded into each router at bootup.

Note that if a centrally stored script is changed, it is not automatically re-loaded onto the router. The reload must be forced by executing the **shutdown / no shutdown** commands on the affected URL(s).

## Determining the Subscriber Profile and SLA Profile of a Host

Figure 49 describes the data flow while determining which subscriber profile and SLA profile to use for a certain subscriber host based on a snooped/relayed DHCP ACK for that subscriber host.



OSSG086

**Figure 49: Data Flow in Determining Subscriber Profile and SLA Profile**

An incoming DHCP ACK (relayed or snooped) is processed by the script provisioned in the sub-ident-policy defined in the SAP on which the message arrived. This script outputs one or more of the following strings:

- sub-ident — Identifies the subscriber (always needed).
- sub-profile — Identifies the subscriber class (optional).
- sla-profile — Identifies the SLA Profile for this subscriber host (optional).

These strings are used for a lookup in one or more maps to find the names of the sub-profile and sla-profile to use. If none of the maps contained an entry for these strings, the names will be determined based on a set of defaults.

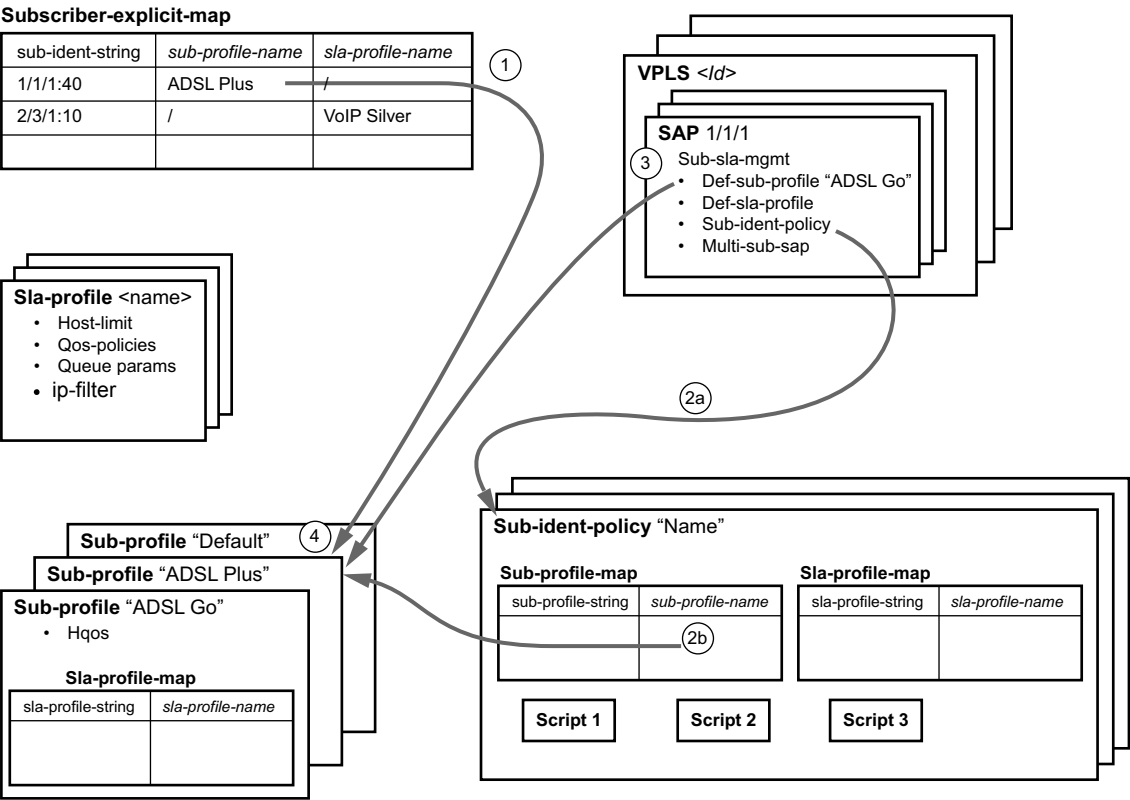
Only when the names for both the sub-profile and sla-profile are known, the subscriber host can be instantiated. If even no default is found for either profile, the DHCP ACK is dropped and the host will not gain network access.

Determining the Subscriber Profile

All hosts (devices) belonging to the same subscriber will be subject to the same HQoS processing. The HQoS processing is defined in the sub-profile. A sub-profile refers to an existing scheduler policy and offers the possibility to overrule the rate of individual schedulers within this policy.

Because all subscriber hosts of one subscriber use the same scheduler policy instance, they must all reside on the same I/O module.

The figure below shows how the sub-profile is derived, based on the sub-ident string, the sub-profile string and/or the provisioned data structures. The numbers associated with the arrows pointing toward the subscriber profiles indicate the precedence of the checks.



OSSG087

Figure 50: 7450 ESS Determining the Subscriber Profile

1. A lookup in the **explicit-subscriber-map** is done with the sub-ident string returned by the script. If a matching entry is found, the sub-profile-name (if defined) is taken. Otherwise:
2. 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, then:
3. If provisioned, the sub-profile-name is taken from the **def-sub-profile** attribute on the SAP. If not provisioned, then:
4. The **sub-profile** with the name “default” is selected (if provisioned). If this is not provisioned, there are no other alternatives, the ACK is dropped, and the host will not gain access.

## Determining the SLA Profile

For each host that comes on-line, the router also needs to determine which SLA profile to use. The SLA profile will determine for this host:

- The QoS-policies to use:
  - classification
  - queues/policers
  - queue mapping
- The egress scheduling policies to use:
  - egress HQoS
- The IP filter to use.

The SLA profile also has a host-limit attribute which limits the number of hosts (belonging to the same subscriber) on a certain SAP that can be using this SLA profile.

The classification and the queue mapping are shared by all the hosts on the same forwarding complex that use the same QoS policy (by their SLA profile).

The queues/policers are shared by all the hosts (of the same subscriber) on the same SAP that are using the same SLA profile. In other words, queues/policers are instantiated when, on a given SAP, a host of a subscriber is the first to use a certain SLA profile. This instantiation is referred to as an SLA profile instance. Ingress queues can be parented to a scheduler referenced in the ingress of a subscriber profile. Egress queues can be parented to a scheduler referenced in the egress of a subscriber or SLA profile, or to a port scheduler.

A scheduler policy can be applied to the egress an SLA profile, allowing its schedulers to be the parent for its queues and for its tier 1 schedulers to be parented to a scheduler in a scheduler policy applied to the egress of a subscriber profile or a vport, or to a port scheduler applied to a port or vport. Configuring scheduler overrides is allowed for SLA profile egress schedulers. The configuration of a scheduler policy in the egress of an SLA profile is supported for all host types only on Ethernet interfaces on FP2 and higher hardware. It is not supported for ESM over MPLS pseudowires, nor is HQoS adjustment and host tracking supported on its schedulers.

The following show/monitor/clear commands are available related to the SLA profile scheduler:

```
show qos scheduler-hierarchy subscriber sub-ident-string sla-profile sla-profile-name sap  
sap-id [scheduler scheduler-name] [detail]
```

The **show qos scheduler-hierarchy subscriber** command (shown above) displays the scheduler hierarchy with the SLA profile scheduler as the root. Note that if the SLA profile scheduler is orphaned (that is when the scheduler has a parent which does not exist) then the hierarchy is only shown when the **show** command includes the **sla-profile** and **sap** parameters.

**Note:** If the SLA profile scheduler is orphaned (that is when the scheduler has a parent which does not exist) then the hierarchy is only shown when the show command includes the sla-profile and SAP parameters.

```
monitor qos scheduler-stats subscriber sub-ident-string [interval seconds] [repeat repeat]
```

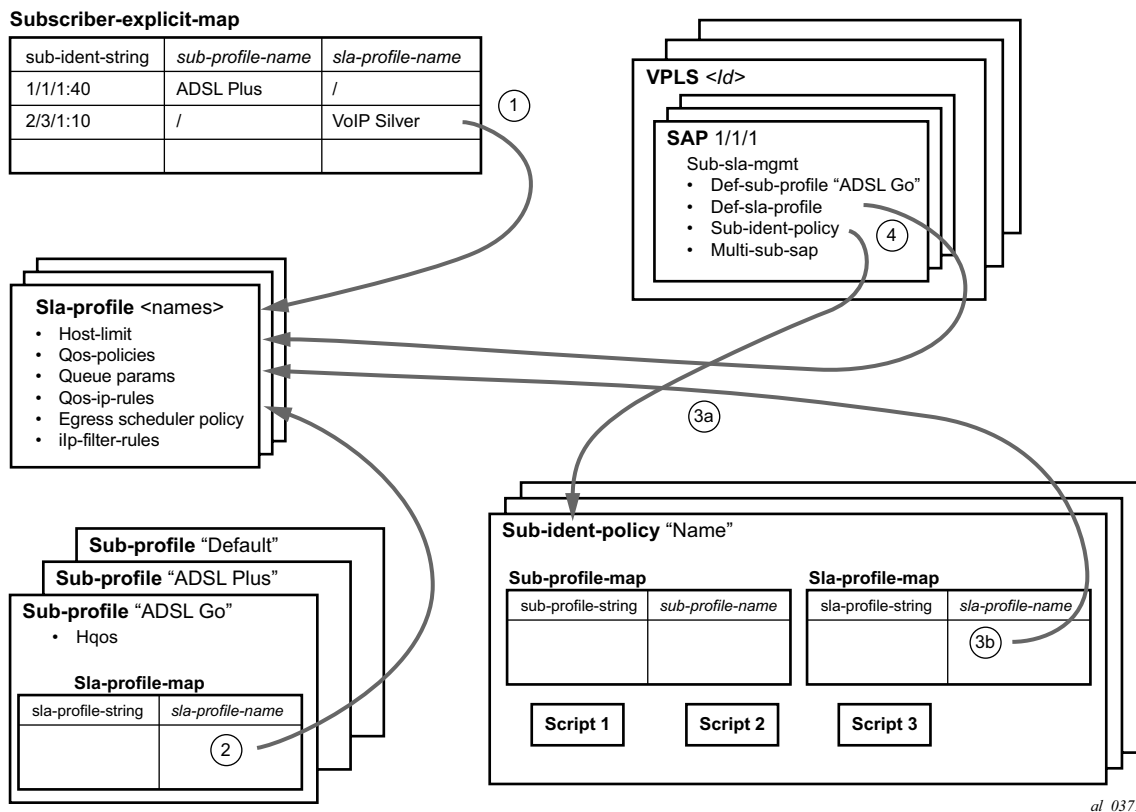


```
[absolute|rate] sap sap-id sla-profile sla-profile-name
```

```
show qos scheduler-stats subscriber sub-ident-string sap sap-id sla-profile sla-profile-name [scheduler scheduler-name]
```

```
clear qos scheduler-stats subscriber subscriber sub-ident-string sap sap-id sla-profile sla-profile-name [scheduler scheduler-name]
```

The figure below shows a graphical description of how the SLA profile is derived based on the subscriber identification string, the SLA profile string and the provisioned data structures. The numbers on the arrows towards the SLA profile indicate the *priority* of the provisioning (the lower number means the higher priority).



**Figure 51: 7450 ESS Determining the SLA Profile**

1. A lookup is done with the sub-ident string returned by the script in the **explicit-subscriber-map**. If a matching entry is found, the sla-profile-name is taken from it – if defined. Otherwise:

2. A lookup with the sla-profile string from the script is done in the **sla-profile-map** of the sub-profile found earlier. The sla-profile-name from the found entry is taken. If no entry was found, then:
3. A lookup is done with the sla-profile string in the **sla-profile-map** of the **sub-ident-policy** configured on the SAP. The sla-profile-name from the found entry is taken. If no **sub-ident-policy** was configured on the SAP or no entry was found, then:
4. If provisioned, the sla-profile-name is taken from the **def-sla-profile** attribute on the SAP. If not provisioned, there are no more alternatives, the ACK is dropped, and the host will not gain access.

## SLA-Based Egress QoS Marking

The egress QoS marking for subscriber-host traffic is derived from SAP-egress QoS policy associated with a corresponding SAP, rather than from the SLA profile associated with the corresponding subscriber-host. As a consequence, no egress QoS marking (or Dot1p marking is set to 0, the dscp/prec field is kept unchanged) is performed for traffic transmitted on a managed-SAP because per default, sap-egress policy 1 is attached to every managed-SAP.

The default value of the “qos-marking-from-sap” flag is enabled. This means that the qos-marking defined in the SAP egress QoS policy associated with the SAP will be used. The default setting of this flag in a combination with managed-SAP will result in the same behavior as in the current system (dot1p=0, dscp/prec is unchanged).

If “no qos-marking-from-sap” is executed, then both the Dot1p marking (all IOMs) and DSCP marking (IOM2/3 only) are derived from the sla-profile.

Changing the flag setting in the SLA profile being used by any subscriber-hosts (this includes subscriber-hosts on managed-SAPs as well) will be allowed.

The following MC traffic characteristics apply:

- On Layer 3 subscriber-interfaces, MC is not supported so it is impossible to enable it at the SAP level or at the sla-instance level.
- On Layer 2 SAPs IGMP snooping is supported while it is not supported on the sla-instance level. Therefore, any MC traffic transmitted at egress belongs to a SAP (meaning it will use SAP queues), rather than to sla-instance.
- The special case are SAPs with a profiled-traffic-only flag enabled. Although it is possible to define an sla-profile applicable to a Layer 2-host, this will not be taken as reference for marking mc-traffic, but rather SAP settings will be used.

## Auto-Sub ID

The subscriber ID name (sub-id) is a mandatory object that binds all hosts of a given subscriber together. Briefly, the sub-id name represents a residential household. Many management/troubleshooting and even billing operations rely on the sub-id name entity. The sub-id name is required for the host creation process, and it can be supplied by RADIUS or LUDB. It is derived from the sap-id or is statically provisioned in the form of a string.

In many ESM deployments with RADIUS, it is desirable that the sub-id is auto-generated within the 7x50 rather than burdening the OSS and the RADIUS server with this function. A typical application for auto sub-id is as follows:

- RADIUS server provides the sla-profile string and the sub-profile string but not the sub-id string.
- The sub-id name is auto-generated and formatted based on the configured options.

The following are the properties of auto sub-id generation:

- The auto-generation of the sub-id name can be based on any combination of the following fields:
  - MAC address
  - sap-id
  - circuit-id
  - remote-id
  - session-id

There can be only a single set of subscriber identification fields defined per host type (IPoE or PPPoE) per chassis. If the combination of the fields must be modified, the existing subscribers with an auto-generated sub-id must be manually terminated. Considering that remote termination of the IPoE subscribers by a DHCP server is not supported by all DHCP client vendors through the FORCERENEW DHCP message (RFC 3203, *DHCP reconfigure extension*), changing the subscriber fields while subscribers with auto generated sub-id are active should be avoided.

The sub-id name generation will take place at the end of the host initiation process (as after the authentication phase is completed) and only in case whereby the sub-id had not been already provided by any other more specific means (RADIUS, LUDB). This means that if the sub-id is supplied by other means (RADIUS, LUDB), then the sub-id name will not be auto-generated.

The format of the sub-id name can be either a random 10 characters encoded string or a user-friendly string based on the subscriber identification fields. Note that the maximum length of the sub-id name is 32 characters.

The sub-id name will not be passed in the Access-Request to the RADIUS server since it is generated after the authentication phase.

The sub-id name can be auto-generated regardless of how the sla/sub-profile strings are obtained (RADIUS, LUDB or static).

The subscriber identification fields used in auto-generation of the sub-id name are enabled on the global level.

**CLI Syntax:**

```
configure
subscriber-mgmt
  auto-sub-id-key
    ppp-sub-id-key [mac] [sap-id] [circuit-id] [remote-id] [session-id]
    ipoe-sub-id-key [mac] [sap-id] [circuit-id] [remote-id]
```

If no sub-id-key per host type is configured, then the defaults are:

PPPoE host type: <mac, sap-id, session-id>

IPoE host type : <mac, sap-id>.

The order in which the fields are configured is important because the sub-id name will potentially become a concatenated string of the subscriber host identifiers in the order in which they are provisioned. Note that the sub-id cannot be longer than 32 characters.

- In case that the length of the concatenated fields for the sub-id name is larger than 32 characters, the host creation will fail.
- In case that the circuit-id/remote-id is in the key and they contain non-printable characters, their place in sub-id name will be formatted in hex instead of ASCII. ASCII printable characters contain byte values 0x20..0x7E. All other values are ASCII non-printable and thus are formatted in hex characters.

The following would generate a sub-id name: xx:xx:xx:xx:xx:xx|1/1/3:23|44. The length of such sub-id name would be 29B.

- mac: xx:xx:xx:xx:xx:xx
- sap: 1/1/3:23
- session-id: 44 (16bits length)

In case that the key contains the circuit-id as: 0x610163 (3 bytes), then the sub-id name will be formatted as '610161' (hex) since '01' hex is non printable in ASCII. In this case the sub-id name will be of length 6B.

However, if the circuit-id is 0x616263 (3 bytes), then the string will be formatted as ASCII string 'abc' (3 characters). The sub-id name is 3B long.

The assignment of the sub-id to dynamic hosts is as follows:

- From RADIUS (sub-ident-policy including use-direct-map-as-default)
- From LUDB (sub-ident-policy including use-direct-map-as-default)

- Configured (explicit) defaults:
  - use-sap-id: sap-id
  - auto-id: combination of sub-id identifiers specified in auto-sub-id-key.  
The sub-id name will be in a human friendly format, i.e. concatenation of the fields in the pppoe|ipoe-sub-id-key command separated by a “|” character.
  - string: custom string
- Non-configured (implicit) defaults:

PPPoE host types: random 10 character string based on fields defined in the

- **ppp-sub-id-key** command. If no such fields are explicitly defined, the default ones will be assumed: <mac, sap-id, session-id>.
- IPoE host types: random 10 character string based on the **ipoe-sub-id-key** command. If no such fields are explicitly defined, the defaults will be assumed: <mac, sap-id >.

The way in which the default sub-id is generated is configured under the SAP level in the following manner:

**CLI Syntax:**

```
configure
service ies/vprn
    subscriber-interface <sub-if-name>
        group-interface <grp-if-name>
            sap <sap-id>
                sub-sla-mgmt
                    def-sub-id use-sap-id|use-auto-id|string
```

Under the msap-policy:

**CLI Syntax:**

```
configure
subscriber-mgmt
    msap-policy <name> (msap-policy referenced in msap-de-
        faults under the capture sap)
        sub-sla-mgmt
            def-sub-id <use-sap-id|use-auto-id|string <sub-
                id>
```

The **use-auto-id** keyword parameter of the def-sub-id string consists of concatenated auto-sub-id-keys separated by a ‘|’ character. In the absence of the **use-auto-id** keyword, the sub-id name will be a random 10 characters encoded string based on the ipoe|ppp-sub-id-keys. This random encoded 10 character string is unique per chassis as well as in dual-homed environment.

This command will have no effect if it is configured directly under the capture SAPs in VPLS (in the **config>service>vpls>sap>sub-sla-mgmt** context). Managed SAPs in ESM are instantiated by a capture SAP and the msap-policy in this case is mandatory. An auto-id keyword in case of managed SAP will be looked only under the msap-policy.

Static subscribers are required to have the sub-id manually configured.

---

## Sub-id Identifiers

The sub-id can be based on any combination of the following identifiers:

- The sap-id, in combination with any other allowable identifier, will be used as the search key. This assumes a 1:1 (subscriber per SAP) deployment model.
  - The circuit-id, in combination with any other allowable identifier, will be used to identify subscribers. This can be used in 1:1 deployment model, or in service per SAP deployment model. Circuit-id is applicable to IPoE v4 type hosts (option 82), to IPoE v6 type hosts (option 18 – interface-id) and PPPoE hosts (remote agent option signaled by PPPoE tags). The format of circuit-id is identical for IPv4 and IPv6 hosts.
  - The remote-id, in combination with any other allowable identifier, will be used to identify subscribers. This can be used in 1:1 deployment model, or in service per SAP deployment model. The remote-id is applicable to IPoE v4 type hosts (option 82), to IPoE v6 type hosts (option 37) and PPPoE hosts (remote agent option signaled via PPPoE tags).
  - The mac address (in combination with any other allowable identifier will be used to identify subscribers. This assumes a 1:1 deployment model.
  - The PPPoE session id, in combination with any other allowable identifier, is applicable only to PPPoE hosts. The session-id used will be of the first host that is instantiated for the subscriber.
- 

## Dual Stack Hosts

Autogeneration of sub-id names for subscribers with a single dual stack hosts (IPoE and PPPoE) is enabled by default by not explicitly provisioning anything for the def-sub-id. The sub-id name would be semi-randomly generated based on the <mac, sap-id, session-id> for PPPoE hosts and the <mac, sap-id> combination for IPoE host.

---

## Mixing Hosts with Auto-Generated IDs and non Auto-Generated IDs

Hosts with different sub-id names but identical auto-sub-id keys are not linked into the same subscriber. Such scenarios can arise with hosts with the same auto-sub-id keys but different methods for obtaining the sub-id name. For example, one host relying on auto-generated sub-id name while the other is using explicit configuration methods (sap-id, string, RADIUS or LUDb). If the auto-generated sub-id name and explicit sub-id name are the same, the host will be tied into the same subscriber.

For example:

The default auto-sub-id for the following two hosts are <mac, sap-id>.

Host X on SAP 1/1/1:1 with MAC 00:00:00:00:00:01 obtains sub-id through RADIUS.

Host Y on SAP 1/1/1:1 with MAC 00:00:00:00:00:01 has sub-id auto-generated.

Regardless of which host comes up first, those two hosts at the end will belong two different subscribers as long as their sub-ids are different.

---

## PPPoA/PPPoEoA Considerations

PPPoA/PPPoEoA hosts will adhere to the same rules as PPPoE. Fields that are supported in PPPoE but not PPPoA/PPPoEoA will be simply ignored.

Fields that are not supported in PPPoA are:

- Remote-id
- Circuit-id
- MAC address

Fields that are most likely not applicable to PPPoEoA are:

- Remote-id
  - Circuit-id
- 

## Deployment Considerations

The following is a possible deployment example scenario.

**CLI Syntax:**

```
configure
subscriber-mgmt
  auto-sub-id-key
  ppp-sub-id-key sap-id
  ipoe-sub-id-key mac circuit-id
```

**CLI Syntax:**

```
configure
service vprn 10
  subscriber-interface <sub-if-name>
  authentication-policy <auth-pol-name>
  group-interface <grp-if-name>
    sap 1
      sub-sla-mgmt
        def-sub-id use-sap-id
        sub-ident-policy <ident-pol-name>
    sap 2
      sub-sla-mgmt
```



```

        def-sub-id auto-id
        sub-ident-policy <ident-pol-name>
sap 3
    sub-sla-mgmt
        def-sub-id "sub3"
        sub-ident-policy <ident-pol-name>
sap 4
    sub-sla-mgmt
        sub-ident-policy <ident-pol-name>

```

Assume the following cases:

1. RADIUS returns the sub-id on all four SAPs.
2. RADIUS does not return the sub-id string on any of the SAPs.

In the first case where RADIUS returns the sub-id string, the following will occur:

- On all 4 SAPs the sub-id string will be assigned by the RADIUS server. Defaults have no effect, and neither do identifiers specified under the auto-sub-id-key node.

In the second case, the effects are the following:

- On SAP1 the sub-id name will be the <sap-id> (1/1/1:3)
- On SAP 2 the sub-id name will be <sap-id> for PPPoE hosts and <mac>-<circuit-id> concatenation for IPoE type hosts.

Example:

1/1/1:100 for PPPoE

AC:AB:AA:AD:AE:AE-AN-id eth 1/1/1:2 for IPoE

(circuit-ID format is: Access-Node-Identifier atm slot/port:vpi.vci or Access-Node-Identifier eth slot/port:[vlan-id]).

Note that the circuit-ID can itself be 63B in length whereas the length of the sub-id name is limited to 32 Bytes. So in the above case, the sub-id name length would be 38 Bytes (>32B) and the host instantiation would fail.

- On SAP3 the sub-id name will be the literal 'sub3' for PPPoE and IPoE hosts.
- On SAP4 the sub-id name will be a semi-random value based on <sap-id> for PPPoE hosts and the <mac, circuit-id> combination for IPoE hosts.

## **Caveats**

Only a single combination of the subscriber fields used to auto generate sub-id is allowed per host type (IPoE or PPPoE) and per chassis. In case that the combination of the fields needs to be changed, the existing subscribers with an auto-generated sub-id must be manually terminated. Considering that remote termination of the IPoE subscribers by DHCP server is not supported by all DHCP client vendors through FORCERENEW DHCP message (RFC 3203), changing the subscriber fields while subscribers with auto generated sub-id are active should be avoided.

## Limiting Subscribers and Hosts on a SAP

A number of configuration parameters are available to control the maximum amount of subscribers and/or hosts that can be simultaneously active on a SAP:

- multi-sub-sap — Limits the number of subscribers (dynamic + static) on a SAP
- lease-populate — Limits the number of dynamic hosts on a SAP
- host-limit — Limits the number of hosts (dynamic + static) per SLA profile instance.

If any of these limits are reached, a new host will be denied access and the DHCP ACK will be dropped. The only exception is when **host-limit** command is configured with the keyword **remove-oldest** specified, then the oldest active host is dropped and the new host is granted access. The dynamic host with the least remaining lease time will be considered the oldest host.

---

## Static Subscriber Hosts

While it is typically preferred to have all hosts provisioned dynamically through DHCP snooping, it may be needed to provide static access for specific hosts (those that do not support DHCP).

Since a subscriber identification policy is not applicable to static subscriber hosts, the subscriber identification string, subscriber profile and SLA profile must be explicitly defined with the host's IP address and MAC address (if Enhanced Subscriber Management is enabled).

If an SLA profile instance associated with the named SLA profile already exists on the SAP for the subscriber, the static subscriber host is placed into that SLA profile instance. If an SLA profile instance does not yet exist, one will be created if possible. If the SLA profile cannot be created, or the host cannot be placed in the existing SLA profile instance (the **host-limit** was exceeded), the static host definition will fail.

## QoS for Subscribers and Hosts

---

### QoS Parameters in Different Profiles

QoS aspects for subscribers and hosts can be defined statically on a SAP or dynamically using Enhanced Subscriber Management. For example, in a VLAN-per-service model, different services belonging to a single subscriber are split over different SAPs, and thus the overall QoS (such as a scheduler policy) of this subscriber must be assigned using Enhanced Subscriber Management.

QoS parameters are shared among the subscriber profile and SLA profile as follows:

- The subscriber profile refers to HQoS ingress and egress scheduler policies which define the overall treatment for hosts of this subscriber when queues are used. If the subscriber is using policers, the subscriber profile also refers to CFHP ingress and egress policer-control-policies which define the overall treatment for hosts of this subscriber.
- The SLA profile refers to specific queue/policer settings for each host (BTV, VoIP, PC) using SAP ingress and SAP egress QoS policies. The SLA profile can also refer to an egress HQoS scheduler policy which defines the scheduling from the queues of the related host.

The primary use of the subscriber profile is to define the ingress and egress scheduler policies/policer-control-policies used to govern the aggregate SLA for all hosts associated with a subscriber. To be effective, the queues/policers defined in the SLA profile's QoS policies will reference a scheduler/arbitrer from the scheduler policy/policer-control-policy respectively as their parent.

---

### QoS Policy Overrides

Generic QoS queue/policer parameters could be specified for the SAP in a QoS policy and overridden for some customers by queue/policer parameters defined in the SLA profile. This allows for a single SAP ingress and SAP egress QoS policy to be used for many subscribers, while providing individual subscriber parameters for queue/policer operation.

## ESM Subscriber Hierarchical Traffic Control

ESM subscribers can make use of both queues and/or policers for both the ingress and egress traffic. The queue and policers are configured within SAP ingress/egress policies applied to the SLA profile. The queue can parent to different levels/cir-levels with different weights/cir weights of a virtual scheduler configured within a scheduler policy, and to an egress port scheduler configured in a port scheduler policy, to achieve hierarchical traffic control. The policers can parent to different levels with different weights of an arbiter configured within a policer control policies to achieve hierarchical traffic control.

---

### Subscriber HQoS

Hierarchical QoS (HQoS) corresponds to scheduling bandwidth distribution to queues/schedulers and is applied using scheduler policies at ingress and egress of the subscriber profile for a subscriber, and at egress in the SLA profile for a host, together with a port scheduler at both the port and vport level.

Each scheduler policy can contain up to three tiers of schedulers with lower level schedulers being able to parent to higher level schedulers in the same scheduler policy.

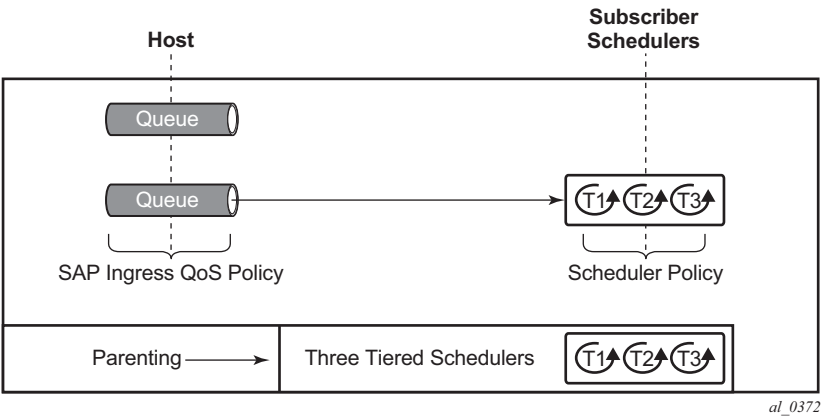
Queues can parent to any scheduler in their related scheduler policy hierarchy (except vport at egress) and also at the egress to a port scheduler.

Schedulers can parent to any higher level scheduler in their related scheduler policy hierarchy and, at the egress to a port scheduler configured within the port or vport. When an egress port scheduler is used, an aggregate rate limit can be applied at the subscriber profile and vport levels instead of using a scheduler. To extend the hierarchy further at egress, a tier 1 scheduler within a scheduler policy can parent to any scheduler in a scheduler policy at a higher level.

The scheduling levels are comprised of:

- Ingress and egress queues
- Egress SLA-profile schedulers
- Ingress and egress subscriber profile schedulers
- Egress vport schedulers
- Port schedulers

The ingress hierarchical parenting relationship options are shown in [Figure 52](#).



**Figure 52: Ingress Scheduling Hierarchy Options**

The egress hierarchical parenting relationship options are shown in [Figure 53](#). Note that not all combinations can be configured concurrently, and some uses of port parent could be equally achieved using a scheduler parent and a child parent-location.

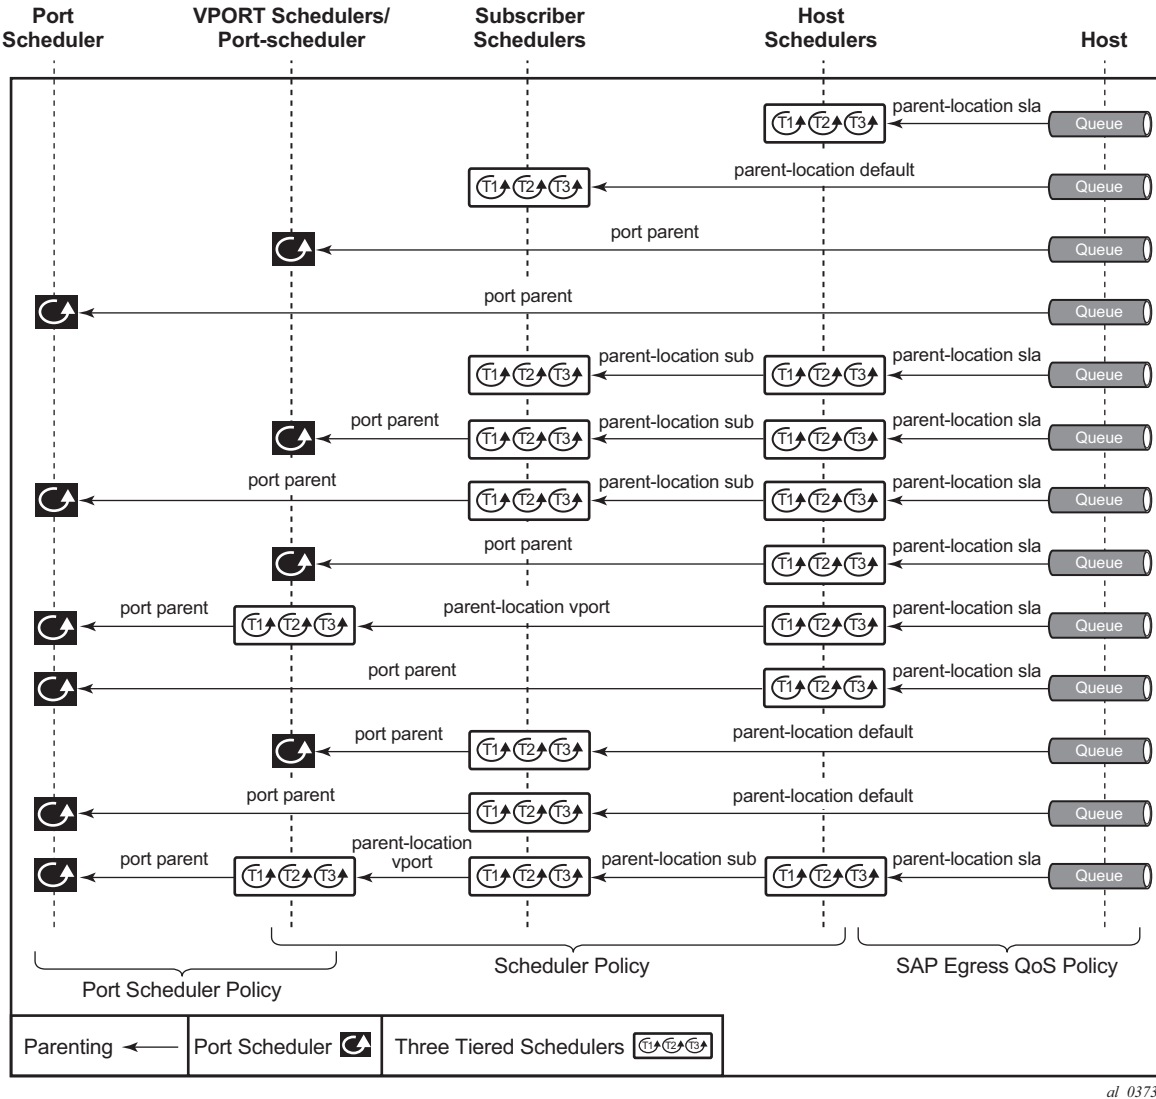


Figure 53: Egress Scheduling Hierarchy Options

The **parent** command is used to specify the name of the parent scheduler when parenting a queue or scheduler, together with the level/cir-level and weight/cir-weight at which to connect.

```
config>qos>sap-ingress>queue# parent
- parent scheduler-name [weight weight] [level level]
[cir-weight cir-weight] [cir-level cir-level]

config>qos>sap-egress>queue# parent
- parent scheduler-name [weight weight] [level level]
[cir-weight cir-weight] [cir-level cir-level]
```

```
config>qos>scheduler-policy>tier>scheduler# parent  
- parent scheduler-name [weight weight] [level level]  
[cir-weight cir-weight] [cir-level cir-level]
```

The location of the parent scheduler (in which applied scheduler policy it exists) for a queue defaults to a scheduler in the subscriber ingress or egress scheduler policy. Parents of schedulers themselves must be explicitly configured and by default must be within the same scheduler policy.

At egress, the scheduler parenting relationship is determined using the parent-location command:

- By default, egress queues parent to any scheduler in subscriber egress scheduler policy.

```
config>qos>sap-egress# parent-location default
```

- Egress queues can parent to any scheduler within the scheduler policy applied to the egress of an SLA profile.

```
config>qos>sap-egress# parent-location sla
```

- By default, a tier 1 scheduler in the scheduler policy is not allowed to be parented to another scheduler.

```
config>qos>scheduler-policy>tier# parent-location none
```

- A tier 1 scheduler in the scheduler policy applied to the egress of an SLA profile can parent to a scheduler applied to the egress of a subscriber profile.

```
config>qos>scheduler-policy>tier# parent-location sub
```

- A tier 1 scheduler in the scheduler policy applied to the egress of a subscriber profile can parent to a scheduler applied to the egress of a Vport.

```
config>qos>scheduler-policy>tier# parent-location vport
```

The configuration of a **parent-location** and frame-based accounting in a scheduler policy is mutually exclusive in order to ensure consistency between the different scheduling levels.

Note that the parent-location command is supported only on Ethernet interfaces on FP2 and higher hardware. It is not supported for ESM over MPLS pseudowires.

Both egress queues and egress schedulers can port parent using directly to different levels/cir-levels, with different weights/cir weights, to a port egress port scheduler. Egress schedulers can also port parent directly to different levels/cir-levels, with different weights/cir weights, to a vport egress port scheduler.



## Subscriber CFHP

Class Fair Hierarchical Policing (CFHP) corresponds to the policing control of traffic by policers/ arbiters. This uses policer control policies and can be applied for ingress and egress capacity control for the subscriber in the subscriber profile.

Each policer control policy can contain up to three tiers of arbiters with lower level arbiters being able to parent to higher level arbiters in the same scheduler policy.

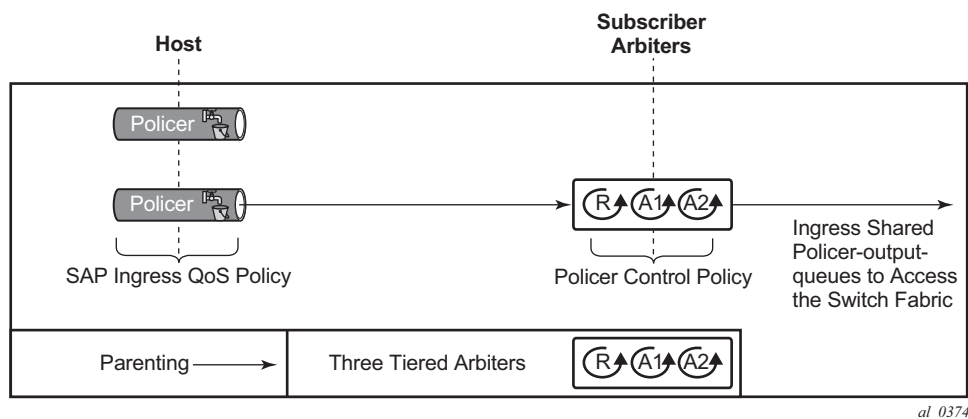
Policers can parent to any arbiter in their related policer control policy hierarchy.

The policing levels are comprised of:

- Ingress and egress policers
- Ingress and egress subscriber arbiters

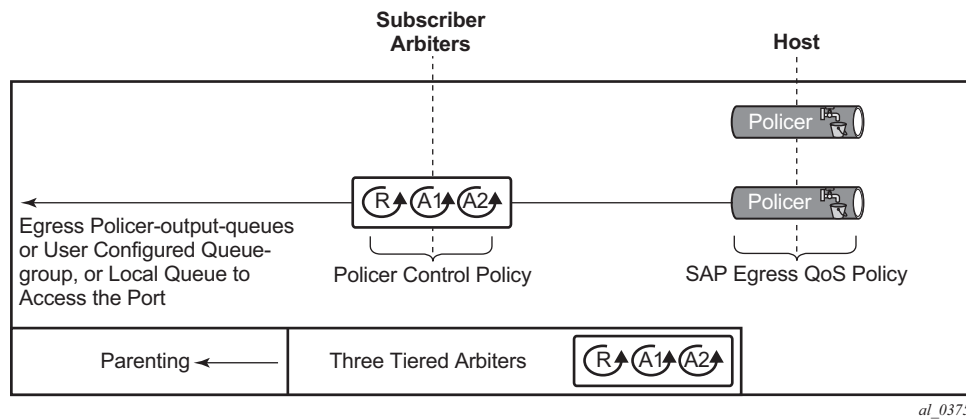
Note that ingress policed traffic uses the shared policer-output-queues to access the switch fabric. At egress, the policed traffic accesses the egress port through a queue group queue (by default the policer-output-queues queue group, though user configurable queue groups can also be used) or a locally configured subscriber queue.

The ingress hierarchical parenting relationship options are shown in [Figure 54](#).



**Figure 54: Ingress Policing Hierarchy Options**

The egress hierarchical parenting relationship options are shown in [Figure 55](#).



**Figure 55: Egress Policing Hierarchy Options**

The **parent** command is used to specify the name of the parent arbiter when parenting a policer or arbiter, together with the level and weight at which to connect.

```
config>qos>sap-ingress>policer$ parent
- parent arbiter-name [weight weight-level] [level level]

config>qos>sap-egress>policer$ parent
- parent arbiter-name [weight weight-level] [level level]

config>qos>plcr-ctrl-plcy>tier>arbiter# parent
- parent arbiter-name [weight weight-level] [level level]
```

## ATM/Ethernet Last-Mile Aware QoS for Broadband Network Gateway

This feature allows the user to perform hierarchical scheduling of subscriber host packets such that the packet encapsulation overhead and ATM bandwidth expansion (when applicable) due to the last mile for each type of broadband session, that is, PPPoEoA LLC/SNAP and VC-Mux, IPoE, IPoEoA LLC/SNAP and VC-Mux, etc., is accounted for by the 7x50 acting as the Broadband Network Gateway (BNG).

The intent is that the BNG distributes bandwidth among the subscriber host sessions fairly by accounting for the encapsulation overhead and bandwidth expansion of the last mile such that packets are less likely to be dropped downstream in the DSLAM DSL port.

The last mile encapsulation type can be configured by the user or signaled using the Access-loop-encapsulation sub-TLV in the Vendor-Specific PPPoE Tags or DHCP Relay Options as per RFC 4679.

Furthermore, this feature allows the BNG to shape the aggregate rate of each subscriber and the aggregate rate of all subscribers destined to a given DSLAM to prevent congestion of the DSLAM. The subscriber aggregate rate is adjusted for the last mile overhead. The shaping to the aggregate rate of all subscribers of a given destination DSLAM is achieved via a new scheduling object, referred to as Virtual Port or vport in CLI, which represents the DSLAM aggregation node in the BNG scheduling hierarchy

## Broadband Network Gateway Application

An application of this feature in a BNG is shown in [Figure 56](#).

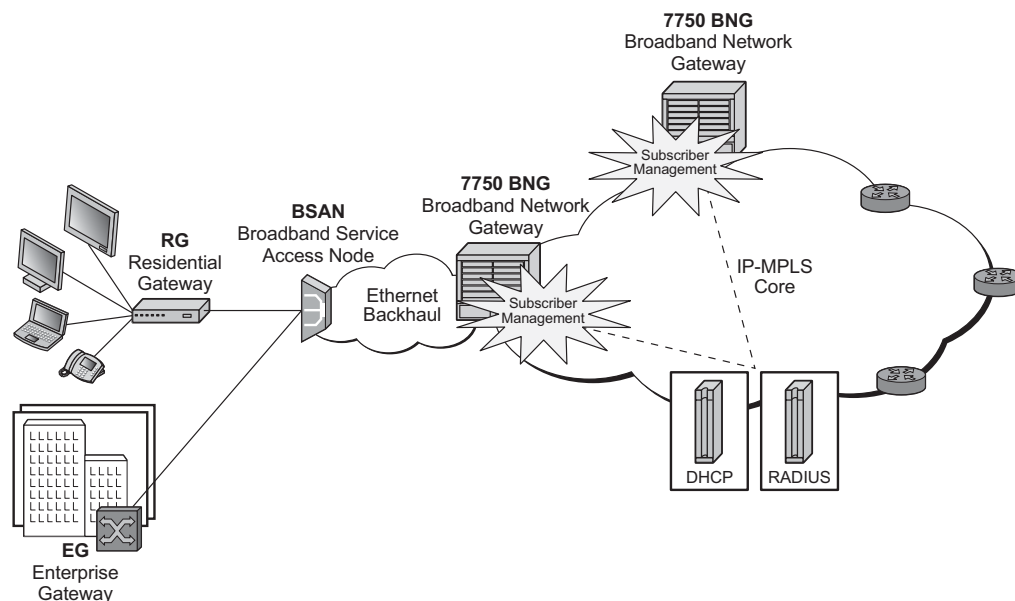


Figure 56: BNG Application

Residential and business subscribers use PPPoEoA, PPOA, IPoA, or IPoEoA based session over ATM/DSL lines. Each subscriber host can use a different type of session. Although Figure 1 illustrates ATM/DSL as the subscriber last mile, this feature supports both ATM and Ethernet in the last mile.

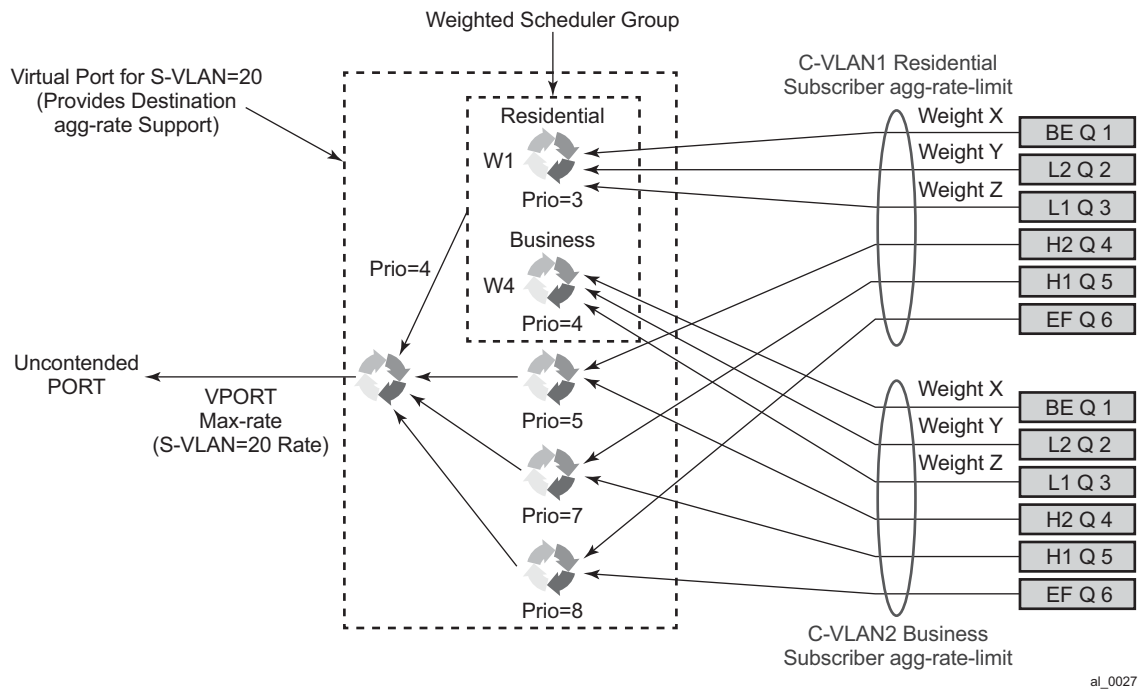
A subscriber SAP is auto-configured via DHCP or RADIUS authentication process, or is statically configured, and uses a Q-in-Q SAP with the inner C-VLAN identifying the subscriber while the outer S-VLAN identifies the Broadband Service Access Node (BSAN) which services the subscriber, i.e., the DSLAM. The SAP configuration is triggered by the first successfully validated subscriber host requesting a session. Within each subscriber SAP, there can be one or more hosts using any of the above session types. The subscriber SAP terminates on an IES or VPRN service on the BNG. It can also terminate on a VPLS instance.

When the 7750 BNG forwards IP packets from the IP-MPLS core network downstream towards the Residential Gateway (RG) or the Enterprise Gateway (EG), it adds the required PPP and Ethernet headers, including the SAP encapsulation with C-VLAN/S-VLAN. When the BSAN node receives the packet, it strips the S-VLAN tag, strips or overwrites the C-VLAN tag, and adds padding to minimum Ethernet size if required. It also adds the LLC/SNAP or VC-mux headers plus the fixed AAL5 trailer and variable AAL5 padding (to next multiple of 48 bytes) and then segments the resulting PDU into ATM cells when the last mile is ATM/DSL. Thus the packet size will undergo a fixed offset due to the encapsulation change and a variable expansion due to the AAL5 padding when applicable. Each type of subscriber host session will require a different amount of fixed offset and may require a per packet variable expansion depending of the encapsulation used by the session. The BNG node learns the encapsulation type of each subscriber host session by inspecting the Access-loop-encapsulation sub-TLV in the Vendor-Specific PPPoE Tags as specified in RFC 4679. The BNG node must account for this overhead when shaping packets destined to subscriber.

---

## Queue Determination and Scheduling

[Figure 57](#) illustrates the queuing and scheduling model for a BNG using the Ethernet/ATM last-mile aware QoS feature.



**Figure 57: BNG Queuing and Scheduling Model**

**CLI Syntax:** A set of per FC queues are applied to each subscriber host context to enforce the packet rate within each FC in the host session as specified in the subscriber's host SLA profile. A packet is stored in the queue corresponding the packet's FC as per the mapping of forwarding class to queue-id defined in the sap-egress QoS policy used by the host SLA profile. In the BNG application however, the host per FC queue packet rate is overridden by the rate provided in the RADIUS access-accept message. This rate represents the ATM rate that will be seen on the last mile, that is, it includes the encapsulation offset and the per packet expansion due to ATM segmentation into cells at the BSAN.

In order to enforce the aggregate rate of each destination BSAN, a scheduling node, referred to as virtual port, and vport is in the CLI. The vport operates exactly like a port scheduler with the difference that multiple vport objects can be configured on the egress context of an Ethernet port. The user adds a vport to an Ethernet port using the following command:

**CLI Syntax:** `configure>port>ethernet>access>egress>vport vport-name create`

The vport is always configured at the port level even when a port is a member of a LAG. The vport name is local to the port it is applied to but must be the same for all member ports of a LAG. It however does not need to be unique globally on a chassis.

**CLI Syntax:** `configure>port>ethernet>access>egress>vport vport-name create`

The vport is always configured at the port level even when a port is a member of a LAG. The vport name is local to the port it is applied to but must be the same for all member ports of a LAG. It however does not need to be unique globally on a chassis.

The user applies a port scheduler policy to a vport using the following command:

**CLI Syntax:** `configure>port>ethernet>access>egress>vport>port-scheduler-policy port-scheduler-policy-name`

A vport cannot be parented to the port scheduler when it is using a port scheduler policy itself. It is thus important the user ensures that the sum of the **max-rate** parameter value in the port scheduler policies of all vport instances on a given egress Ethernet port does not oversubscribe the port's hardware rate. If it does, the scheduling behavior degenerates to that of the H/W scheduler on that port. A vport which uses an agg-rate can be parented to a port scheduler. This is explained in [Applying Aggregate Rate Limit to a Vport on page 843](#). Note that the application of the agg-rate rate, port-scheduler-policy and scheduler-policy commands under a vport are mutually exclusive.

Each subscriber host queue is port parented to the vport which corresponds to the destination BSAN using the existing port-parent command:

**CLI Syntax:** `configure>qos>sap-egress>queue>port-parent [weight weight] [level level] [cir-weight cir-weight] [cir-level cir-level]`

This command can parent the queue to either a port or to a vport. These operations are mutually exclusive in CLI as explained above. When parenting to a vport, the parent vport for a subscriber host queue is not explicitly indicated in the above command. It is determined indirectly. The determination of the parent vport for a given subscriber host queue is described in [Vport Determination and Evaluation on page 843](#).

Furthermore, the weight (cir-weight) of a queue is normalized to the sum of the weights (cir-weights) of all active subscriber host queues port-parented at the same priority level of the vport or the port scheduler policy. Since packets of ESM subscriber host queues are sprayed among the link of a LAG port based on the subscriber-id, it is required that all subscribers host queues mapping to the same vport, such as having the same destination BSAN, be on the same LAG link so that the aggregate rate towards the BSAN is enforced. The only way of achieving this is to operate the LAG port in active/standby mode with a single active link and a single standby link.

The aggregate rate of each subscriber must also be enforced. The user achieves this by applying the existing agg-rate-limit command to the egress context of the subscriber profile:

**CLI Syntax:** `configure>subscriber-mgmt>sub-profile>egress>agg-rate-limit agg-rate`

In the BNG application however, this rate is overridden by the rate provided in the RADIUS access-accept message. This rate represents the ATM rate that will be seen on the last mile, that is, it includes the encapsulation offset and the per packet expansion due to ATM segmentation into cells at the BSAN.

## Weighted Scheduler Group

The existing port scheduler policy defines a set of eight priority levels with no ability of grouping levels within a single priority. In order to allow for the application of a scheduling weight to groups of subscriber host queues competing at the same priority level of the port scheduler policy applied to the vport, or to the Ethernet port, a new group object is defined under the port scheduler policy:

**CLI Syntax:** `configure>qos>port-scheduler-policy>group group-name rate pir-rate [cir cir-rate]`

Up to eight groups can be defined within each port scheduler policy. One or more levels can map to the same group. A group has a rate and optionally a cir-rate and inherits the highest scheduling priority of its member levels. For example, the scheduler group shown in the vport consists of level priority 3 and level priority 4. It thus inherits priority 4 when competing for bandwidth with the standalone priority levels 8, 7, and 5.

In essence, a group receives bandwidth from the port or from the vport and distributes it within the member levels of the group according to the weight of each level within the group. Each priority level will compete for bandwidth within the group based on its weight under congestion situation. If there is no congestion, a priority level can achieve up to its rate (cir-rate) worth of bandwidth.

The mapping of a level to a group is performed as follows:

**CLI Syntax:** `configure>qos>port-scheduler-policy>level priority-level rate pir-rate [cir cir-rate] group group-name [weight weight-in-group]`

Note that CLI will enforce that mapping of levels to a group are contiguous. In other words, a user would not be able to add priority level to group unless the resulting set of priority levels is contiguous.

When a level is not explicitly mapped to any group, it maps directly to the root of the port scheduler at its own priority like in existing behavior.

## Queue and Subscriber Aggregate Rate Configuration and Adjustment

---

### Software-Based Implementation (8.0R4)

The subscriber aggregate rate is adjusted and it will be based on an average frame size.

The user enables the use of this adjustment method by configuring the following option in the egress context of the subscriber profile:

**CLI Syntax:** `configure>subscriber-management>sub-profile>egress>encap-offset [type type]`

This command allows the user to configure a default value to be used by all hosts of the subscriber in the absence of a valid signaled value. The following is a list of the configurable values:

<b>Values</b>	pppoa-llc, pppoa-null, pppeoa-llc, pppeoa-llc-fcs, pppeoa-llc-tagged, pppeoa-llc-tagged-fcs, pppeoa-null, pppeoa-null-fcs, pppeoa-null-tagged, pppeoa-null-tagged-fcs ipoa-llc, ipoa-null, ipoeoa-llc, ipoeoa-llc-fcs, ipoeoa-llc-tagged, ipoeoa-llc-tagged-fcs, ipoeoa-null, ipoeoa-null-fcs, ipoeoa-null-tagged, ipoeoa-null-tagged-fcs, pppoe, pppoe-tagged, ipoe, ipoe-tagged
---------------	--

Otherwise, the fixed packet offset will be derived from the encapsulation type value signaled in the Access-loop-encapsulation sub-TLV in the Vendor-Specific PPPoE Tags as explained in Section Signaling of Last Mile Encapsulation Type. Only signaling using PPPoE Tags is supported in the software based implementation. The last signaled valid value is then applied to all active hosts of this subscriber. If no value is signaled in the subscriber host session or the value in the fields of the Access-loop-encapsulation sub-TLV are invalid, then the offset applied to the aggregate rate of this subscriber will use the last valid value signaled by a host of this subscriber if it exists, or the user entered default type value if configured, or no offset is applied.

The user also configures the average frame size value to be used for this adjustment:

**CLI Syntax:** `configure>subscriber-management>sub-profile>egress>avg-frame-size bytes`

The value entered by the user must include the FCS but not the Inter-Frame Gap (IFG) or the preamble. If the user does not explicitly configure a value for the **avg-frame-size** parameter, then it will also be assumed the offset is zero regardless of the signaled or user-configured value.

The computation of the subscriber aggregate rate consists of taking the average frame size, adding the encapsulation fixed offset including the AAL5 trailer, and then adding the variable offset consisting of the AAL5 padding to next multiple of 48 bytes. The AverageFrameExpansionRatio is then derived as follows:

$$\text{AverageFrameExpansionRatio} = (53/48 \times (\text{AverageFrameSize} + \text{FixedEncapOffset} + \text{AAL5Padding})) / (\text{AverageFrameSize} + \text{IFG} + \text{Preamble}).$$

When the last mile is Ethernet, the formula simplifies to:



$$\text{AverageFrameExpansionRatio} = (\text{AverageFrameSize} + \text{FixedEncapOffset} + \text{IFG} + \text{Preamble}) / (\text{AverageFrameSize} + \text{IFG} + \text{Preamble}).$$

The following are the frame size and rate applied to the subscriber queue and scheduler:

Subscriber Host Queue (no change):

$$\text{Size} = \text{ImmediateEgressEncap} + \text{Data}$$

$$\text{Rate} = \text{ImmediateEgressEncap} + \text{Data}$$

Subscriber Aggregate Rate Scheduler:

$$\text{Size} = \text{ImmediateEgressEncap} + \text{Data}$$

$$\text{Rate} = \text{sub-agg-rate} / \text{AverageFrameExpansionRatio}$$

Note that the CPM applies the *AverageFrameExpansionRatio* adjustment to the various components used in the determination of the net subscriber operational aggregate rate. It then pushes these adjusted components to IOM which then makes the calculation of the net subscriber operational aggregate rate.

The formula used by the IOM for this determination is:

$$\text{sub-oper-agg-rate} = \min(\text{sub-policy-agg-rate} / \text{AverageFrameExpansionRatio}, \text{anep\_rate} / \text{AverageFrameExpansionRatio}) + (\text{igmp\_rate\_delta} / \text{AverageFrameExpansionRatio}),$$

where *sub-policy-agg-rate* is either the value configured in the **agg-rate-limit** parameter in the subscriber profile or the resulting RADIUS override value. In both cases, the CPM uses an internal override to download the adjusted value to IOM.

The value of *sub-oper-agg-rate* is stored in the IOM's subscriber table.

The following are the procedures for handling signaling changes or configuration changes affecting the subscriber profile:

1. If a new RADIUS update comes in for the aggregate subscriber rate, then a new subscriber aggregate ATM adjusted rate is computed by CPM using the last configured **avg-frame-size** and then programmed to IOM.
2. If the user changes the value of the **avg-frame-size** parameter, enables/disables the **encap-offset** option, or changes the parameter value of the **encap-offset** option, the CPM will immediately trigger a re-evaluation of subscribers using the corresponding subscriber profile and an update the IOM with the new subscriber aggregate rate.
3. If the user changes the value of the **agg-rate-limit** parameter in a subscriber profile which has the **avg-frame-size** configured, this will immediately trigger a re-evaluation of subscribers using the corresponding subscriber profile. An update to the subscriber aggregate rate is performed for those subscribers which rate has not been previously overridden by RADIUS.
4. If the user changes the **type** value of the encap-offset command, this will immediately trigger a re-evaluation of subscribers using the corresponding subscriber profile. An update to the subscriber aggregate rate is performed for those subscribers which are currently using the default value.

5. If two hosts of the same subscriber signal two different encapsulation types, the last one signaled gets used at the next opportunity to re-evaluate the subscriber profile.
6. If a subscriber has a DHCP host, a static host or an ARP host, the subscriber aggregate rate will continue to use the user-configured default encapsulation type value or the last valid encapsulation value signaled in the PPPoE tags by other hosts of the same subscriber. If none was signaled or configured, then no rate adjustment is applied.

## Hardware-Based Implementation

The data path will compute the adjusted frame size real-time for each serviced packet from a queue by adding the actual packet size to the fixed offset provided by CPM for this queue and variable AAL5 padding.

Like in the software based implementation, the user enables the use of the fixed offset and per packet variable expansion by configuring the following option in the egress context of the subscriber profile:

**CLI Syntax:** `configure>subscriber-management>sub-profile>egress>encap-offset [type type]`

When this command is enabled, the fixed packet offset will be derived from the encapsulation type value signaled in the Access-loop-encapsulation sub-TLV in the Vendor-Specific PPPoE Tags or DHCP Relay Options as explained in Section Signaling of Last Mile Encapsulation Type.

If the user specifies an encapsulation type with the command, this value will be used as the default value for all hosts of this subscriber until a host session signaled a valid value. The signaled value is applied to this host only and the remaining hosts of this subscriber continue to use the user entered default type value if configured, or no offset is applied. Note however that hosts of the same subscriber using the same SLA profile and which are on the same SAP will share the same instance of FC queues. In this case, the last valid encapsulation value signaled by a host of that same instance of the SAP egress QoS policy will override any previous signaled or configured value.

The procedures for handling signaling changes or configuration changes affecting the subscriber profile are the same as in the Software based implementation with except for the following:

1. The **avg-frame-size** parameter in the subscriber profile is ignored.
2. If the user specifies an encapsulation type with the command, this value will be used as the default value for all hosts of this subscriber until a host session signaled a valid value. The signaled value is applied to this host and other hosts of the same subscriber sharing the same SLA profile and which are on the same SAP. The remaining hosts of this subscriber continue to use the user entered default type value if configured, or no offset is applied.
3. If the user enables/disables the **encap-offset** option, or changes the parameter value of the **encap-offset** option, the CPM will immediately trigger a re-evaluation of subscribers hosts using the corresponding subscriber profile and an update the IOM with the new fixed offset value.
4. If subscriber host session signals an encapsulation type at the session establishment time and subsequently sends a DHCP renewal message using a Layer 2 DHCP relay which does not insert option82 in a unicast message, the encapsulation type for this host will not change. Note that TR-101 states that option82 is mandatory for DHCP broadcast messages).
5. If a subscriber has a static host or an ARP host, the subscriber host will continue to use the user-configured default encapsulation type value or the last valid encapsulation value signaled in the PPPoE tags or DHCP relay options by other hosts of the same subscriber which

use the same SLA profile instance. If none was signaled or configured, then no rate adjustment is applied.

6. The encapsulation type value signaled in DHCP relay options or PPPoE tags are not cross-checked against the host type. So, a host signaling PPPoA/LLC encapsulation type via DHCP relay options will be handled as if the packet included a PPPoE header when forwarded over the local Ethernet port. This results in applying an encap-offset in the data path which assumes the PPPoE header is added to forwarded packets over the local Ethernet port.

The **encap-offset** option forces all the rates to be either last-mile frame over the wire or local port frame over the wire, referred to as **LM-FoW** and **FoW** respectively. The system maintains a running average frame expansion ratio for each queue to convert queue rates between these two formats as explained in [Frame Size, Rates, and Running Average Frame Expansion Ratio on page 842](#). Here are the details of the queue and scheduler operation:

1. When the **encap-offset** option is configured in the subscriber profile, the subscriber host queue rates, that is, CLI and operational PIR and CIR as well as queue bucket updates, the queue statistics, that is, forwarded, dropped, and HQoS offered counters use the **LM-FoW** format. The scheduler policy CLI and operational rates also use **LM-FoW** format. The port scheduler **max-rate** and the priority level rates and weights, if a Weighted Scheduler Group is used, are always entered in CLI and interpreted as **FoW** rates. The same is true for an **agg-rate-limit** applied to a vport. Finally the subscriber **agg-rate-limit** is entered in CLI as **LM-FoW** rate. When converting between **LM-FoW** and **FoW** rates, the queue running average frame expansion ratio value is used.
  - If the user enabled **frame-based-accounting** in a scheduler policy or **queue-frame-based-accounting** with subscriber **agg-rate-limit** and a port scheduler policy, the queue operational rate will be capped to a user configured **FoW** rate. The scheduler policy operational rates will also be in the **FoW** format. Note that a user configured queue **avg-frame-overhead** value is ignored since the running average frame expansion ratio is what is used when the **encap-offset** option is enabled.
  - If the user configured queue **packet-byte-offset** value, it is ignored and is not accounted for in the net packet offset calculation.
2. When **no encap-offset** is configured in the subscriber profile, that is, default and pre-R9.0 behavior, queue CLI and operational PIR and CIR rates, as well as queue bucket updates, the queue statistics, use data format. The scheduler policy CLI and operational rates also use data format. The port scheduler **max-rate** and the priority level rates and weights, if a Weighted Scheduler Group is used, and the subscriber **agg-rate-limit** are entered in CLI and interpreted as **FoW** rates. When converting between **FoW** and data rates, the queue **avg-frame-overhead** value is used and since this an Ethernet port, it is not user-configurable but constant and is equal to +20 bytes (IFG and preamble).
  - If the user enabled **frame-based-accounting** in a scheduler policy or **queue-frame-based-accounting** with subscriber **agg-rate-limit** and a port scheduler policy, the queue operational rate will be capped to a user configured **FoW** rate in

CLI which is then converted into a data rate using the queue **avg-frame-overhead** constant value of +20 bytes. The scheduler policy operational rates will also be in the **FoW** format.

- If the user configured queue **packet-byte-offset** value, it adjusts the immediate packet size. This means that the queue rates, that is, operational PIR and CIR, and queue bucket updates use the adjusted packet size. In addition, the queue statistics will also reflect the adjusted packet size. Scheduler policy rates, which are data rates, will use the adjusted packet size. The port scheduler **max-rate** and the priority level rates and weights, if a Weighted Scheduler Group is used, as well as the subscriber **agg-rate-limit** are always **FoW** rates and thus use the actual frame size

## Frame Size, Rates, and Running Average Frame Expansion Ratio

The following are the details of the rates and frame sizes applied to the subscriber host queues, the subscriber aggregate rate, and the vport root scheduler for the scheduling model and when the **encap-offset** option is enabled in the subscriber profile.

Subscriber Host Queue:

Size = LastMileFrameOverWireEncap + Data

Rate = (48/53)\* x (LastMileFrameOverWireEncap + Data)

\*Applicable to ATM last-mile only.

Subscriber Aggregate Rate:

Size = LastMileFrameOverWireEncap + Data

Rate = (48/53)\* x (LastMileFrameOverWireEncap + Data)

\*Applicable to ATM last-mile only.

Vport/Port Port Scheduler and Weighted Scheduler Group

Size = FrameOverWireEncap + Data

Rate = FrameOverWireEncap + Data

When a frame arrives at the queue, its size will be *ImmediateEgressEncap+Data*. This size is stored as the *OfferedFrameSize* so that the queue offered stats used in HQoS calculations are correct. Let us refer to the HQoS offered statistics as Offered.

This size is then adjusted by removing the *ImmediateEgressEncap* and adding the *LastMileFrameOverWireEncap*. This new adjusted frame size, let us refer to it as *LastMileOfferedFrameSize*, is then used for checking compliance of the frame against the queue PIR and CIR bucket sizes and for updating the queue forwarded and dropped stats.

The *LastMileOfferedFrameSize* value is computed dynamically for each packet serviced by the queue.

A new HQoS stat counter *OfferedLastMileAdjusted* is maintained for the purpose of calculating the running average frame expansion ratio, which is the ratio of the accumulated *OfferedLastMileAdjusted* and Offered of each queue:

$$\text{RunningAverageFrameExpansionRatio} = \text{OfferedLastMileAdjusted} / \text{Offered}$$

The vport/port port scheduler will hand out its **FoW** bandwidth in terms of Fair Information Rate (FIR) bandwidth to each subscriber queue. This queue FIR must be converted into **LM-FoW** format to cap it by the queue PIR (*adminPIR*) and to make sure the sum of *FIRs* of all queues of the same subscriber does not exceed the subscriber **agg-rate-limit** which is also expressed in **LM-FoW** format. The conversion between these two rates makes use of the cumulative *RunningAverageFrameExpansionRatio* value.

Note that a queue **LM-foW** AdminPIR value will always be capped to the value of the local port **FoW** rate even if the conversion based on the current *RunningAverageFrameExpansionRatio* value indicates that a higher AdminPIR may be able to fill in the full line rate of the local port.

## Vport Determination and Evaluation

In the BNG application, host queues of all subscribers destined to the same downstream BSAN, for example, all SAPs on the egress port matching the same S-VLAN tag value, are parented to the same vport which matches the destination ID of the BSAN.

The BNG determines the parent vport of a subscriber host queue, which has the **port-parent** option enabled, by matching the destination string **dest** string associated with the subscriber with the string defined under a vport on the port associated with the subscriber.

The user configures the dest string match under the egress context of the Ethernet port associated with the subscriber:

**CLI Syntax:** `configure>port>ethernet>access>egress>vport>host-match dest string create`

If a given subscriber host queue does not have the **port-parent** option enabled, it will be foster-parented to the vport used by this subscriber and which is based on matching the **dest** string. If the subscriber could not be matched with a vport on the egress port, the host queue will not be bandwidth controlled and will compete for bandwidth directly based on its own PIR and CIR parameters.

By default, a subscriber host queue with the **port-parent** option enabled is scheduled within the context of the port's port scheduler policy. In order to indicate the option to schedule the queue in the context of a port scheduler policy associated with a vport, the user enters the following command in SLA profile used by the subscriber host:

**CLI Syntax:** `configure>subscriber-mgmt>sla-profile>egress>qos sap-egress-qos-policy-id vport-scheduler`

This command is persistent meaning that the user can re-enter the **qos** node without specifying the **vport-scheduler** argument each time and the system will remember it. The user can revert to the default setting without deleting the association of the SLA profile with the SAP egress QoS policy by explicitly re-entering the command with the following new argument:

**CLI Syntax:** `configure>subscriber-mgmt>sla-profile>egress>qos sap-egress-qos-policy-id port-scheduler`

---

## Applying Aggregate Rate Limit to a Vport

The user can apply an aggregate rate limit to the vport and apply a port scheduler policy to the port.

This model allows the user to oversubscribe the Ethernet port. The application of the **agg-rate** option is mutually exclusive with the application of a port scheduler policy, or a scheduler policy to a vport.

When using this model, a subscriber host queue with the **port-parent** option enabled is scheduled within the context of the port's port scheduler policy. However, the user must still indicate to the

system that the queues are managed by the aggregate rate limit instance of a vport by enabling the **vport-scheduler** option in the subscriber host SLA profile:

**CLI Syntax:** `configure>subscriber-mgmt>sla-profile>egress>qos sap-egress-qos-policy-id vport-scheduler`

A subscriber host-queue which is port-parented will be parented to the port scheduler policy of the port used by the subscriber and aggregate rate limited within the instance of the vport used by this subscriber and which is based on matching the **dest** string and **org** string. If the vport exists but the port does not have a port scheduler policy applied, then the host queue will be orphaned and no aggregate rate limit can be enforced.

If a given subscriber host queue does not have the port-parent option enabled, it will be foster-parented to the port used by this subscriber and aggregate rate limited within the instance of the vport used by this subscriber. If the vport exists but the port does not have a port scheduler policy applied, then the host queue will be orphaned and no aggregate rate limit can be enforced.

---

### Applying a Scheduler Policy to a Vport

The user can apply a scheduler policy to the vport. This allows scheduling control of subscriber tier 1 schedulers in a scheduler policy applied to the egress of a subscriber or SLA profile, or to a PW SAP in an IES or VPRN service.

The advantage of using a scheduler policy under a vport, compared to the use of a port scheduler (with or without an agg-rate rate), is that it allows a port parent to be configured at the vport level.

Bandwidth distribution from an egress port scheduler to a vport configured with a scheduler policy can be performed based on the level/cir-level and weight/cir-weight configured under the scheduler's port parent. The result is in allowing multiple vports, for example representing different DSLAMs, to share the port bandwidth capacity in a flexible way that is under the control of the user.

The configuration of a scheduler policy under a vport is mutually exclusive with the configuration of a port scheduler policy or an aggregate rate limit.

A scheduler policy is configured under a vport as follows:

**CLI Syntax:** `config>port>ethernet>access>egress>vport# scheduler-policy scheduler-policy-name`

When using this model, a tier 1 scheduler in a scheduling policy applied to a subscriber profile or SLA profiles must be configured as follows:

**CLI Syntax:** `config>qos>scheduler-policy>tier# parent-location vport`

If the vport exists, but port does not have a scheduler policy applied, then its schedulers will be orphaned and no port level QOS control can be enforced.



The following **show/monitor/clear** commands are available related to the vport scheduler:

```
show qos scheduler-hierarchy port port-id vport name [scheduler scheduler-name] [detail]
```

```
show qos scheduler-stats port port-id vport name [scheduler scheduler-name] [detail]
```

```
monitor qos scheduler-stats port port-id vport name [interval seconds ] [repeat repeat ]  
[absolute|rate]
```

```
clear qos scheduler-stats port port-id vport name [scheduler scheduler-name] [detail]
```

HQoS adjustment and host tracking are not supported on schedulers that are configured in a scheduler policy on a vport, so the configuration of a scheduler policy under a vport is mutually exclusive with the configuration of the egress-rate-modify parameter.

ESM over MPLS pseudowires are not supported when a scheduler policy is configured on a vport.

Signaling of Last Mile Encapsulation Type

A subscriber host session can signal one of many encapsulation types each with a different fixed offset in the last mile. These encapsulation types are described in RFC 4679 and are illustrated in Figure 58 and Figure 59. The BNG node learns the encapsulation type of each subscriber host session by inspecting the Access-loop-encapsulation sub-TLV in the Vendor-Specific PPPoE Tags as specified in RFC 4679. When Ethernet is the last mile, the encapsulation type will result in a fixed offset for all packet sizes. When ATM/DSL is the last mile, there will be an additional expansion due to AAL5 padding to next multiple of 48 bytes and which varies depending on the packet size.

Both ATM and Ethernet access using PPP encapsulation options are supported in the software and hardware based implementations. Thus both provide support for the Access-loop-encapsulation sub-TLV in the Vendor-Specific PPPoEv4/PPPoEv6 Tags with the ATM encapsulation values and Ethernet encapsulation values.

ATM and Ethernet access using IP encapsulation are only supported using default encapsulation offset configuration in the subscriber profile in the software based implementation. Support for signaling the Access-loop-encapsulation sub-TLV in the DHCPv4/DHCPv6 Relay Options is included in the hardware based implementation. There is no support for DHCPv6 relay options.

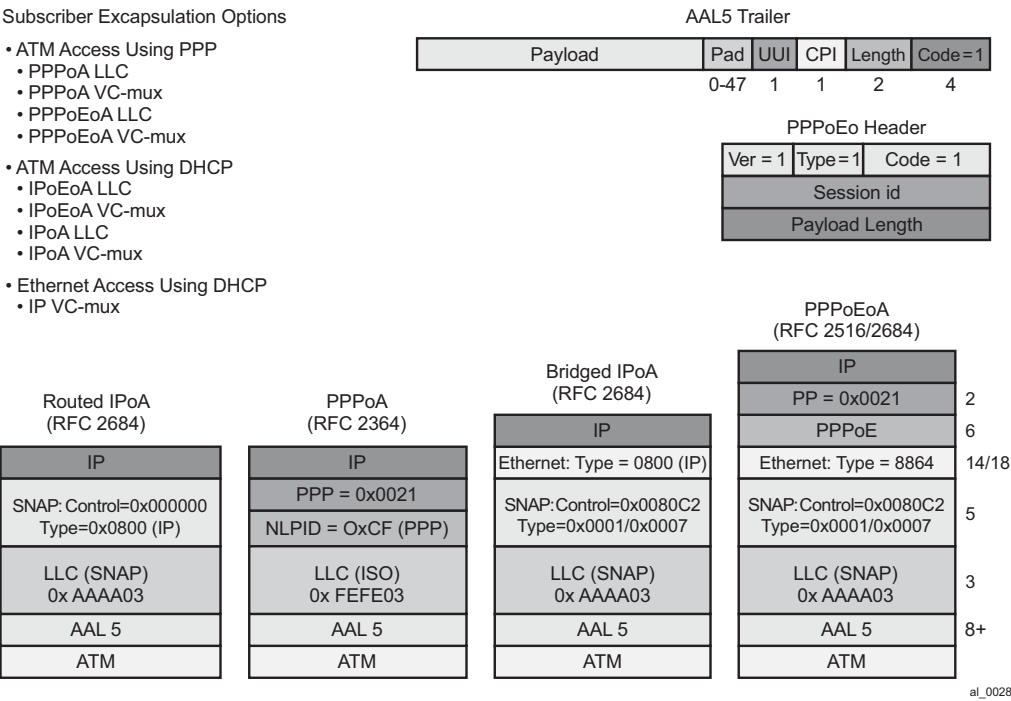



Figure 58: Subscriber Host Session Encapsulation Types

## Encapsulation combinations (RFC 4679)



## Encapsulation combinations

Valid values for the sub-fields are as follows:

## Data Link

```

0x00 AAL5
0x01 Ethernet

```

## Encaps 1

```

0x00 NA - Not Available
0x01 Untagged Ethernet
0x02 Single-Tagged Ethernet

```

## Encaps 2

```

0x00 NA - Not Available
0x01 PPPoA LLC
0x02 PPPoA Null
0x03 IPoA LLC
0x04 IPoA Null
0x05 Ethernet over AAL5 LLC with FCS
0x06 Ethernet over AAL5 LLC without FCS
0x07 Ethernet over AAL5 Null with FCS
0x08 Ethernet over AAL5 Null without FCS

```

- AAL5
  - PPPoA LLC/Null
  - IPoA LLC/Null
  - Ethernet over ATM x 4
    - Tagged/Untagged PPP
    - Tagged/Untagged DHCP
  - Total of 20 AAL5 combinations
- Ethernet
  - Tagged/Untagged PPP
  - Tagged/Untagged DHCP
  - Total of four Ethernet combinations
- *Total of 24 access combinations*

**Figure 59: Access-Loop-Encapsulation Sub-TLV**

The operational last-mile values for hosts on the same sap, having the same SLA profile are displayed in following the show-command:

**CLI Syntax: show>service active-subscribers>ale-adjust**

The data-link can have values: atm, other and unknown. If no offset is supplied it will be set to unknown. Other is used when the data-link is non-atm, otherwise it will state atm.

Operational per-queue values can also be found in the show-command:

**CLI Syntax: show>qos>scheduler-hierarchy**

Here, one can see whether the queue is operating in last-mile mode. Note that this command is not available on HSMDAv2,

**Last Mile ATM**

```

*A:Dut-C# /show service active-subscribers ale-adjust
=====
Active Subscriber Access Loop Encapsulation adjustment
=====
Subscriber
SAP                               SLA profile
Data-link Offset (bytes)

```

```

-----
hpolSub81
  1/1/11:2000.1
  atm          -10
                                hpolSlaProf1
-----

No. of Access Loop Encapsulation adjustments: 1
=====

*A:Dut-C# show qos scheduler-hierarchy subscriber "hpolSub81"
=====
Scheduler Hierarchy - Subscriber hpolSub81
=====
Ingress Scheduler Policy:
Egress Scheduler Policy :
-----

Root (Ing)
|
No Active Members Found on slot 1

Root (Egr)
| slot(1)
|--(Q) : Sub=hpolSub81:hpolSlaProf1 2000->1/1/11:2000.1->8->ATM (Port 1/1/11)
|
|--(Q) : Sub=hpolSub81:hpolSlaProf1 2000->1/1/11:2000.1->7->ATM (Port 1/1/11)
|
|--(Q) : Sub=hpolSub81:hpolSlaProf1 2000->1/1/11:2000.1->6->ATM (Port 1/1/11)
|
|--(Q) : Sub=hpolSub81:hpolSlaProf1 2000->1/1/11:2000.1->5->ATM (Port 1/1/11)
|
|--(Q) : Sub=hpolSub81:hpolSlaProf1 2000->1/1/11:2000.1->4->ATM (Port 1/1/11)
|
|--(Q) : Sub=hpolSub81:hpolSlaProf1 2000->1/1/11:2000.1->3->ATM (Port 1/1/11)
|
|--(Q) : Sub=hpolSub81:hpolSlaProf1 2000->1/1/11:2000.1->2->ATM (Port 1/1/11)
|
|--(Q) : Sub=hpolSub81:hpolSlaProf1 2000->1/1/11:2000.1->1->ATM (Port 1/1/11)
|

```

### Last Mile Ethernet

```

*A:Dut-C# show service active-subscribers ale-adjust
=====
Active Subscriber Access Loop Encapsulation adjustment
=====
Subscriber
  SAP                                SLA profile
  Data-link Offset (bytes)
-----
hpolSub81
  1/1/11:2000.1
  other          +12
                                hpolSlaProf1
-----

No. of Access Loop Encapsulation adjustments: 1
=====

*A:Dut-C# show qos scheduler-hierarchy subscriber "hpolSub81"

```

```

=====
Scheduler Hierarchy - Subscriber hpolSub81
=====
Ingress Scheduler Policy:
Egress Scheduler Policy :
-----
Root (Ing)
|
No Active Members Found on slot 1

Root (Egr)
| slot(1)
|--(Q) : Sub=hp01Sub81:hp01SlaProf1 2000->1/1/11:2000.1->8->Eth (Port 1/1/11)
|
|--(Q) : Sub=hp01Sub81:hp01SlaProf1 2000->1/1/11:2000.1->7->Eth (Port 1/1/11)
|
|--(Q) : Sub=hp01Sub81:hp01SlaProf1 2000->1/1/11:2000.1->6->Eth (Port 1/1/11)
|
|--(Q) : Sub=hp01Sub81:hp01SlaProf1 2000->1/1/11:2000.1->5->Eth (Port 1/1/11)
|
|--(Q) : Sub=hp01Sub81:hp01SlaProf1 2000->1/1/11:2000.1->4->Eth (Port 1/1/11)
|
|--(Q) : Sub=hp01Sub81:hp01SlaProf1 2000->1/1/11:2000.1->3->Eth (Port 1/1/11)
|
|--(Q) : Sub=hp01Sub81:hp01SlaProf1 2000->1/1/11:2000.1->2->Eth (Port 1/1/11)
|
|--(Q) : Sub=hp01Sub81:hp01SlaProf1 2000->1/1/11:2000.1->1->Eth (Port 1/1/11)
|

```

### Next Mile ATM

```

*A:Dut-C# show service active-subscribers ale-adjust subscriber "hp01Sub321"
=====
Active Subscriber Access Loop Encapsulation adjustment
=====
Subscriber
  SAP                               SLA profile
  Data-link Offset (bytes)
-----
hp01Sub321
  5/1/1:100/1                       hp01SlaProf1
  atm                               +8
-----
No. of Access Loop Encapsulation adjustments: 1
=====

*A:Dut-C# show qos scheduler-hierarchy subscriber "hp01Sub321"
=====
Scheduler Hierarchy - Subscriber hp01Sub321
=====
Ingress Scheduler Policy:
Egress Scheduler Policy :
-----
Root (Ing)
|
No Active Members Found on slot 5

```

## ESM Subscriber Hierarchical Traffic Control

```
Root (Egr)
| slot(5)
|--(Q) : Sub=hpolSub321:hpolslaProf1 2000->5/1/1:100/1->8:ATM (Port 5/1/1)
|
|--(Q) : Sub=hpolSub321:hpolslaProf1 2000->5/1/1:100/1->7:ATM (Port 5/1/1)
|
|--(Q) : Sub=hpolSub321:hpolslaProf1 2000->5/1/1:100/1->6:ATM (Port 5/1/1)
|
|--(Q) : Sub=hpolSub321:hpolslaProf1 2000->5/1/1:100/1->5:ATM (Port 5/1/1)
|
|--(Q) : Sub=hpolSub321:hpolslaProf1 2000->5/1/1:100/1->4:ATM (Port 5/1/1)
|
|--(Q) : Sub=hpolSub321:hpolslaProf1 2000->5/1/1:100/1->3:ATM (Port 5/1/1)
|
|--(Q) : Sub=hpolSub321:hpolslaProf1 2000->5/1/1:100/1->2:ATM (Port 5/1/1)
|
|--(Q) : Sub=hpolSub321:hpolslaProf1 2000->5/1/1:100/1->1:ATM (Port 5/1/1)
|
=====
```

## Configuration Example

The following CLI configuration achieves the specific use case shown in [Figure 57](#).

```

config
  qos
    port-scheduler-policy "dslam-vport-scheduler"
      group res-bus-be create
        rate 1000
        level 3 rate 1000 group res-bus-be weight w1
        level 4 rate 1000 group res-bus-be weight w4
        level 5 rate 1000 cir-rate 100
        level 7 rate 5000 cir-rate 5000
        level 8 rate 500 cir-rate 500
        max-rate 5000

    sap-egress 100 // residential policy
      queue 1 // be-res
      port-parent weight x level 3
      queue 2 // l2-res
      port-parent weight y level 3
      queue 3 // l1-res
      port-parent weight z level 3
      queue 4 // h2-res
      port-parent level 5
      queue 5 // h1-res
      port-parent level 7
      queue 6 // ef-res
      port-parent level 8
      fc be queue 1
      fc l2 queue 2
      fc l1 queue 3
      fc h2 queue 4
      fc h1 queue 5
      fc ef queue 6

    exit
    sap-egress 200 // business policy
      queue 1 // be-bus
      port-parent weight x level 4
      queue 2 // l2-bus
      port-parent weight y level 4
      queue 3 // l1-bus
      port-parent weight z level 4
      queue 4 // h2-bus
      port-parent level 5
      queue 5 // h1-bus
      port-parent level 7
      queue 6 // ef-bus
      port-parent level 8
      fc be queue 1
      fc l2 queue 2
      fc l1 queue 3
      fc h2 queue 4
      fc h1 queue 5
      fc ef queue 6

    exit
  exit

```

## ESM Subscriber Hierarchical Traffic Control

```
config
  sub-mgmt
    sla-profile "residential"
      egress
        qos 100 vport-scheduler
      exit
    exit
    sla-profile "business"
      egress
        qos 200 vport-scheduler
      exit
    exit
    sub-profile "residential"
      egress
        encap-offset
        avg-frame-size 1500
        agg-rate-limit 100
      exit
    exit
    sub-profile "business"
      egress
        encap-offset type pppoeoa-llc-tagged-fcs
        avg-frame-size 500
        agg-rate-limit 200
      exit
    exit
  exit
exit

config
  port 1/1/1
  ethernet
    access
      egress
        vport "dslam-1" create
        port-scheduler-policy "dslam-vport-scheduler"
        host-match dest "20" create
      exit
    exit
  exit
exit
exit
```



## Subscriber Volume Statistics

Subscriber volume statistics or octet and packet counters are available through the queues and policers that are instantiated for the subscriber. The queue and policer configuration is defined in the SLA profile via ingress and egress qos policy associations with optional overrides. Subscriber hosts that belong to the same subscriber, that are active on the same SAP and that have the same sla-profile will share the set of queues and policers defined by that SLA profile instance. For HSMDA hardware, the ingress policer/queue and egress queue configuration is defined in the subscriber profile via ingress and egress qos policy associations with optional overrides. All hosts belonging to the subscriber share the same set of queues and policers.

---

## IP (Layer 3) Volume Accounting

Subscriber volume statistics by default count Layer 2 frame sizes optionally modified by configuration such as packet-byte-offset, last mile aware shaping, etc.

To report subscriber volume statistics as Layer 3 (IP) packet sizes, the volume-stats-type can be configured to **ip** in the subscriber profile:

```
configure
  subscriber-mgmt
    sub-profile <subscriber-profile-name>
      volume-stats-type ip
```

**volume-stats-type ip** affects the subscriber statistics in SNMP, CLI, RADIUS accounting, XML accounting and Diameter Gx usage monitoring. Volume quota for RADIUS or Diameter Credit Control applications are interpreted as Layer 3 quota.

The following restrictions apply for **volume-stats-type ip**:

- Layer-3/IP accounting is not supported in combination with last-mile-aware shaping on HSMDA
- Layer-3/IP accounting is not supported in combination with ESMoPW on HSMDA
- Layer-3/IP accounting is not supported in combination with MLPPP
- Layer-3/IP accounting in combination with ESMoPW and last-mile-aware shaping may be inaccurate if the MPLS encapsulation overhead changes during the lifetime of a subscriber.
- Layer-3/IP accounting is restricted to a single encap per sla-profile instance (queue instance). The first host associated with the sla-profile instance (queue instance) determines the allowed encap. Conflicting encaps are:
  - PPPoE and IPoE on regular Ethernet SAPs
  - PPPoE and IPoE on PW-SAPs
  - PPPoA and PPPoEoA on ATM ports

- PPPoE keep alive packets do not contain IP payload and introduce an error in Layer-3/IP accounting when enabled in combination with L2TP-LAC. A workaround is to isolate the keep alives in a separate queue/policer.
  - Padding of frames smaller than the Ethernet minimum frame size (64B) may introduce an inaccuracy in Layer-3/IP accounting.
  - With ATM in the last mile, last-mile-aware shaping may introduce an inaccuracy in Layer-3/IP accounting.
  - Packet-Byte-Offset (PBO) changes during the lifetime of a subscriber introduces an inaccuracy in Layer-3/IP accounting.
- 

## Separate IPv4 and IPv6 Counters

IPv4 and IPv6 forwarded and dropped subscriber traffic can be counted separately via a **stat-mode v4-v6** command that is configured as a policer or queue qos override in the sla-profile. The **stat-mode v4-v6** command is only applicable for Enhanced Subscriber Management (ESM).

```
configure subscriber-mgmt
    sla-profile "sla-profile-1" create
        ingress
            qos 10
                queue 1
                    stat-mode v4-v6
                exit
            policer 1
                stat-mode v4-v6
            exit
        exit
    egress
        qos 10
            queue 1
                stat-mode v4-v6
            exit
        policer 1
            stat-mode v4-v6
        exit
    exit
exit
exit
```

For policers, the stat-mode command overrides the policer stat-mode configuration as defined in the sap-ingress or sap-egress qos policy. For details on sap-ingress and sap-egress policer stat-mode, refer to the 7750 SR OS Quality of Service Guide. For a policer in stat-mode v4-v6, following counters are available:

- Offered IPv4 octets and packets
- Offered IPv6 octets and packets
- Dropped IPv4 octets and packets

- Dropped IPv6 octets and packets
- Forwarded IPv4 octets and packets
- Forwarded IPv6 octets and packets

When a policer's stat-mode is changed while the sla profile is in use, any previous counter values are lost and any new counters are set to zero.

For queues, a stat-mode is only available for use in Enhanced Subscriber Management (ESM) context to enable separate IPv4/IPv6 counters. For a queue in stat-mode v4-v6, following counters are available:

- Offered High Priority, Low Priority, Uncolored, Managed octets and packets
- Dropped IPv4 octets and packets
- Dropped IPv6 octets and packets
- Forwarded IPv4 octets and packets
- Forwarded IPv6 octets and packets

A queue's stat-mode cannot be changed while the sla profile is in use.

Note that there will be no in-profile/out-of-profile forwarded and dropped counters for policers and queues in **stat-mode v4-v6**.

Non IP traffic (for example PPPoE LCP frames) is counted against the IPv4 counters.

The separate IPv4 and IPv6 forwarded and dropped counters are reported in

- SNMP
- CLI

```
show service active-subscribers detail
- - - snip - - -
-----
SLA Profile Instance statistics
-----

```

	Packets	Octets
Off. HiPrio	: 0	0
Off. LowPrio	: 1102685	1102685000
Off. Uncolor	: 0	0
Off. Managed	: 0	0
Queueing Stats (Ingress QoS Policy 10)		
Dro. HiPrio	: 0	0
Dro. LowPrio	: 0	0
For. InProf	: 0	0
For. OutProf	: 0	0
Dro. V4	: 0	0
Dro. V6	: 0	0
For. V4	: 367543	367543000
For. V6	: 735142	735142000
Queueing Stats (Egress QoS Policy 10)		
Dro. InProf	: 0	0
Dro. OutProf	: 0	0

## Subscriber Volume Statistics

For. InProf	: 0	0
For. OutProf	: 0	0
Dro. V4	: 0	0
Dro. V6	: 0	0
For. V4	: 367543	367543000
For. V6	: 735088	735088000

---

### SLA Profile Instance per Queue statistics

---

	Packets	Octets
Ingress Queue 1 (Unicast) (Priority) (Stats mode: v4-v6)		
Off. HiPrio	: 0	0
Off. LowPrio	: 1102685	1102685000
Dro. V4	: 0	0
Dro. V6	: 0	0
For. V4	: 367545	367545000
For. V6	: 735146	735146000

Egress Queue 1 (Stats mode: v4-v6)		
Dro. V4	: 0	0
Dro. V6	: 0	0
For. V4	: 367547	367547000
For. V6	: 735096	735096000

---

### SLA Profile Instance per Policer statistics

---

	Packets	Octets
Ingress Policer 1 (Stats mode: v4-v6)		
Off. V4	: 0	0
Off. V6	: 0	0
Dro. V4	: 0	0
Dro. V6	: 0	0
For. V4	: 0	0
For. V6	: 0	0
Egress Policer 1 (Stats mode: v4-v6)		
Off. V4	: 0	0
Off. V6	: 0	0
Dro. V4	: 0	0
Dro. V6	: 0	0
For. V4	: 0	0
For. V6	: 0	0

- RADIUS accounting

When a queue or policer is configured in stat-mode v4-v6, existing VSA's are re-used in RADIUS detailed per queue / per policer accounting (**configure subscriber-mgmt radius-accounting-policy <name> include-radius-attribute detailed-acct-attributes**):

- in-profile counter VSA's map to IPv4 octets/packets
- ingress queue high priority dropped counter VSA's map to IPv4 octets/packets
- out-of-profile counter VSA's map to IPv6 octets/packets
- ingress queue low priority dropped counter VSA's map to IPv6 octets/packets

In addition the [26-6527-107] Alc-Acct-I-statmode / [26-6527-127] Alc-Acct-O-statmode is sent with value set to "v4-v6".

Optionally a set of VSA's can be included in RADIUS accounting to report the aggregate IPv6 forwarded octets and packets of queues and policers with stat-mode v4-v6 enabled (configure subscriber-mgmt radius-accounting-policy <name> include-radius-attribute detailed-acct-attributes v6-aggregate-stats):

```
[26-6527-194] Alc-IPv6-Acct-Input-Packets
[26-6527-195] Alc-IPv6-Acct-Input-Octets
[26-6527-196] Alc-IPv6-Acct-Input-GigaWords
[26-6527-197] Alc-IPv6-Acct-Output-Packets
[26-6527-198] Alc-IPv6-Acct-Output-Octets
[26-6527-199] Alc-IPv6-Acct-Output-Gigawords
```

Refer to the 7750 SR-OS RADIUS Attributes Reference Guide for a detailed description of all counter attributes.

- XML accounting

The complete-subscriber-ingress-egress and custom-record-subscriber XML records use following fields to represent IPv4 and IPv6 forwarded/dropped octets and packets for queues or policers with **stat-mode v4-v6** enabled:

```
v4po    - IPv4PktsOffered (policer only)
v4oo    - IPv4OctetsOffered (policer only)
v6po    - IPv6PktsOffered (policer only)
v6oo    - IPv6OctetsOffered (policer only)
v4pf    - IPv4PktsForwarded
v6pf    - IPv6PktsForwarded
v4pd    - IPv4PktsDropped
v6pd    - IPv6PktsDropped
v4of    - IPv4OctetsForwarded
v6of    - IPv6OctetsForwarded
v4od    - IPv4OctetsDropped
v6od    - IPv6OctetsDropped
```

For custom records, the following CLI is re-used to include v4/v6 counters if the queue is configured in **stat-mode v4-v6**:

```
i-counters
    all-packets-offered-count      # n/a
    all-octets-offered-count      # n/a
    high-packets-offered-count    # n/a
    low-packets-offered-count     # n/a
    uncoloured-packets-offered-count # n/a
    high-octets-offered-count     # n/a
    low-octets-offered-count      # n/a
    uncoloured-octets-offered-count # n/a
    all-packets-offered-count      # n/a
    all-octets-offered-count      # n/a
    high-packets-discarded-count  # IPv4
    low-packets-discarded-count   # IPv6
    high-octets-discarded-count   # IPv4
    low-octets-discarded-count    # IPv6
    in-profile-packets-forwarded-count # IPv4
    out-profile-packets-forwarded-count # IPv6
    in-profile-octets-forwarded-count # IPv4
    out-profile-octets-forwarded-count # IPv6

e-counters
    in-profile-packets-forwarded-count # IPv4
    in-profile-packets-discarded-count # IPv4
    out-profile-packets-forwarded-count # IPv6
    out-profile-packets-discarded-count # IPv6
    in-profile-octets-forwarded-count   # IPv4
    in-profile-octets-discarded-count   # IPv4
    out-profile-octets-forwarded-count  # IPv6
    out-profile-octets-discarded-count  # IPv6
```

On HSMDB, **stat-mode v4-v6** is configured as a policer or queue qos override in the subscriber profile:

```
configure subscriber-mgmt
    sub-profile "sub-profile-2" create
        hsmdb
            egress-qos
                qos 10
                    queue 1
                        stat-mode v4-v6
                    exit
                exit
            exit
        ingress-qos
```

```
        qos 10
            policer 1
                stat-mode v4-v6
            exit
        queue 1
            stat-mode v4-v6
        exit
    exit
exit
exit
exit
```

The stat-mode on egress hsmda queues is always enabled per subscriber: when enabling **stat-mode v4-v6** for one hsmda queue, it is automatically enabled for all hsmda queues for that subscriber profile.

## Configuring IP and IPv6 Filter Policies for Subscriber Hosts

This section applies to the 7750 SR and 7450 ESS.

Access Control Lists (ACLs) for subscriber traffic are defined as IP and IPv6 filter policies and are configured in the SLA-profile associated with the subscriber. For information about IP and IPv6 filter policy configurations, refer to the 7750 SR-OS Router Configuration Guide.

**CLI Syntax:**

```
config>subscr-mgmt>sla-prof
      sla-profile sla-profile-1 create
      ingress
        ip-filter 100
        ipv6-filter 300
      exit
      egress
        ip-filter 200
        ipv6-filter 400
      exit
exit
```

Traffic from different hosts of a single subscriber and associated with the same sla-profile instance, is subject to the filter policies defined in the SLA profile.

The IP or IPv6 filter policy configuration of subscriber hosts can be dynamically updated using different mechanisms:

1. Assign a new SLA profile.

This can be done dynamically by, for example, a RADIUS CoA message. As the SLA profile also defines the QoS configuration for the subscriber hosts, this change may result in a discontinuity in accounting.

Note: changing the ip-filter policy in an SLA profile in use by an active subscriber is allowed in the CLI, but not recommended. Changing the IPv6 filter policy in an SLA profile in use by an active subscriber is prevented in the CLI.

2. Override the IP and IPv6 filter policies

Alternatively, it is also possible to dynamically override the IP and IPv6 filter policies per subscriber-host through a RADIUS Access-Accept or CoA message. Following VSA should be included in the RADIUS message:

3. Insert subscriber host specific filter entries

A subscriber host specific entry is dynamically created from a RADIUS access accept or CoA message using the NAS-Filter-Rule or Alc-Ascend-Data-Filter-Host-Spec attribute (see also [IP Filter Attribute Format Details on page 867](#) for a detailed description of the attribute format):



Attribute ID	Attribute Name	Type	Limits	SR OS Format
92	NAS-Filter-Rule	String	max. 10 attributes per message or max. 10 filter entries per message	<p>The format of a NAS-Filter-Rule is defined in rfc-3588 section-4.3. A single filter rule is a string of format "&lt;action&gt; &lt;direction&gt; &lt;protocol&gt; from &lt;source&gt; to &lt;destination&gt; &lt;options&gt;". Multiple rules should be separated by a NUL (0x00). A NAS-Filter-Rule attribute may contain a partial rule, one rule, or more than one rule. Filter rules may be continued across attribute boundaries.</p> <p>A RADIUS message with NAS-Filter-Rule attribute value equal to 0x00 or " " (a space) removes all host specific filter entries for that host.</p> <p>See also <a href="#">IP Filter Attribute Format Details on page 867</a>.</p> <p>For example: Nas-Filter-Rule = "permit in ip from any to 10.1.1.1/32"</p>
26-6527-159	Alc-Ascend-Data-Filter-Host-Spec	octets	<p>max. 10 attributes per message or max. 10 filter entries per message.</p> <p>min. length 22 bytes (IPv4), 46 bytes (IPv6)</p> <p>max. length: 110 bytes (IPv4), 140 bytes (IPv6)</p>	<p>a string of octets with fixed field length (type (ipv4/ipv6), direction (ingress/egress), src-ip, dst-ip, ...). Each attribute represents a single filter entry. See <a href="#">IP Filter Attribute Format Details on page 867</a> for a description of the format.</p> <p>For example: # "permit in ip from any to 10.1.1.1/32" Alc-Ascend-Data-Filter-Host-Spec = 0x010101000000000000a0101010020000000000000000000</p>

The Alc-Subscriber-Filter VSA is a comma separated list of strings:

Field	Use
Ingr-v4:<number>	Ingress ipv4 filter

Field	Use
Egr-v4:<number>	Egress ipv4 filter
Ingr-v6:<number>	Ingress ipv6 filter
Egr-v6:<number>	Egress ipv6 filter

The filter number can have following values:

<number>	Result
1..65535	Ignore filter from sla-profile configuration and assign corresponding pre-configured filter
0	Ignore filter from sla-profile configuration and do not assign a new filter (only allowed if no dynamic subscriber host specific rules are present)
-1	No change in filter configuration
-2	Restore filter from sla-profile configuration

Notes:

- Not relevant fields (IPv4 filters for an IPv6 host) will be ignored.
- RADIUS CoA message: if the ingress or egress field is missing in the VSA, there will be no change for that direction.
- RADIUS Access-Accept message: if the ingress or egress field is missing in the VSA, then the IP filters as specified in the SLA profile will be active for that direction.

An SLA profile IP filter override is applicable to all dynamic host types, including L2TP LNS but excluding L2TP LAC.

#### 4. Insert subscriber host specific filter entries

A subscriber host specific entry is a filter entry where the match criteria is automatically extended with the subscriber host IP or IPv6 address as source (ingress) or destination (egress) IP. They represent a per host customization of a generic filter policy: only traffic to/from the subscriber host will match against these entries.

A subscriber host specific entry is dynamically created from a RADIUS access accept or CoA message using the NAS-Filter-Rule attribute:

The format used to specify host specific filter entries (NAS-Filter-Rule format or Alc-Ascend-Data-Filter-Host-Spec format) cannot change during the lifetime of the subscriber host. A RADIUS message can only contain a single format for host specific filter entries.

Up to 10 host specific filter rules can be specified in a single RADIUS message. Each new RADIUS CoA message containing host specific filter attributes overwrites the previous subscriber host-specific filter entries for that host provided that there are enough free entries in the reserved range.

Subscriber host specific filter entries can be removed with a RADIUS CoA message with NAS-Filter-Rule attribute value equal to 0x00 or " " (a space).

When the subscriber host session terminates or is disconnected the corresponding subscriber host specific filter entries are also deleted.

Note that subscriber host-specific filter entries are moved if the subscriber host filter policy is changed (new SLA profile or ip filter policy override) and the new filter policy contains enough free reserved entries (sub-insert-RADIUS).

A range of entries must be reserved for subscriber host specific entries in a filter policy:

**CLI Syntax: config>filter**

ip-filter 100 create

sub-insert-radius start-entry 1000 count 100

High and low watermarks can be configured to raise an event when the thresholds of free entries in the reserved range are reached:

**CLI Syntax: config>filter>ip-filter# sub-insert-wmark ?**

- no sub-insert-wmark

- sub-insert-wmark low <low-watermark> high <high-watermark>

<low-watermark> : [0..100]

<high-watermark> : [0..100]

Use following show commands to check filter policy details and the filter configuration for a subscriber host:

```
# show filter ip <ip-filter-id> type <entry-type>
# show filter ipv6 <ipv6-filter-id> type <entry-type>
    <entry-type>          : fixed|radius-insert|credit-control-insert

# show service active-subscribers filter [subscriber <sub-ident-string>] [origin <origin>]
    <sub-ident-string>    : [32 chars max]
    <origin>              : radius|credit-control
```

## 5. Insert shared filter entries

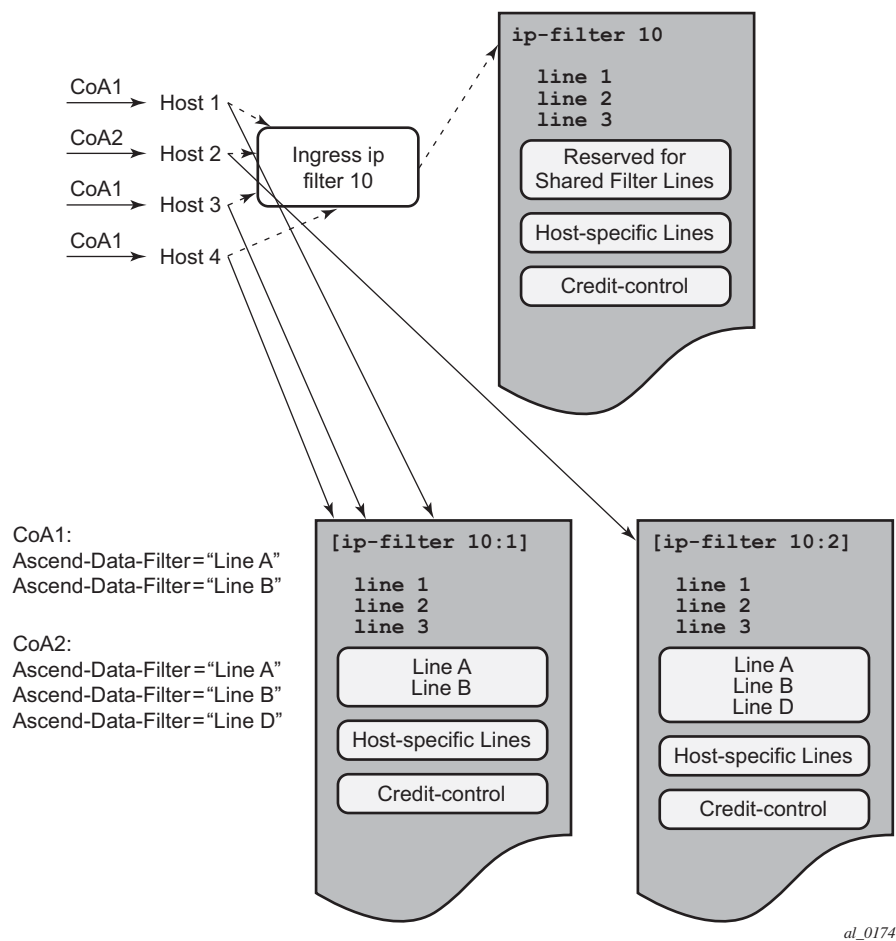
The target application for RADIUS shared filter entries is operators that have a predefined limited number of different filter lists that each are shared with multiple subscriber hosts and that are to be managed and activated from RADIUS at authentication.

A local configured ip or ipv6 filter associated with a host (sla-profile or host filter override) can be enhanced with dynamic filter entries that can be shared with multiple subscriber hosts. The shared dynamic filter entries are inserted with a set of RADIUS attributes "[242]"

Ascend-Data-Filter" or "[26-6527-158] Alc-Nas-Filter-Rule-Shared" received in a RADIUS Access-Accept or CoA message. A CoA message containing a set of one of those attributes overrides the previous set of shared filter entries active for that subscriber host.

For each unique set of dynamic filter entries received per type (ipv4/ipv6) and direction (ingress/egress), a copy is made of the local filter with the dynamic entries included at a preconfigured insert point. If the same set of dynamic filter entries is sent to subscriber hosts that have the same associated local filter, then they will share the same filter copy. When there are no more subscriber hosts associated with a filter copy, then the filter copy is deleted. A filter copy is identified as *local filter id:number*. For example: show filter ip 10:2

Shared filter entries are moved if the subscriber host filter policy is changed (new SLA profile or ip filter policy override) and if the new filter policy contains enough free reserved entries.



**Figure 60: Insert Shared Filters**

A range of entries must be reserved for shared entries in a filter policy:

**CLI Syntax:** **config>filter**

ip-filter 10 create

sub-insert-shared-radius start-entry 100 count 10

High and low watermarks can be configured to raise an event when the thresholds of dynamic filter copies are reached:

**CLI Syntax:** **config>filter>ip-filter# shared-radius-filter-wmark ?**

- no shared-radius-filter-wmark

- shared-radius-filter-wmark low *low-watermark* high *high-watermark*

low-watermark : [0..8000]

high-watermark : [0..8000]

The format used to specify shared filter entries (Alc-Nas-Filter-Rule-Shared format or Ascend-Data-Filter format) cannot change during the lifetime of the subscriber host. A RADIUS message can only contain a single format for shared filter entries.

Shared filter entries can be removed with a RADIUS CoA message with Alc-Nas-Filter-Rule-Shared attribute value equal to 0x00 or “ ” (a space).

Attribute ID	Attribute Name	Type	Limits	SR OS Format
242	Ascend-Data-Filter	Octets	multiple attributes per RADIUS message allowed. min. length 22 bytes (IPv4), 46 bytes (IPv6) max. length: 110 bytes (IPv4), 140 bytes (IPv6)	a string of octets with fixed field length (type (ipv4/ipv6), direction (ingress/egress), src-ip, dst-ip, ...). Each attribute represents a single filter entry. See <a href="#">IP Filter Attribute Format Details on page 867</a> for a description of the format. For example: # "permit in ip from any to 10.1.1.1/32" Ascend-Data-Filter = 0x010101000000000000a0101010020000000000000000000
26-6527-158	Alc-Nas-Filter-Rule-Shared	string	Multiple attributes per RADIUS message allowed.	The format is identical to [92] NAS-Filter-Rule and is defined in rfc-3588 section-4.3. A single filter rule is a string of format "<action> <direction> <protocol> from <source> to <destination> <options>". Multiple rules should be separated by a NUL (0x00). An Alc-Nas-Filter-Rule-Shared attribute may contain a partial rule, one rule, or more than one rule. Filter rules may be continued across attribute boundaries. A RADIUS message with Alc-Nas-Filter-Rule-Shared attribute value equal to 0x00 or " " (a space) removes the shared filter entries for that host. See also <a href="#">IP Filter Attribute Format Details on page 867</a> . For example: Alc-Nas-Filter-Rule-Shared = "permit in ip from any to 10.1.1.1/32"

## IP Filter Attribute Format Details

The format for [92] Nas-Filter-Rule and [26-6527-158] Alc-Nas-Filter-Rule-Shared is a string formatted as: “*action direction protocol from source to destination options*”. Refer to the table below for details on the respective fields.

Action or Classifier	Value		Corresponding SR-OS Filter Function
<direction>	in		ingress
	out		egress
<protocol>	ip		
	any number [0..255]		
	ip		
	any number [1..42]		
	any number [45..49]		
	any number [52..59]		
	any number [61..255]		
	any number 43 44 50 51 60		
from <source>	any	100	ingress: src-ip = host-ip-address; src-port eq 100 egress: src-ip = 0.0.0.0/0   ::/0; src-port eq 100
		200-65535	ingress: src-ip = host-ip-address; src-port range 200 65535 egress: src-ip = 0.0.0.0/0   ::/0; src-port range 200 65535
	ip-prefix/ length	100	ingress: src-ip = host-ip-address; src-port eq 100 egress: src-ip = ip-prefix/length; src-port eq 100
		200-65535	ingress: src-ip = host-ip-address; src-port range 200 65535 egress: src-ip = ip-prefix/length; src-port range 200 65535
	any	100	ingress: dst-ip = 0.0.0.0/0   ::/0; dst-port eq 100 egress: dst-ip = host-ip-address; dst-port eq 100
		200-65535	ingress: dst-ip = 0.0.0.0/0   ::/0; dst-port range 200 65535 egress: dst-ip = host-ip-address; dst-port range 200 65535

## Configuring IP and IPv6 Filter Policies for Subscriber Hosts

Action or Classifier	Value		Corresponding SR-OS Filter Function
	ip-prefix/ length	100	ingress: dst-ip = ip-prefix/length; dst-port eq 100 egress: dst-ip = host-ip-address; dst-port eq 100
		200-65535	ingress: dst-ip = ip-prefix/length; dst-port range 200 65535 egress: dst-ip = host-ip-address; dst-port range 200 65535
<options: frag>	frag		fragment true (ipv4 only)
<options: ipoptions>	ssrr		ip-option 9 / ip-mask 255
	lsrr		ip-option 3/ ip-mask 255
	rr		ip-option 7/ ip-mask 255
	ts		ip-option 4/ ip-mask 255
	!ssrr		not supported
	!lsrr		not supported
	!rr		not supported
	!ts		not supported
	ssrr,lsrr,rr, ts		not supported
<options: tcptions>	mss		not supported
	window		not supported
	sack		not supported
	ts		not supported
	!mss		not supported
	!window		not supported
	!sack		not supported
	!ts		not supported
	mss>window,sack,ts		not supported
<options: established>	established		not supported
			not supported
			not supported



Action or Classifier	Value		Corresponding SR-OS Filter Function
<options: setup>	setup		tcp-syn true
			tcp-ack false
			protocol tcp
<options: tcpflags>	syn		tcp-syn true
	!syn		tcp-syn false
	ack		tcp-ack true
	!ack		tcp-ack false
	fin		not supported
	rst		not supported
	psh		not supported
	urg		not supported
<options: icmpypesv4>	echo reply		protocol 1 / icmp-type 0
	destination unreachable		protocol 1 / icmp-type 3
	source quench		protocol 1 / icmp-type 4
	redirect		protocol 1 / icmp-type 5
	echo request		protocol 1 / icmp-type 8
	router advertisement		protocol 1 / icmp-type 9
	router solicitation		protocol 1 / icmp-type 10
	time-to-live exceeded		protocol 1 / icmp-type 11
	IP header bad		protocol 1 / icmp-type 12
	timestamp request		protocol 1 / icmp-type 13
	timestamp reply		protocol 1 / icmp-type 14
	information request		protocol 1 / icmp-type 15
	information reply		protocol 1 / icmp-type 16
	address mask request		protocol 1 / icmp-type 17
	address mask reply		protocol 1 / icmp-type 18

## Configuring IP and IPv6 Filter Policies for Subscriber Hosts

Action or Classifier	Value	Corresponding SR-OS Filter Function
	-	protocol 1 / icmp-type [0..255]
	3-9 ( range)	not supported
	3,5,8,9 ( comma seperated)	not supported
<options: icmptypesv6>	destination unreachable	icmp-type 1
	time-to-live exceeded	icmp-type 3
	IP header bad	icmp-type 4
	echo request	icmp-type 128
	echo reply	icmp-type 129
	router solicitation	icmp-type 133
	router advertisement	icmp-type 134
	redirect	icmp-type 137

The format for [242] Ascend-Data-Filter and [26-6527-159] Alc-Ascend-Data-Filter-Host-Spec is an octet string with fixed length fields. Refer to the table below for details on the respective fields.

Field	Length	Value
Type	byte	1 = IPv4
		3 = IPv6
Filter or forward	1 byte	0 = drop
		1 = accept
Indirection	1 byte	0 = egress
		1 = ingress
Spare	1 byte	ignored
Source IP address	IPv4 = 4 bytes	IP address of the source interface
	IPv6 = 16 bytes	
Destination IP address	IPv4 = 4 bytes	IP address of the destination interface
	IPv6 = 16 bytes	

Field	Length	Value
Source IP prefix	1 byte	Number of bits in the network portion
Destination IP prefix	1 byte	Number of bits in the network portion
Protocol	1 byte	Protocol number. Note: match the inner most header only for IPv6
Established	1 byte	ignored (not implemented)
Source port	2 bytes	Port number of the source port
Destination port	2 bytes	Port number of the destination port
Source port qualifier	1 byte	0 = no compare
		1 = less than
		2 = equal to
		3 = greater than
		4 = not equal to (not supported)
destination port qualifier	1 byte	0 = no compare
		1 = less than
		2 = equal to
		3 = greater than
		4 = not equal to (not supported)
Reserved	2 bytes	ignored

### Checking Filter Policy Details

Use following show commands to check filter policy details and the filter configuration for a subscriber host:

**CLI Syntax:**

```
show filter ip ip-filter-id detail
show filter ipv6 ip-filter-id detail
show filter ip ip-filter-id type entry-type
show filter ipv6 ipv6-filter-id type entry-type
    entry-type : fixed | radius-insert | credit-control-insert | radius-shared
show service active-subscribers filter [subscriber sub-ident-string] [origin origin]
    sub-ident-string : [32 chars max]
    origin : radius | credit-control"
```

## ESM PPPoA/PPPoEoA

This section applies to the 7750 SR and 7450 ESS.

The main goal of PPP in the subscriber context is to provide authentication, to negotiate link layer parameters (such as MTU) and to negotiate IP parameters (IP address, WINS, DNS, Default Gateway, etc.).

Each PPP session is carried over a single ATM VC over to the BNG. In PPPoA environment, PPP session is directly encapsulated over ATM transport on a DSL Customer Premise Equipment (CPE).

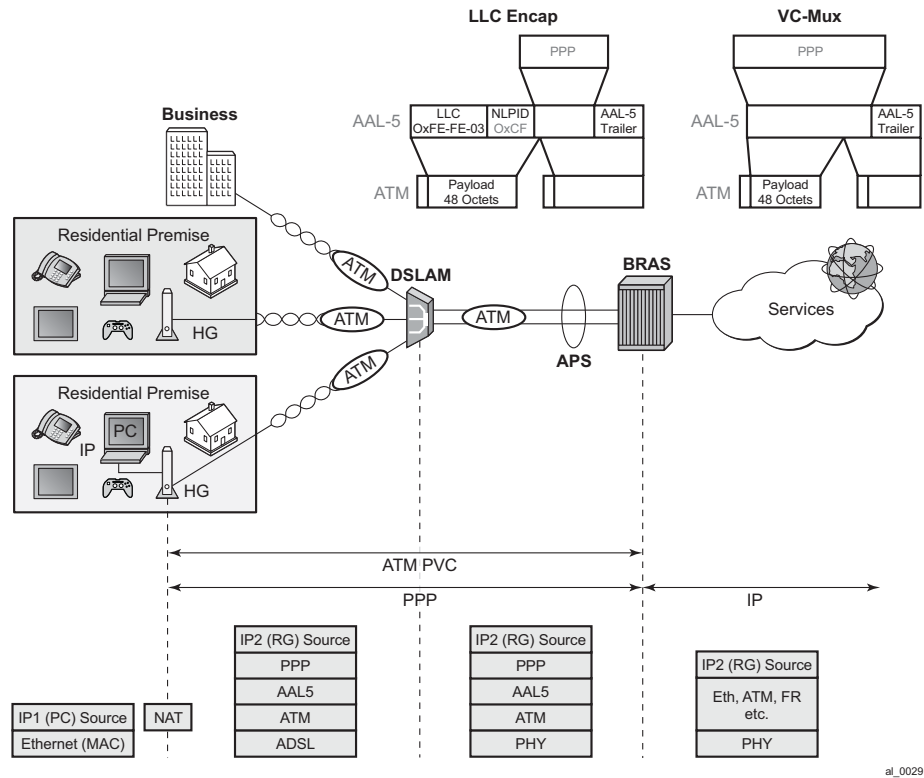
In the PPPoEoA environment, an additional layer is added to accommodate Ethernet medium. In this fashion, a PPP session can be directly terminated on any host within the customer Ethernet network and then transported over an ATM network to the BNG. Multiple PPPoE sessions can be carried over a single ATM PVC.

However, the majority of current implementations in ATM transport networks have PPPoE session terminated at the DSL CPE and not on a host within the customer network. This PPPoE session is then transported over ATM to the BNG. Although it seems unnecessary to add the overhead associated with Ethernet on an ATM-equipped DSL CPE, mainly for historical reasons PPPoE became engrained in the home network and as such has moved into DSL CPE.

---

## PPPoA

In a PPPoA environment, services are offered to residential and business customers. DSL CPE and BNG are the originating and terminating points of a PPP connection. In a residential example, a DSL CPE is a Home Gateway (HG) as shown in [Figure 61](#).



**Figure 61: PPPoA Architecture and Packet Encapsulation**

An ATM VC must exist between a CPE and a BNG before a PPP session can be established. From a QoS perspective, this VC is one of the following service categories: CBR, rt-VBR, nrt-VBR or UBR(+). CPE starts negotiating PPP link/session parameters over the VC. Once the IP address is obtained from the Service Provider (SP) side, the customer is ready for data transfer. There is a 1:1 mapping between a PPPoA session and a VC.

There is no need for the customer to run PPP within its own network. CPE, as a default gateway, accepts Ethernet IP packets, strips off the Ethernet header, performs a NAT function, and encapsulates the IP packets into PPPoA before it sends them on to the BNG.

In most residential cases, this PPPoA session is used for pure data transport (Internet access). As such, the BNG side would require a single service queue per VC.

In certain cases, customers use PPPoA for VoIP. In this case, they use an IAD (Integrated Access Device) to gain access to the ATM network.

It must be noted here that nothing in this architecture precludes customers from using multiple services over a single PPPoA session. Services would be differentiated through DSCP bits and each service in this case would require a separate queue. This is most likely a scenario with business customers where a customer runs multiple services over a PPPoA session. DSL CPE

would play a role here in differentiating and appropriately marking the services, either through a separate physical port per service or through some other means.

## PPPoEoA

In the PPPoE model the originating PPP point could extend beyond the DSL CPE and into the customer Ethernet network where any host can originate a PPP session. In such case, ATM is generally used as a transport to carry PPPoE sessions that are originated by hosts within the customer network (beyond DSL CPE). The DSL CPE would operate in a bridged mode. This is shown in [Figure 62](#).

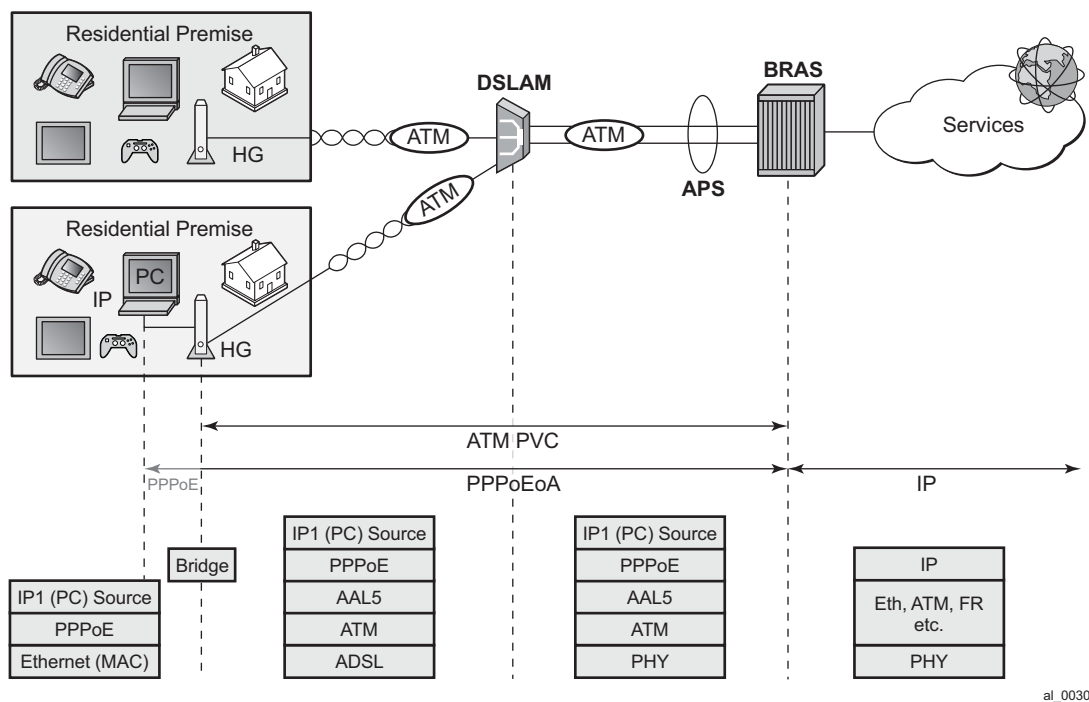
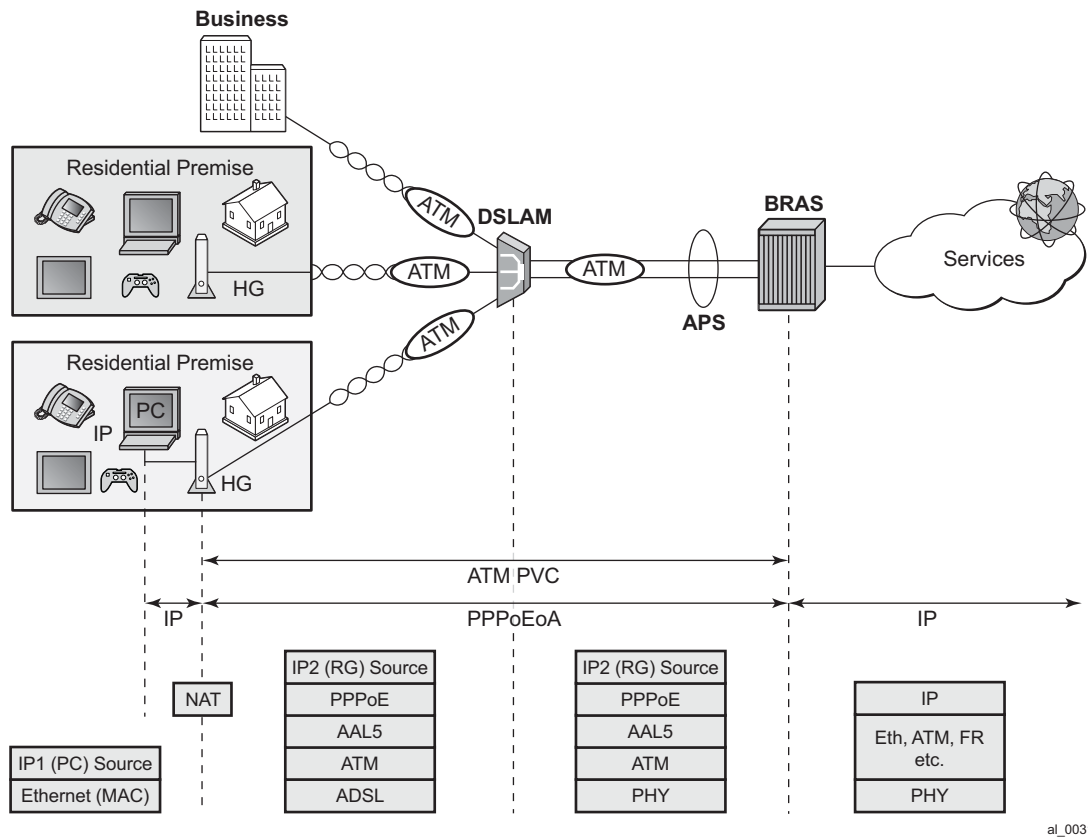


Figure 62: PPOEoA Host Terminated Session

However, the majority of deployments have PPPoE sessions terminated in DSL CPE. Although it is inefficient to add an extra Ethernet encapsulation layer over ATM-equipped DSL CPE, the evolution of PPP in broadband is the chief reason for this deployment scenario as shown in Figure 63.



**Figure 63: PPPoEoA DSL CPE Terminated Session**

PPPoEoA implementation must allow multiple PPPoE sessions of the same subscriber to be carried over a single ATM PVC.



## Hardware Support

This feature is supported on ATM MDA on:

- 7750SR (SR-12, SR-7) platforms
- 7750 SR-c4/12 platforms
- 7450 platforms in Mixed Mode.

This feature is implemented only on IOM3 based hardware.

ATM MDAs:

- 16 port ATM OC-3/STM-1 (single rate) –
- 4 port ATM OC-3/12c/STM-1/4c (dual rate on a per port basis; oc3/stm1 or oc12/stm4; port speed can only be changed in groups of four ports).

Chassis modes B,C and D are supported.

PPPoEoA/PPPoA will NOT be supported on the following modules:

ASAP MDAs:

- 4 port channelized OC-3/STM-1 (IOM2-20G and IOM3-XP)
- 1 port channelized OC-12/STM-4
- 12 port channelized DS3/E3 (coax)
- 4 port channelized DS3/E3 (coax)

ATM CMA:

- 8 port T1/E1 ATM (RJ-48)

The 7750-c4/12 currently supports only the four port ATM MDA.

## Termination Points within 7x50

PPPoA/PPPoEoA sessions are terminated on the access ATM SAP in IES and IP-VPN service context through subscriber/group interfaces.

However, for the wholesale/retail deployment scenarios, the ATM VCs are terminated on the LAC while the PPP(oE) sessions will be terminated on the LNS.

PPPoA/PPPoEoA is not supported in the wholesale/retail VRF model (wholesale VRF + retailer VRF). However it is supported in wholesale/retail MSAP model (capture SAP in VPLS that is mapped into a VRF).

---

## PPPoA Encapsulation


PPP frames are transported over ATM using ATM AAL5 framing mechanism. This is defined in RFC 2364. In short, each PPP packet is appended by an 8-octet AAL5 trailer with some control information (16-bit length field and a 32-bit CRC being the most important). This new frame is not self-identifying, in other words, kind of payload it carries (PPP, IPv4, IPv6, ARP, MPLS, etc.) can not be identified. This means that if you want to send it as such, it can carry only a single protocol type which must be agreed upon in advance by configuration at each side of a PVC connection. Only then will both ends of the connection be able to recognize the payload type inside of it.

To add more flexibility to the payload type and allow multiple protocol types to be multiplexed within a single VC session, an additional header must be defined in the AAL-5 packet. This header which allows protocol multiplexing over ATM VC is called Link Layer Control (LLC) header. PPP transport over ATM using LLC is defined in RFC 2364, *PPP over AAL5*. A more extensive version for multiplexing protocols over VC is defined in RFC 2684 “Multiprotocol Encapsulation over ATM AAL5” (LLC/SNAP encap).

7x50 supports these two types of PPPoA encapsulation:

1. LLC (RFC 2364)

LLC header is extended with a NLPID (Network Layer Protocol Identifier) which identifies the protocol type inside of the AAL5 frame. NLPID for PPP in LLC is 0xCF.



**Figure 64: PPPoA LLC Encapsulation**

Possible CLI syntax:

```
encapsulation aal5nlpid-ppp
```

The **ppp** keyword in the **aal5nlpid-ppp** command is used to indicate that ppp is the only encapsulation that currently supported in NLPID.

SNAP is an extension to LLC for protocols that are not defined in LLC NLPID. PPP protocol identifier is defined within NLPID and therefore it does not need additional SNAP header (this reduces the overall AAL5 header overhead).


LLC/SNAP encapsulation is defined in RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*, and it is needed for PPPoEoA encapsulation.

## 2. VC-MUX

In VC-MUX mode there is no additional (LLC/SNAP) header used for protocol multiplexing. Instead, VC endpoints must agree before hand on the payload type that they will transport. For PVCs this is done during the provisioning phase on each side of the connection. For example:

```
encapsulation aal5mux-ppp
```

aal5mux-ppp will tell each end of the VC in advance that the payload inside of the AAL5 frame is PPP.



**Figure 65: PPPoA AAL5MUX Encapsulation**

PPPoEoA Encapsulation

Similar to PPPoA, two types of encapsulation are defined for PPPoEoA:

- Protocol multiplexing (Layer 2 bridging defined by additional headers - LLC/SNAP).
- VC-Multiplexing where payload type within a VC is agreed upon in advance by static configuration.

Both types are supported in our implementation:

1. “Multiprotocol Encapsulation over ATM AAL5” (LLC/SNAP) defined in RFC 2684.

This encapsulation adds support for protocols that are currently not defined in the LLC header. There are two basic types of encapsulations defined under this RFC:

- Routed
- Bridged

PPPoEoA encapsulation is defined as ‘Bridged’ where Layer 2 information is preserved through transitioning between two different Layer 2 network types (Ethernet -> ATM) which is shown in Figure 66.

0xAA-AA-03 in the LLC header indicates the presence of the SNAP header.

0x00-80-C2 in the OUI indicates that a bridged PDU is encapsulated.

0x00-01 or 0x00-07 in PID indicates that the encapsulated Layer 2 network type is 802.3 Ethernet with or without preserved FCS.

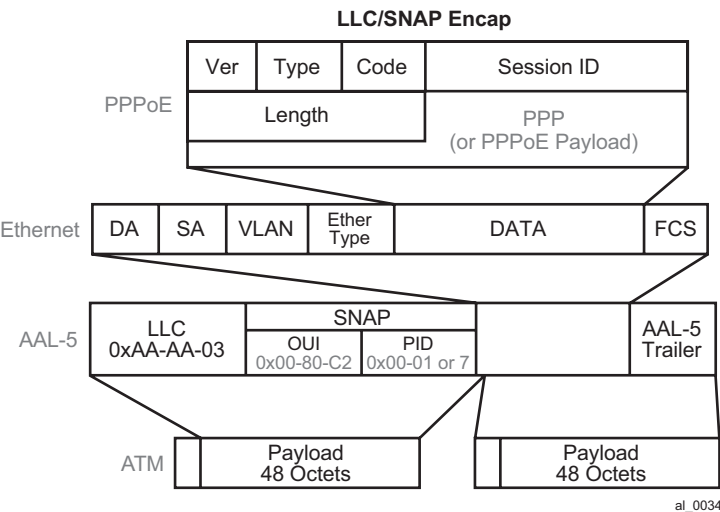


Figure 66: PPPoEoA Bridged LLC/SNAP Encapsulation

The CLI syntax is:

```
encapsulation aal5snap-bridged
```

## 2. VC-MUX

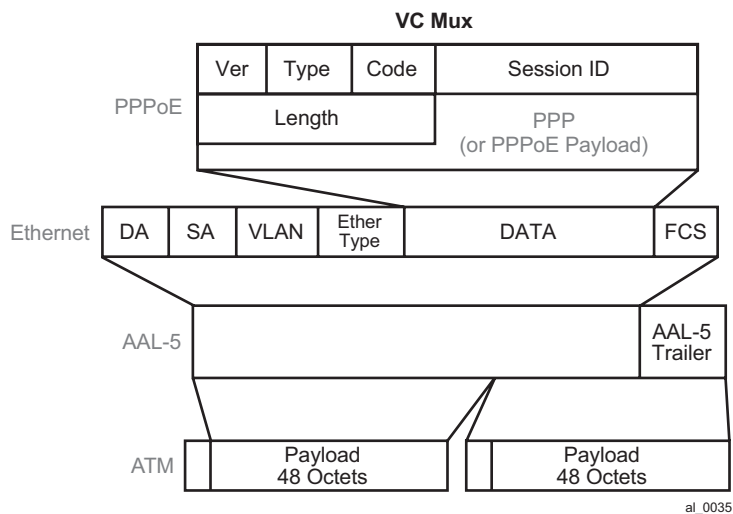
In the VC-MUX mode there is no additional (LLC/SNAP) header used for protocol multiplexing. Instead, the VC endpoints must agree before hand on the payload type that they will transport. For PVCs this is done during the provisioning phase on each side of the connection. For example:

```
encapsulation aal5mux-bridged-eth-nofcs
```

`aal5mux-bridged-eth-nofcs` tells each end of the VC in advance that the payload inside of the AAL5 frame is an Ethernet frame. In this case, it accepts the frame and treat it as an Ethernet frame inside AAL5. The EtherType within the frame must be set to 0x8863 (PPPoE Discovey Phase) or 0x8864 (PPPoE Session Phase).

The '-nofcs' portion indicates that the FCS is not supported on those Ethernet frames.

PPPoEoA encapsulation is shown in [Figure 67](#).



**Figure 67: PPPoA AAL5MUX Encapsulation**

Encapsulation Summary

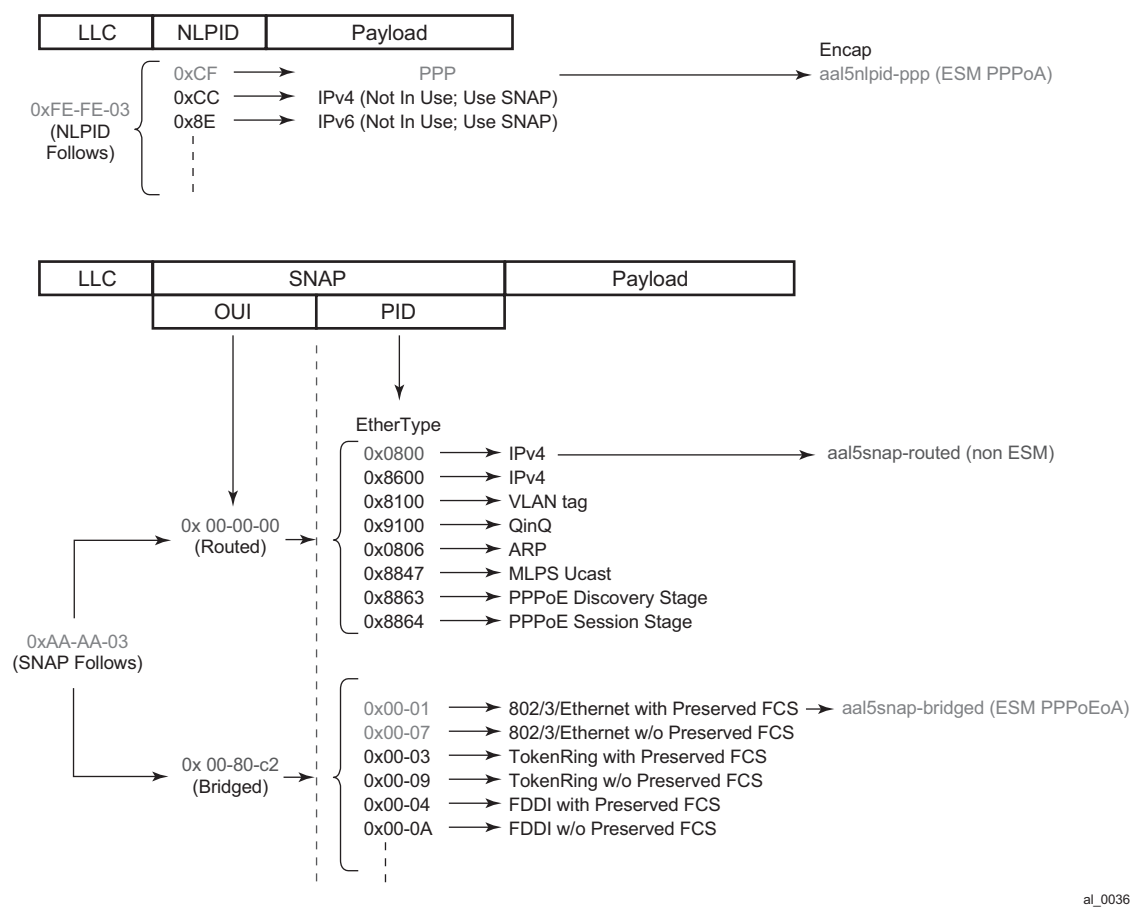


Figure 68: LLC/SNAP Encapsulation

There is a 0-2Byte additional padding (optional) in the SNAP Bridged encapsulation header that is not shown in Figure 68. This padding, according to RFC 2684, is necessary to align the info field (payload of the Layer 2 encapsulated frame) on a 4-Byte boundary.

At the end the following types of ATM encapsulation are supported on 7x50 in ESM:

aal5snap-bridged

This encapsulation type is used for PPPoEoA encapsulation with FCS or without FCS on ingress. On egress only frames without MAC FCS can be sent. PPPoE session type will be determined based on the EtherType in the Ethernet frame.

`aal5mux-bridged-eth-nofcs`

This encapsulation type is used for mux PPPoEoA sessions without MAC FCS.

`aal5nlpid-ppp`

This encapsulation type is used for LLC/NLPID PPPoA encapsulated packets.

`aal5mux-ppp`

This type encapsulation is used for PPPoA traffic without LLC/SNAP header (VC-MUX).

`aal5auto` (new command)

This encapsulation type is supported in ESM and it is used for auto detecting the encapsulation type. This is also called autosensing.

All hosts for the same subscriber will use the same encapsulation type.

---

## Concurrent Support for Different Service Types on the Same Port

An ATM port on 7x50 can concurrently support various service types. One service type can be mapped only to one VC. This would normally be the case if there is an aggregation network in front of 7x50. A Service Provider could run a variety of services over this aggregation network (ATM switches) connected on the same physical port within the 7x50.

For example, the following services can be run on the same physical port:

- PPPoA/PPPoEoA sessions connecting residential/business customers through a DSLAM to 7x50
  - Providing access for xPIPE services
  - Plain aal5snap-ip or vc-mux-ip PVC for Internet access, etc.
- 

## Restrictions in Scaled ATM MDA Mode

**Note:** ATM concatenation mode for Apipe is not supported in the 16K VC mode.

In the concatenated mode, cells are delayed so that they can be concatenated and delivered in a single packet over to pseudowires to the other side. Without concatenation, each cell is transported individually. The implication is that each cell is individually encapsulated into Eth/MPLS which results in wasted bandwidth on the link.

Support for the concatenated ATM pseudowires is not removed from the CLI in the 16K mode of operation:

```
configure
  service apipe <id> [vc-type <cell-type>]
    spoke-sdp <sdp:pw> cell-concatenation
      [no] aal5-frame-aware
      [no] clp-change
      [no] max-cells
      [no] max-delay
```

Instead:

- Adding an ATM port or a connection profile on an MDA in 16-VC mode to an APIPE is disabled if the vc-type is set to atm-cell AND cell-concatenation is enabled.

```
A:BNG>config>service>apipe# sap x/y/z:cp.w create
```

```
MINOR: SVCMMGR #2603 Cell-concatenation is not allowed on 16k VC-mode ATM MDAs
```

- cell-concatenation on an APIPE of the vc-type atm-cell is disabled if it already contains an ATM port or a connection profile.

```
A:BNG>config>service>apipe>spoke-sdp>cell-concat# max-delay X
```

```
MINOR: SVCMMGR #2603 Cell-concatenation is not allowed on 16k VC-mode ATM MDAs
```

Note that regular VPI/VCI SAPs (sap:vpi/vci) are not allowed to be configured on an Apipe of vc-type atm-cell.

Cell-concatenation is supported on Apipe services with a VC on an ATM MDA in the 8-VC mode.

AAL5 SDU mode is continued to be supported.

## QoS Implementation

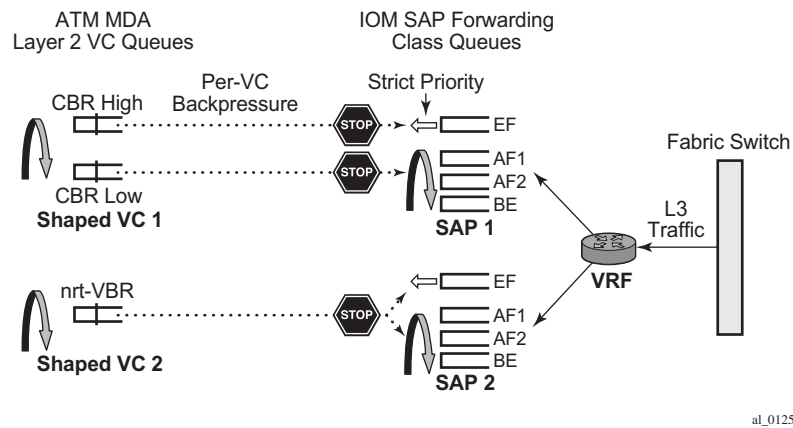
In addition to our system QoS that is provided in the “Q” chip on IOM, the ATM MDA offers QoS capabilities at the ATM cell level. In the context of this document, the system QoS is referred to as ‘IOM QoS’ and to ATM MDA provided QoS simply as ‘ATM QoS’.

Both, IOM QoS as well as ATM QoS defined by Traffic Descriptors working at the cell level, play a role in the overall QoS for the SAP (or virtual circuit). ATM QoS defines rates of each VC stream and defines the behavior under port congestion. IOM QoS defines bandwidth allotment and the scheduling scheme for each service within a VC stream.



In general, both MDAs, ATM and ASAP, support four traffic categories:

- CBR
- Rt-VBR
- Nrt-VBR
- UBR



**Figure 69: Scheduling on ATM MDA**

Each ATM traffic category is defined by a set of parameters, such as PIR, SIR, MIR, MBS and CDTV.

Currently, policing is supported only on ingress for CBR and VBR traffic classes. CBR traffic class police at PIR, while rt/nrt-VBR police at SIR.

Shaping is supported only on egress for CBR and rt/nrt-VBR traffic classes. CBR shapes traffic at PIR, while rt/nrt-VBR shape traffic at SIR. Egress shaping can be disabled only for nrt-VBR traffic class.

Scheduling at the ATM layer is shown in [Figure 69](#). Shaped CBR and VBR traffic classes have two queues, an HP and an LP queue. Packets from IOM are marked according to scheduling priority of the Forwarding Class (expedited or best-effort) from which they were sent and are accepted into the ATM VC queue (HP | LP) accordingly. For example, at the IOM level packets from an expedited FC (queue) are marked as HP, and the packets from a best-effort FC (queue) are marked as LP. When these packets arrive to MDA, they will be admitted into appropriate queues, pending VC (or port) buffer availability.

Non-shaped VBR and UBR have only one queue at the MDA level.


In terms of the ATM QoS scheduling, CBR has the highest scheduling priority (strict priority) followed by rt-VBR (also strict priority). The remaining traffic classes (non-shaped VBR and UBR) are serviced in WRR fashion, the weight being their configured SIR (non shaped nrt-VBR) or MIR (UBR) rate.

There are several areas of QoS that are addressed in relation to integration of IOM QoS and ATM QoS:

- Association between the subscriber and ATM VC traffic descriptor.
- Rate adjustments between Layer 3 and ATM QoS rates due to difference in frame/cell sizes on which they operate.
- Per VP shaping.

## Association Between the Subscriber and ATM VC Traffic Descriptor (QoS)

Each PPPoA PVC is of a certain traffic class – CBR, rt/nrt-VBR or UBR/UBR+MIR which along with other parameters is defined in the ATM traffic descriptor profile. For ATM service categories and traffic descriptors. There should be a uniform mapping between service offered and L3 and ATM QoS.



**Figure 70: ATM Traffic Descriptor Association with Subscriber**

The ATM traffic descriptor (atm-td) is applied to a VC under the SAP CLI hierarchy. In this fashion, MDA related QoS (ATM QoS) is referenced outside of the subscriber context (SUB/SLA-profiles).

The following describes the operability:

atm-td for the VC is applied under the SAP:

```
configure
  services ies/vprn
    subscriber-interface <sub-if-name>
      group-interface <grp-if-name>
        sap <sap-id>
          atm
            ingress
              traffic-desc <id>
            egress
              traffic-desc <id>
```

and for MSAP:

```
con figure
  subscriber-management
    msap-policy <name>
      atm
        ingress
          traffic-desc <id>
        egress
          traffic-desc <id>
```

msap-policy is then invoked via LUDB, RADIUS or the default-msap-policy under the capture SAP.

This allows 7x50 to have preconfigured MSAP policies, each corresponding to a specific VC type with its own traffic-class parameters (CBR/rt-nrt-VBR/UBR.). Each subscriber with corresponding hosts is then associated with a msap-policy that determines the VC type.

A more flexible way to go about this would be to allow the subscriber to reference the atm-td directly by the sub-host at the host creation time, independently of the msap-policy. This is supported in the following manner:

- atm-td is referenced via a RADIUS VSA in the Access-Accept message.
- There are two VSAs, one for ingress and one for egress atm-td:
  - alc-ingress-atm-td  
(ATTRIBUTE Alc-ATM-Ingress-TD-Profile 128 integer)
  - alc-egress-atm-td  
(ATTRIBUTE Alc-ATM-Egress-TD-Profile 129 integer)
- Only the first host of the subscriber can overwrite the atm-td that is defined under the msap-policy or under the static SAP.
- Consecutive hosts (second, third, etc.) of the same subscriber will not have any effect on the atm-td of the VC. For example, if the second host tries to overwrite the existing atm-td with a different atm-td, it will fail. Once the atm-td is set via the Access-Accept message for the first host, it cannot be changed as long as the subscriber is active in the system. This implies that all hosts of the same subscriber will have the same atm-td.


- If another hosts of the same subscriber initiates a session with a different atm-td name in the Access-Accept message, the host will be created but the atm-td for the VC will not be changed. A trap/syslog will be generated with a 'VC-parameter-mismatch:VSA ignored' info message.
- atm-td change will not be supported through CoA.

The following guidelines should be followed in configuring ATM traffic descriptors:

- ATM policing is the only function that can discard ATM cells on an ATM MDA
- Ingress:
  - ATM policing can only be enabled on ingress.
  - Policing is disabled by default.
  - Policing can only be enabled for the following traffic categories: CBT, rt-VBR and nrt-VBR.
  - CBR police at PIR. CBR is CLP transparent (it does not look at the CLP bit – aggregate traffic policing).
  - rt-VBR and nrt-VBR police at SIR and can operate either in CLP Significant (CLP marking) or CLP transparent mode (only relevant for Apipes and not for ESM).
  - In CLP significant mode with policing enabled, traffic is policed at the PIR value, but it can be marked with CLP 1 if it exceeds the SIR rate (only relevant for Apipes and not for ESM).
  - Ingress classification: the service category and CLP bits only affect Apipe traffic. If ATM traffic is terminated in the BNG, rely on the existing QoS classification implemented in the IOM.
- Egress:
  - Shaping is supported only on egress. Traffic shaping is supported for CBR, rt/nrt-VBR service categories.
  - ATM MDA is cell “lossless” – It sends backpressure to IOM which drops packets.

## Per VP Shaping

VCs traversing the same DSLAMs will typically use the same VPI. To prevent overrunning DSLAM capacity (intermediate destinations) in case an aggregation network is in place (between BNG and a DSLAM), per VP shaping will be implemented.



al\_0038

**Figure 71: VP Shaper**

The shaping rate per VP and the VP service type is provisioned manually via a traffic descriptor on a per port level. The cli syntax is:

```
config
  port <port-id>
    sonet-sdh
    path
    atm
    vp <vpi> egress-traffic-desc <atm-td-profile-id>
```

Where the vpi is a VPI identifier and the egress-traffic-desc is the traffic descriptor id. Only traffic descriptors with service-category of cbr, rt-vbr and nrt-vbr can be used in VP Shapers. However, only the rate in the traffic descriptor can be changed on-the-fly, and not the service category (nrtVBR, rtVBR, CBR). A VP shaper can be added to or removed from active VCs.

CBR VP Shaper shapes cells at the exact PCR rate. There is no burst concept in CBR shaping. Excessive traffic is back-pressured towards the IOM (Q-chip).

The IOM will never send more frames to the ATM MDA than the MDA cannot buffer. This is implemented through a combination of software and hardware backpressure mechanisms. This backpressure mechanism utilizes a combination of hardware and software. Software backpressure aims to have around 100ms of traffic queued against a VC based on its configured shaping/scheduling rate, but being a software mechanism that is only a guideline. As ATM MDA detects more traffic that it can accept, a hardware backpressure is exerted.

A rt/nrt-VBR type VP Shaper has three parameters associated with it: PCR (peak cell rate), SCR (sustained cell rate) and MBS (maximum burst size at a peak rate). As long as there is enough MBR credit, traffic will be shaped at the PCR rate. Once all MBR credit (burst) is exhausted, traffic will be shaped at the SCR rate. Bursting above the SCR is configurable via the MBS parameter.

In both cases cells will be spaced at  $1/PCR$  or  $1/SCR$  as perfectly as possible with minimum jitter.

VP shaping is supported only when the ATM MDA is in max16k-vc mode. The maximum number of VP shapers per MDA to 128.

The maximum number of VCs that can feed into a single VP shaper is 16K. This includes the sum of all VC-ranges on the VP plus any statically configured VC on that VP.

VCs within the VP tunnel is serviced by a single scheduler assigned to each VP tunnel. The ATM VP shaper will condition the aggregate traffic for all ATM VCs within the VP tunnel. VCs within the shaped VP tunnel are degraded from the originally assigned service category to a common UBR service category (default traffic descriptor). If the VP shaper is removed from the VCs, the VCs will be reverted to their original service category. Scheduling between VCs will be WRR based with a weight parameter that is explicitly configured. The weights assigned to VCs within the VP tunnel are in range 1-255. By default, VCs are assigned a priority based on the originally assigned service category:

- VC degraded from CBR = weight 10
- VC degraded from rt-VBR = weight 7
- VC degraded from nrt-VBR = weight 5
- VC degraded from UBR+ = weight 2
- VC degraded from UBR = weight 1

The weight parameter is user configurable under the traffic descriptor hierarchy.

```
configure
  qos
    atm-td-profile <td-profile-id> [create]
      weight <weight>
```

```
<weight> : 1-255.
```

If weight is not specifically configured, the defaults are taken as described above.

The explicitly configured weight parameter is honored only on ATM MDA in the max16k-vc mode. On all other ATM capable MDAs (ASAP or ATM MDA in max8K-VC mode), the weight parameter is ignored.

Note that in the current ATM implementation there is already a WRR scheme in place based on internally calculated weights. This WRR scheme is used to service traffic from the VC queues of equal priority (where there are two queues per VC - a HiPrio and a LowPrio queue). Weights are assigned to VCs automatically based on the rate of the VC,

---

## ATM/IOM QoS Integration

There are major differences between the QoS mode of operation at the ATM MDA level and the IOM level.

ATM QoS operates on fixed size cells that contain additional transport overhead. In addition, ATM shaping is very accurate so that traffic is paced into the ATM network with nodes that are sensitive to bursts (ingress policing). Buffering and service differentiation (number of queues) at the ATM layer is not as flexible as it is on the IOM level.

On the other hand, HQoS in the IOM is less accurate and less responsive to sudden traffic fluctuations (bursts). It operates on Layer 2 frame lengths. Bursts of traffic are usually let into the network more freely than the ATM network would like to accept.

To combine the extensibility of IOM HQoS with the stringent ATM QoS requirements, the two modes of operation are integrated.

The congestion in the ATM network is treated by using extensive IOM HQoS.

When the VC ATM queue becomes congested, it exerts backpressure to the subscriber queues in the IOM on the corresponding VC.

To avoid the condition where a VC becomes overly oversubscribed and excessive in exerting the backpressure, our HQoS in IOM has to be proactive and has to be able to detect congestion on the IOM level based on the ATM VC configured rate or the VP configured rate. IOM HQoS must deal with this congestion before it actually becomes the problem in the ATM layer. Occasional and short-lived backpressure from VCs, but persistent QoS backpressure and congestion at the ATM layer is avoided. A prerequisite for treating ATM congestion at the IOM level is to adjust the frame size in HQoS calculations (on all levels - including queues and aggregate rate limits) so that it reflects the ATM overhead associated with ATM transport (AAL5 encap, cellification).



## Intermediate Node Rate Limit/Shaper

In many cases it is desired or even necessary to shape traffic per DSLAM. In the ATM network, a DSLAM corresponds to a VP (Virtual Path) which is at the ATM cell level identified by a VPI (Virtual Path Identifier). The VPI/VCI pair represents the addressing mechanism in the ATM network.

Another level of hierarchy is introduced into our HQoS model - an aggregate rate limit in the IOM that will correspond to the ATM VP shaper. The purpose of this new construct in the IOM is to help detect congestion at the IOM level before it becomes a severe problem at the ATM layer.

As a final solution, our HQoS hierarchy in IOM is a 4 tier hierarchy consisting of:

- subscriber queues that are parented to the port scheduler
- egress agg-rate-limit that represents the egress aggregate subscriber rate
- tier-2 aggregate-rate-limit that represents the ATM VP rate limiter (related to DSLAM). This is a pure rate limiter and not a scheduler (it does not receive any scheduling opportunity).
- port-scheduler to which are subscriber queues parented. The port-scheduler delegates bandwidth to its children based on the priority levels associated with the children queues.

The four tiered hierarchy is needed to deal with:

- port congestion. The port bandwidth can be overbooked the desired service levels can still be ensured.
- ATM VP congestion. An ATM VP can be effectively protected from prolonged congestion that would result into significant backpressure to the IOM queues.
- subscriber congestion. Within the subscriber, bandwidth is managed within the subscriber bandwidth limit.

The four-tiered hierarchy looks like [Figure 72](#).




Figure 72: Tier HQoS

The key point in such HQoS model is that the port-scheduler delegates its available bandwidth to the subscriber queues directly according to the queue priority on a configured level. Higher priority queue are served over all subscribers before any lower priority queues, up to the limits imposed by the tier 1 and tier 2 aggregate rate limits.

## Provisioning Aspects

The tier 2 aggregate rate limit that corresponds to the VP shaper on the IOM level is provisioned in the context of a vport. The vport is a container that is configured directly under the port and it can contain either an aggregate rate limit or another port scheduling policy (port scheduler). These two constructs (vport aggregate and V-port port scheduling policy) are mutually exclusive. In addition, if a V-port port scheduling policy is configured instead of the V-port aggregate, then the port scheduler cannot be used (currently two port schedulers applied under the same port cannot be used in a parent/children relationship).

The V-port aggregate have to have the agg-rate-limit defined explicitly. In other words, the rate can NOT be implicitly inherited by configuration from the VP shaper traffic descriptor.


Example:

```
config
  port <port-id>
    sonet-sdh
      path
        egress-scheduler-policy <port-scheduler-policy-name>
      access
        egress
          vport <vport-name>
            description <description-string>
            host-match dest <dest-string>
            agg-rate-limit <agg-rate>
            port-scheduler-policy <port-scheduler-policy-name>

  atm
    vp <vpi> egress-traffic-desc <atm-td-profile-id>
```

The association between the vport aggregate and the subscriber host is done in three ways:

- RADIUS attribute — a dest-string VSA that is matched against the string defined under the corresponding vport. This dest string VSA is returned via RADIUS at the subhost instantiation time.
- LUDB – similar to the RADIUS method. The dest-string comes from the LUDB instead from RADIUS.
- VP identifier (VPI) - The VPI is known to the CPM from the beginning of the session initiation process – raw (unknown encapsulation) packets and passed to the CPM. These raw packets contain VPI,VCI identifiers. The vport container for the subscriber is referenced implicitly via the VPI in the following fashion ([Figure 73](#)):
  - CPM determines the VPI from the first packet for the session.
  - ATM VP Shaper mapping: The VPI is used to make the subscriber association with the VP Shaper which is defined under the VP Shaping node: **port>sonet-sdh>path>atm>vp**. The VP Shaping node name must be the VPI number.
  - Vport mapping: The vport-name in the **port>sonet-sdh>path>access>egress>vport <vport-name>** hierarchy is matched against the VPI number.



**Figure 73: VPI Based V-Port <-> Subscriber Association**

The association method (automatic via VPI or based on RADIUS/LUDB) between the subscriber host and the vport is defined under the SAP (or MSAP) where the subscriber resides. In most cases subscriber management is used with MSAPs.

This is the syntax:

```
configure
subscriber-mgmt
msap-policy <name>
sub-sla-mgmt
def-inter-dest-id string <inter-dest-string>
def-inter-dest-id {use-top-q | use-vpi}

configure
services ies/vprn
subscriber-interface <sub-if-name>
group-interface <grp-if-name>
sap <sap-id>
sub-sla-mgmt
def-inter-dest-id string <inter-dest-string>
def-inter-dest-id {use-top-q | use-vpi}
```

The **def-inter-dest-id** stand for a 'default inter-destination identifier'.

If the use-vpi method is used and the VPI derived from the incoming traffic points to a non-existing VP container, the association between the subhost and the V-port container will fail and a message/trap is logged. This however will not prevent the creation of the subhost that can be parented to the port-scheduler.

If the use-vpi is used on an Ethernet port where this parameter is not applicable, the parameter is ignored and defaulted to 'string'. A message/trap is logged.

Note that if the vport contains an aggregate-rate-limit, then there is no need for the indication of a vport construct in the sub-profile or sla-profile of the subscriber. On the contrary, in case that the vport contains a port-scheduling-policy, the sla-profile template must contain the indication that the subscriber is tied to a vport (**configure>subscriber-mgmt>sla-profile>egress>qos sap-egress-qos-policy-id vport-scheduler.**)

## HQoS Combinations

These are the possible HQoS combinations that are supported:

- port-scheduler assigned to path, agg-rate-limit assigned to vport, agg-rate-limit assigned to subscriber.
- port-scheduler assigned to path, no vport, agg-rate-limit or scheduler-policy assigned to subscriber.
- no port-scheduler assigned to path, port-scheduler assigned to V-port, agg-rate-limit or scheduler-policy assigned to subscriber.
- no port-scheduler assigned to path, no vport, scheduler-policy assigned to subscriber.

In case that the vport contains the agg-rate-limit, any subscriber host queue that is parented to a virtual scheduler will not be rate-limited by the vport aggregate rate. The queue will compete for bandwidth directly on the port's port scheduler, at the priority level and weighted scheduler group the virtual scheduler is port-parented to. If the virtual scheduler is not port-parented or if there is no port scheduler policy on the port, the host queue will be orphaned and will compete for bandwidth directly based on its own PIR and CIR parameters.

## ATM Rate Adjustment

The difference in cell/frame overhead on the ATM level and IOM Level (Layer 2) lead to inconsistent behavior in integrated QoS. For example, IOM based QoS operates on Layer 2 frames (PPP header is included in rate calculations). On the other hand, ATM QoS operates on ATM 53byte cells. Each PPP frame with AAL5 overhead is segmented into 48-byte chunk cells prepended by a 5 byte ATM header. Assuming that ATM QoS operates on 53-octet cells, a significant rate discrepancy might arise from this difference.

For example, consider two flows:

Flow 1, 1000pps, PPP+IP packet size 190Bytes => L2 (IOM) QoS rate is 1520 kbps

Flow 2, 1000pps, PPP+IP Packet size 232Bytes => L2 (IOM) QoS rate is 1856 kbps

Assuming that VC MUX encap is used, and that AAL5 trailer (8 octets + padding) is used:

Flow 1 = AAL5 length is  $190 + 8 + 48B_{\text{boundary\_padding}} = 240\text{Bytes} \Rightarrow 5 \text{ ATM cells}$

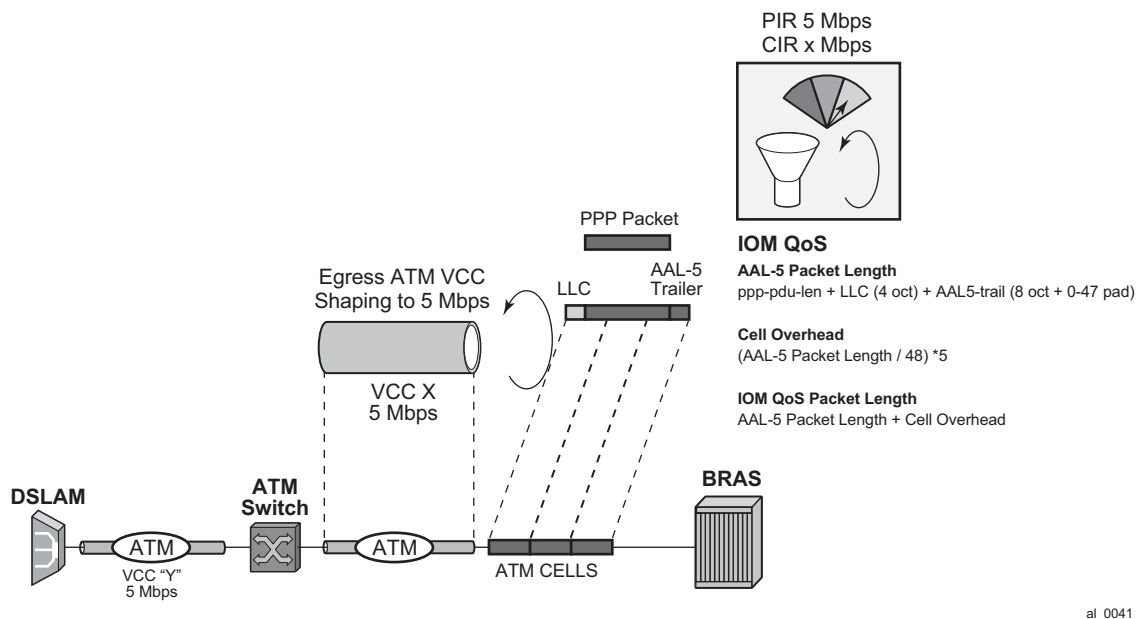
Flow 2 = AAL5 length is  $232 + 8 + 48B_{\text{boundary\_padding}} = 240\text{Bytes} \Rightarrow 5 \text{ ATM cells}$

Each cell has a 5 byte header which translates into a 2120 kbps rate for each flow.

As a result, more bandwidth is required at the ATM layer than the QoS on the IOM level has budgeted. Moreover, the two flows appear to be of the same rate at the ATM layer, even though the IP rate of flow 1 is ~20% less than the flow 2.

From the example above, Flow 1 consumes 40% more bandwidth configured at the ATM layer than it is given in IOM QoS. The rate adjustment depends on the packet size which additionally complicates conversion. This issue is more pronounced with smaller packet sizes.

An example is shown in [Figure 74](#).



**Figure 74: QoS Adjustment**

Depending on the session encapsulation (VC-MUX, LLC/SNAP) the packet length on which IOM QoS operates is adjusted.

The rates are adjusted on a per queue level, per subscriber egress agg-rate-limit level, per V-port aggregate level and on per port-scheduler level. Because the ATM termination points are on the

BNG, there is a direct view of the encapsulation. The encapsulation information is supplied to the forwarding plane via the control plane.

The rate adjustment are examined in the following commands:

- **frame-based-accounting** under the qos>scheduler-policy hierarchy
- **queue-frame-based-accounting** under the sub-profile>agg-rate-limit hierarchy
- **avg-frame-overhead** under the queue hierarchy
- **encap-offset** under the sub-profile hierarchy

In PPPoA/PPPoEoA scenario there is no last mile rate adjustment (due to the difference in encapsulation) performed. It is assumed that the encapsulation in the last mile and in the intermediate mile is the same. In other words, it is assumed that the DSLAM does not add/change any encapsulation but instead it only acts as a VPI/VCI cross-connect. If the last mile encapsulation is PPPoA, then the intermediate encapsulation is considered to be PPPoA as well. Similar is valid for PPPoEoA encapsulation.

There are two configuration scenarios possible:

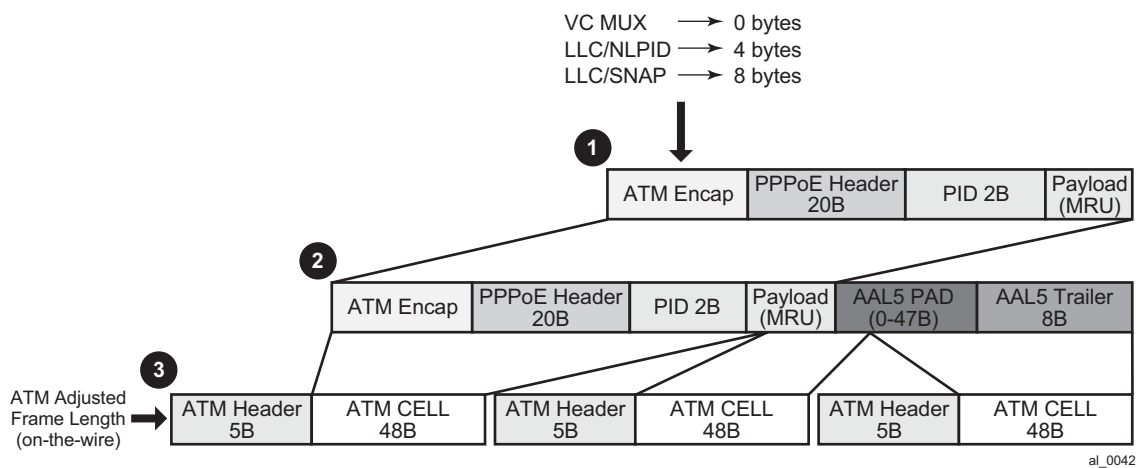
1. The **encap-offset** command in the sub-profile is configured. This command overwrites any other command related to rate conversion that might be configured (**frame-based-accounting**, **queue-frame-based-accounting**, **avg-frame-overhead**, or **avg-frame-size**). The *encap-offset* command forces dynamic wire rates calculation in the intermediate mile (directly connected ports) on all levels in the QoS hierarchy. The wire overhead in the intermediate mile takes into account the length of the fixed ATM encapsulation, the variable length of AAL5 encapsulation (including AAL5 48bit boundary padding) and the ATM cellification overhead. The queue stats are also wire based. All calculations are performed in the data plane using the actual packet size. In other words, this command will ensure that the rates on the queue level and the subscriber aggregate level (either through virtual schedulers or egress aggregate-rate limits) are wire based. Port-scheduler and V-port rates are already by default wire based rates and this cannot be changed.
2. The **encap-offset** command is not configured by default. In this case the other rate conversion related commands are in effect (**frame-based-accounting**, **queue-frame-based-accounting**, **avg-frame-overhead**, or **avg-frame-size**). The behavior is the following:
  - Wire rates in the intermediate mile (directly connected ports) are based on the **avg-frame-overhead** command which is provisioned via CLI. If avg-frame-overhead is not provisioned via CLI, by default it is assumed to be 0[%] and the wire rates effectively become data rates (IP payload + IP header + PPP(oE) header + fixed ATM encapsulation).
  - Queue stats (used in accounting) are always ‘data’ stats. This includes the IP Payload + IP header + PPP(oE) header + fixed ATM encapsulation.
  - *agg-rate-limit* (subscriber or vport) rates are always wire rates (as defined in the first bullet – based on the avg-frame-overhead).
  - Rates in the *port-scheduler-policy* (vport or physical port) are always wire based (rates (as defined in the first bullet – based on the avg-frame-overhead).
  - The **frame-based-accounting** command under the scheduler-policy will affect rate calculation for virtual schedulers and queues:



- If this option is configured, the virtual scheduler and queue rates will be wire based.
  - if this option is NOT configured, the virtual scheduler and queue rates will be data rates
- *queue-frame-based-accounting* configuration option under the subscriber agg-rate-limit command (in sub-profile) will affect rate calculations for queues. If this command is configured, the queue rates will be wire rates, otherwise they will be data rates

Avg-frame-size command in PPPoA/PPPoEoA is ignored.

Currently queue rates and subscriber virtual scheduler rates are allowed to be either data rates (one in [Figure 75](#)) or *on-the-wire-rates* (three in [Figure 75](#)). *Port-scheduler* rates, vport rates and the subscriber *agg-rate-limit* (in sub-profile) are always on *on-the-wire* rates.



**Figure 75: ATM Wire Overhead**


## Subscriber Instantiation Use Cases

Note that multiple subscribers per single VC are not supported. However, multiple VCs per subscriber are supported.

The following displays examples of how subscriber hosts could be instantiated in PPPoA/PPPoEoA environment.

### Case 1

- One host per subscriber.
- One ATM VC per subscriber.
- Authentication is done via RADIUS or LUDB.
- Authentication methods: PAP/CHAP or PADI for PPPoE.
- If there is no RADIUS/LUDB authentication, the subscriber-id should be by default the SAP (with VPI:VCI) with default sub-strings. For MSAP this would be configured under the msap-policy – all subs would have the same service in this case (SLA/SUB profiles).




al\_0043

Figure 76: Subhost per VC

**Case 2**

- Multiple hosts per subscriber
- Single VC per host
- Multiple VCs per subscriber
- PPPoA or PPPoEoA
- Subhosts for the same subscriber can authenticate using a unique username or the same username

**Figure 77: Multiple VCs per Subscriber**

### Case 3

- Multiple hosts per subscriber
- Single VC
- PPPoEoA only
- Per host SLA-PROFILE instantiation

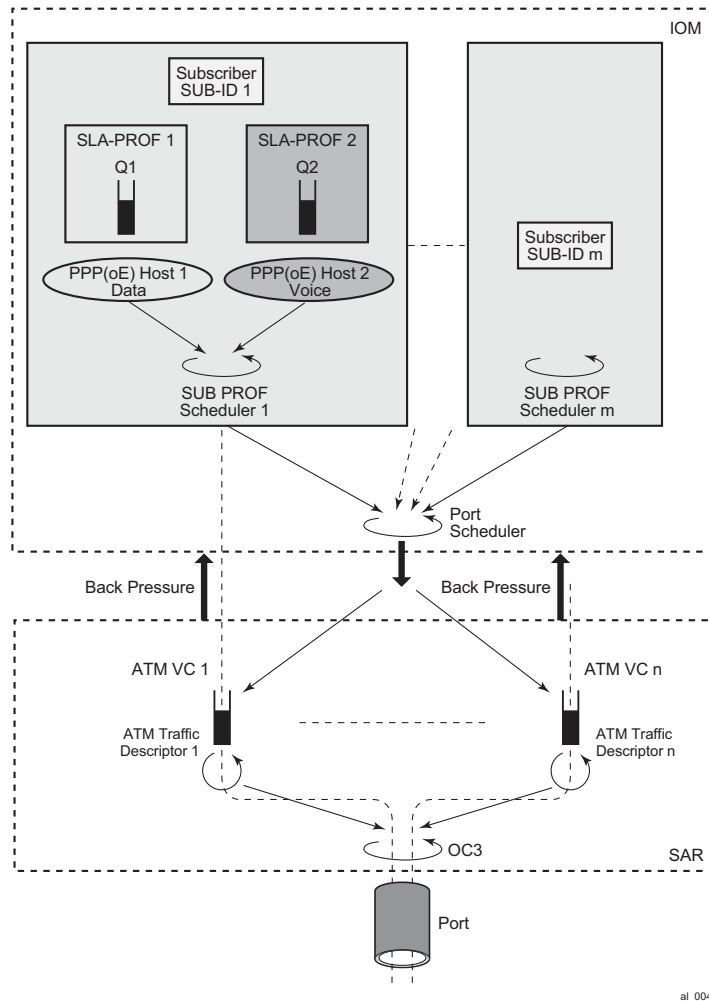
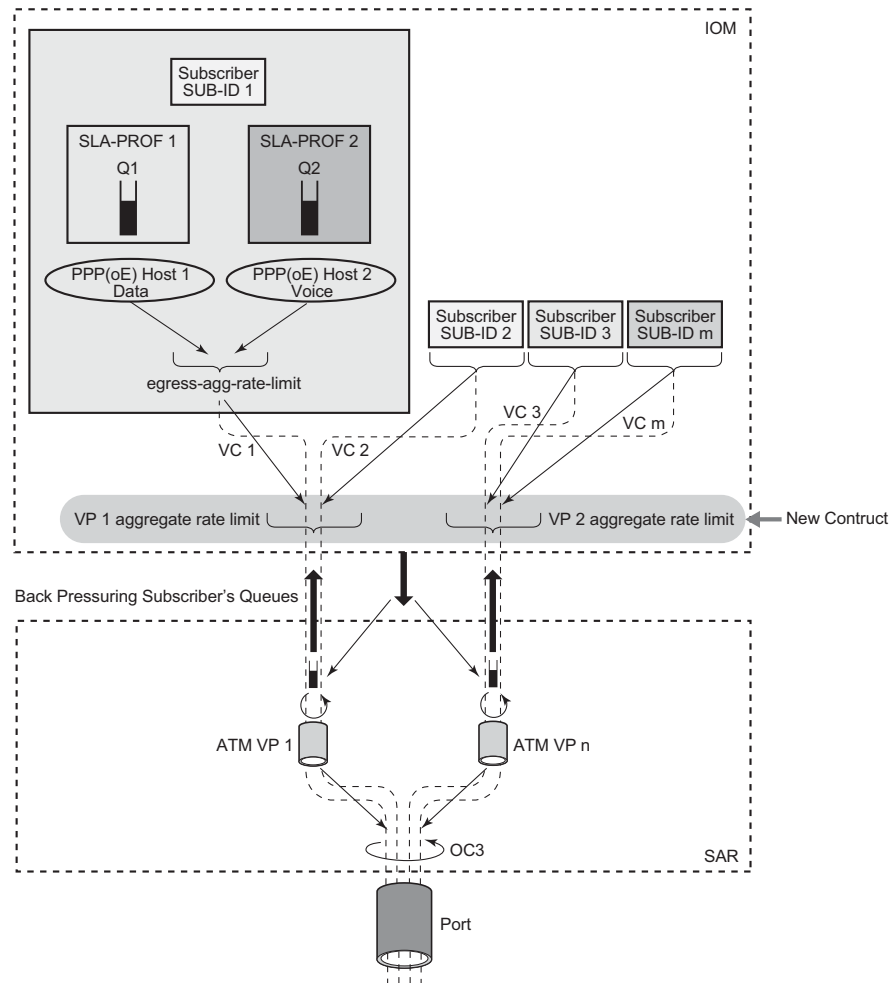


Figure 78: Multiple Hosts per Subscriber, Single VC

**Case 4**

- VP Shaping



al\_0046

**Figure 79: VP Shaping**

## Authentication

Authentication for PPPoEoA is the same as in PPPoE:

- based on PADI
- based on PAP/CHAP
- No authentication

Authentication in PPPoA is based on:

- PAP/CHAP
- No Authentication — There is a small percentage of cases where this might be required.

An example of no authentication configuration for PPPoEoA would be:

```
configure
  service ies/vprn
    subscriber-interface <sub-if-name>
      group-interface <grp-if-name>
        dhcp
          server <ip-address>
            client-application pppoe => make PPPoE session the client of DHCP
server
  local-dhcp-server <dhcp-server-name>
    user-db <ludb-name> => allow DHCP server to query LUDB
  subscriber-management
    local-user-db <ludb-name>
      pppoe
        host <host-name>
          host-identification [circuit-id | mac | remote-id | service-name |
username]
          options
            address
            identification-strings
```

In the case that there is no PAP/CHAP authentication, a PPPoE host can be identified by a MAC address, circuit/remote-id (inserted by the BNG) or service-name in PPPoE.

## LUDB Access via Capture SAP

Access to LUDB via a capture SAP is enabled for this feature:

```
configure
  service vpls <id>
    sap <sap-id> capture-sap
      pppoe-user-db <ludb-name>
      ppp-user-db <ludb-name>
```

Note that if the authentication-policy (RADIUS authentication) is specified under the capture SAP, it (the RADIUS authentication) will take precedence over LUDB.

## Encapsulation Autosensing

As previously discussed, these four static encapsulation types will be supported:

1. aal5nlpid-ppp (LLC/SNAP for PPPoA)
2. aal5snap-bridged (LLC/SNAP for PPPoEoA)
3. aal5mux-ppp (mux PPPoA)
4. aal5-mux-bridged-eth-nofcs (mux PPPoEoA)

These four types of encapsulation are supported for ATM MSAPs as well as for fixed configuration ATM SAPs.

In addition, the LLC/SNAP encapsulation type can be autosensed. This is called autosensing. The keyword for autosensing is **aal5auto**. An option is given to provision encapsulation statically, if needed.

```
sap x/y/z:* capture-sap
atm
encapsulation aal5auto | aal5mux-ppp | aal5nlpid-ppp | aal5mux-pppoe | aal5snap-bridged-
eth-nofcs

sap x/y/z:w/z
atm
encapsulation aal5auto | aal5mux-ppp | aal5nlpid-ppp | aal5mux-pppoe | aal5snap-bridged-
eth-nofcs
```

The aforementioned encapsulation options (including autosensing) is visible only under the SAP hierarchy on group-interfaces and on capture SAPs in VPLS. This is allowed only in 16K-VP mode ATM MDA.

---

## SAP Autoprovisioning

In order to simplify the provisioning of the subscriber access ports in the ESM context, a concept similar to managed SAP(MSAP) on Ethernet is introduced. In Ethernet MSAP, an ingress access SAP is automatically created upon receipt of the first (VLAN) tagged packet from the customer side (pending the authentication process).

In our ESM over PPPoA/PPPoEoA case, a number of PVCs is pre-provisioned that will initially be only in a provisioned (or listening, passive) state. This is sometimes referred as a bulk configuration of VC ranges. Once the initial ESM processing in the CPM is completed (for example, the user is authenticated), the ATM SAP is created in the appropriate context (VRF or GRT) and the ATM VC is activated. ATM VC activation means that the ATM VC is associated with a SAP.



The ratio between the maximum number of provisioned VCs vs the maximum number of active VCs on an MDA supporting PPPoA/PPPoEoA is 2:1. This amounts to 32K ATM VCs in the listening state. Out of the total 32K ATM VCs, 16K ATM VCs can be active simultaneously.

Obviously, there are some differences between the Ethernet MSAP processing and the ATM MSAP processing. On Ethernet there is not need to pre-configure VCs, while on ATM this is necessary due to complexity of ATM layer comparing to Ethernet.

Our current Ethernet based MSAP is configured under VPLS. The same approach is adopted for ATM capture SAP.

Since a combination of ESM over PPPoA/PPPoEoA VCs and other non PPPoA/PPPoEoA VCs on a single physical port is supported, ranges of VCs that will be supported in autoconfiguration are defined. ATM VCs outside this range are available for manual creation. Multiple ranges are necessary in order to address non-contiguous sets of VPI/VCIs.

The following is the CLI syntax:

```
configure
  service vpls <id> customer <customer-id> [create]
    sap x/y/x:*/* capture-sap [create]
      atm
        vc-range <num> vpi-range <vpi-range> vci-range <vci-range>
```

Note that the capture SAP must be configured as ‘\*/\*’ in place of the VPI/VCI identifiers. Any other combination for the VPI/VCI identifiers is not allowed.

Up to 5 ranges are allowed per capture SAP. VCs configured in this way can carry a single PPP session per VC or multiple PPPoE sessions per VC. Ranges are not allowed unless the ATM MDA is in the 16K-VCs mode.

The total number of VCs that can be fed into a VP is 16K. This includes all VC ranges associated with the VP plus any statically configured VC on the VP.

## PPP Nodes and ppp-policy

Differences in operation between PPP and PPPoE warrant creation of a new ppp node under the *group-interface* in the CLI that will cover PPP aspects of operation. The existing pppoe node under the *group-interface* is preserved in the CLI. This allows referencing different *user-dbs* for authentication purposes and different session parameters defined in the *ppp-policy* for each session type (pppoe or plain ppp).

PPP node under the *group-interface* is used to cover PPPoA operation while PPPoE node is used to cover PPPoE and PPPoEoA operation. ATM in PPPoEoA is just a transport and as such does not carry any information relevant to PPP operation (like PADx does in PPPoE).

For dynamic SAPs (managed SAPs), the same PPP(oE) related structures are referenced under the capture SAP hierarchy. Under the capture SAP there are no ppp/pppoe nodes (like they are under the group-interface hierarchy). In order to differentiate between *ppp* and *pppoe* clients, a new ppp-policy command is introduced in addition to the *pppoe-policy* command. The *ppp-policy* under the

capture SAP is needed for the definition of session parameters before the *group-interface* (where normally session parameters are referenced) is determined.

Two commands for LUDb access are available under the same *capture-sap* hierarchy (*pppoe-user-db* and *ppp-user-db*).

The *ppp-policy* under the *subscr-mgmt* hierarchy contains PPP and PPPoE session parameters. PPP parameters are applicable to both session types (PPP and PPPoE) while PPPoE parameters are applicable only to PPPoE session type. The PPPoE parameters are ignored for PPP sessions.

---

## MTU Considerations

MRU configuration option negotiated during the LCP phase in PPPoA is based on the following command:

```
configure>subscr-mgmt>ppp-policy#  
    ppp-mtu
```

By default, this command is disabled and consequently will negotiate the default MRU of 1500B, as long as the ATM port's MTU can accommodate at least 1500B:

```
configure>port>sonnet-sdh>path#  
    mtu
```

The MRU option in PPPoA refers to the PPP packet length (PID+Information+Padding) that is ATM encapsulated.

## PPP(oE) Session Antispoofing

Antispoofing filters need to be in place in order to prevent the hijacking of a PPPoA/PPPoEoA session. For successful anti-spoofing, the following fields are accessible:

- Source IP address (PPPoA and PPPoEoA)
  - VPI/VCI pair which is equivalent to the SAP (PPPoA/PPPoEoA)
  - Source MAC Address (PPPoEoA only)
  - session-ID (PPPoEoA only)
1. For locally terminated PPPoEoA subscriber hosts, access to all fields are required (a, b, c and d). Antispoofing is defined under the following hierarchy:

```
configure
service <svc-name>
subscriber-interface <sub-if-name>
group-interface <grp-if-name>
    pppoe
    anti-spoof [mac-sid | mac-sid-ip]
```

This behavior matches our current PPPoE behavior. The default antispoofing option is set to mac-sid, which means that the incoming traffic is checked against the source MAC address and the session-ID. Antispoofing under this hierarchy cannot be disabled.

The source IP address can be added to the source MAC and session-ID (mac-sid-ip).

Note that the group-interface is the lowest granularity at which options can be enabled. In other words, these two options (mac+sid or mac+sid+ip) cannot be changed at the SAP level.

The VPI/VCI pair is always be checked against the incoming traffic, regardless of the configuration option as the VPI/VCI pair is an intrinsic part of the SAP (similar to VLAN tags on Ethernet).

In case that a subscriber host is a routed host (managed routes), the nh-mac antispoofing option must be enabled under the managed sap (msap-policy) or group-interface->sap level. Otherwise, managed routes for the host would NOT be installed. The nh-mac option forces a lookup of the incoming packet based on the mac+sid of the originally created host (CPE device). Note that only the group-interface->sap>anti-spoof and msap-policy hierarchies contain the nh-mac option, and NOT the pppoe->anti-spoof hierarchy.

For locally terminated PPPoA subscriber hosts, only access to source IP address is available (src MAC and session-ID fields are non-existent in PPP).

The default antispoofing is based on the VPI/VCI pair + IP. This cannot be changed by configuration, except when the session has managed routes or is a LAC session. In the case of managed routes (routed host) a lookup will be done based on the SAP (VPI/VCI) only. The antispoofing for PPPoA should be set to antispoof nh-mac, under msap-policy or grp-if>sap hierarchy even though PPPoA has no MAC address.

For LAC, the behavior is the following:

- For PPPoE traffic, antispoofing is always done based only on the SAP+mac+session-id.
- For PPP traffic, antispoofing is always based on VPI/VCI pair (SAP only). The IOM automatically determines whether IP based antispoofing can be done (e.g. no IP based antispoofing for LAC and managed routes).

There are two other nodes on which antispoofing can be configured:

```
configure
  service ies/vprn
    subscriber-interface <sub-if-name>
      group-interface <grp-if-name>
        sap <sap-id>
          anti-spoofing [ip|ip-mac|nh-mac]
```

A few points in regards to the above hierarchy:

- **ip** - This option can be only used in BSM mode. This is used only for IPoE.
- **ip-mac** – Lookup is performed based on the combination of the IP address and MAC address (incipient host). This is used for IPoE. For PPPoE, it will be overwritten by the configuration option under the group-interface-> ppp node (mac+sid or mac-sid+ip).
- **nh-mac** – Lookup is based on the MAC only for IPoE or mac+sid for PPPoE.


```
configure
  subscriber-managemnt
    msap-policy <msap-pol-name>
      ies-vprn-only-sap-parameters
        anti-spoof [ip-mac|nh-mac]
```

This is similar to the previous case (under the group-interface>sap hierarchy) with the exception that the pure IP option is not supported under the MSAP. The reason is that the IP option can only be used in BSM (under the SAP node) whereas MSAP can only be used in ESM (PPPoX).

Note that for IPoE, the entire control over antispoofing under the (M)SAP node, while for PPPoE, the anti-spoofing control is distributed between the (M)SAP node and the *group-interface->ppp* node.

## Multi-Chassis Synchronization

Figure 80 shows the configuration under which synchronization of subscriber management information is performed. As depicted, a single access node aggregating several subscriber lines is dual-homed to redundant-pair of nodes.



**Figure 80: Dual-Homing Configuration**

Enabling subscriber management features (whether basic subscriber-management (BSM) or enhanced subscriber management (ESM)) causes the node to create and maintain state information related to a given subscriber-host. This information is synchronized between redundant-pair nodes to secure non-stop service delivery in case of the switchover.

## Overview

The synchronization process provides the means to manage distributed database (the Multi-Chassis Synchronization (MCS) database), which contains the dynamic state information created on any of the nodes by any application using its services. The individual entries in the MCS database are always paired by peering-relation, sync-tag and application-id. At any time the given entry is related two the single redundant-pair objects (two SAPs on two different nodes) and hence stored in a local MCS database of the respective nodes.

Internally, peering-relation and sync-tag are translated into a port and encapsulation value identifying the object (SAP) that the given entry is associated with. The application-id then identifies the application which created the entry on one of the nodes. There are three basic operations that the application can perform on MCS database. The MCS database will always synchronize these operations with its respective peer for the given entry.

The following principles apply:

- **add-operation** — Any dynamic-state created in the application is pushed to the MCS database. MCS then creates and synchronizes with the corresponding peer provided (if configured). The application in the peer node is then notified as soon as the entry has been created. Similarly, the application in the local node (the node where the state has been created) is notified that entry has been synchronized (MCS is “in-sync” state). This operation will be also used to modify existing MCS database entry.
- **local-delete** — The MCS database entry is marked as no longer in use locally and this information is sent to the peer node. If the information is no longer used by applications on both nodes (the application in remote-node has already issued local-delete before), it is removed from database.
- **global-delete** — The MCS database entry is removed from both nodes and from the application in the remote node.

The choice of the operation in corresponding situation is driven by the application. The following general guidelines are observed:

- An event which leads to a dynamic-state deletion on a standby chassis will be handled as “local-delete”.
- An event which leads to a dynamic-state deletion on an active chassis will be handled as “global-delete”.
- An exception to above the rules is an explicit “clear” command which will be handled as “global-delete” regardless of where the command was executed.

As previously stated, the MCS process automatically synchronizes any database operation with the corresponding peer. During this time, the MCS process maintains state per peer indicating to the applications (and network operator) the current status, such as in-sync, synchronizing or sync\_down. These states are indicated by corresponding traps.

## Loss of Synchronization and Reconciliation

Each time the connection between the redundant pair nodes is (re)established the MCS database will be re-synchronized. There are several levels of connectivity loss which may have different effect on amount of data being lost. To prevent massive retransmissions when the synchronization connection experiences loss or excessive delay, the MCS process implementation will take provisions to ensure following:

- In the case of a reboot of one or both nodes or establishing the peering for the first time, the full MCS database will be reconciled.
- In the case that the MCS communication is lost and then re-established but neither node rebooted during the connection loss, only the information not synchronized during this time will be reconciled (using sequence numbers helps identify information which was not synchronized).
- In the case that MCS communication is lost because of excessive delay in ACK messages but no information has been effectively lost, the MCS process indicates a loss of synchronization but no reconciliation is performed.

## Subscriber Routed Redundancy Protocol (SRRP)

Subscriber Routed Redundancy Protocol (SRRP) is closely tied to the multi-chassis synchronization protocol used to synchronize information between redundant nodes. An MCS peer must be configured and operational when subscriber hosts have a redundant connection to two nodes. Subscriber hosts are identified by the ingress SAP, the hosts IP and MAC addresses. Once a host is identified on one node, the MCS peering is used to inform the other node that the host exists and conveys the dynamic DHCP lease state information of the host. MCS creates a common association between the virtual ports (SAPs) shared by a subscriber. This association is configured at the MCS peering level by defining a tag for a port and range of SAPs. The same tag is defined on the other nodes peering context for another port (does not need to be the same port-ID) with the same SAP range. In this manner, a subscriber host and Dot1Q tag sent across the peering with the appropriate tag will be mapped to the redundant SAP on the other node.

Once SRRP is active on the group IP interface, the SRRP instance will attempt to communicate through in-band (over the group IP interfaces SAPs) and out-of-band (over the group IP interfaces redundant IP interface) messages to a remote router. If the remote router is also running SRRP with the same SRRP instance ID, one router will enter a master state while the other router will enter a backup state. Since both routers are sharing a common SRRP gateway MAC address that is used for the SRRP gateway IP addresses and for proxy ARP functions, either node may act as the default gateway for the attached subscriber hosts.

For proper operation, each subscriber subnet associated with the SRRP instance must have a gw-address defined. The SRRP instance cannot be activated (no shutdown) unless each subscriber subnet associated with the group IP interface has an SRRP gateway IP address. Once the SRRP instance is activated, new subscriber subnets cannot be added without a corresponding SRRP gateway IP address. [Table 13](#) describes how the SRRP instance state is used to manage access to subscriber hosts associated with the group IP interface.

SRRP instances are created in the disabled state (shutdown). To activate SRRP the no shutdown command in the SRRP context must be executed.

Before activating an SRRP instance on a group IP interface, the following actions are required:

- Add SRRP gateway IP addresses to all subscriber subnets associated with the group IP interface, including subnets on subscriber IP interfaces associated as retail routing contexts (at least one subnet must be on the subscriber IP interface containing the group IP interface and its SRRP instance).
- Create a redundant IP interface and associate it with the SRRP instances group IP interface for shunting traffic to the remote router when master.
- Specify the group IP interface SAP used for SRRP advertisement and Information messaging.

Before activating an SRRP instance on a group IP interface, the following actions should be considered:



- Associate the SRRP instance to a Multi-Chassis Synchronization (MCS) peering terminating on the neighboring router (the MCS peering should exist as the peering is required for redundant subscriber host management)
- Define a description string for the SRRP instance
- Specify the SRRP gateway MAC address used by the SRRP instance (must be the same on both the local and remote SRRP instance participating in the same SRRP context)
- Change the base priority for the SRRP instance
- Specify one or more VRRP policies to dynamically manage the SRRP instance base priority
- Specify a new keep alive interval for the SRRP instance

Table 13 lists the SRRP's state effect on subscriber hosts associated with group IP interfaces.

**Table 13: SRRP State Effect on Subscriber Hosts Associated with Group IP Interface**

SRRP State	ARP	Local Proxy ARP Enabled	Remote Proxy ARP Enabled	Subscriber Host Routing
Disabled	<ul style="list-style-type: none"> <li>Will respond to ARP for all owned subscriber subnet IP addresses.</li> <li>Will not respond to ARP for subscriber subnet SRRP gateway IP addresses.</li> <li>All ARP responses will contain the native MAC of the group IP interface (not the SRRP gateway MAC).</li> </ul>	<ul style="list-style-type: none"> <li>Will respond to ARP for all subscriber hosts on the subscriber subnet.</li> </ul>	<ul style="list-style-type: none"> <li>Will respond to ARP for all reachable remote IP hosts.</li> </ul>	<ul style="list-style-type: none"> <li>All routing out the group IP interface will use the native group IP interface MAC address.</li> <li>The group IP interface redundant IP interface will not be used.</li> <li>Will not accept packets destined to the SRRP gateway MAC received on the group IP interface.</li> </ul>
Becoming Master (In order to enter becoming master state, a master must currently exist)	<ul style="list-style-type: none"> <li>Will respond to ARP for all owned subscriber subnet IP addresses (hardware address and source MAC = group IP interface native MAC).</li> <li>Will respond to ARP for subscriber subnet SRRP gateway IP addresses (hardware address = SRRP gateway IP address, source MAC = group IP interface native MAC).</li> </ul>	<ul style="list-style-type: none"> <li>Will respond to ARP for all subscriber hosts on the subscriber subnet (hardware address = SRRP gateway MAC address, source MAC = group IP interface native MAC).</li> </ul>	<ul style="list-style-type: none"> <li>Will respond to ARP for all reachable remote IP hosts (hardware address = SRRP gateway MAC address, source MAC = group IP interface native MAC).</li> </ul>	<ul style="list-style-type: none"> <li>All routing out the group IP interface will use the native group IP interface MAC address.</li> <li>Subscriber hosts mapped to the redundant IP interface are remapped to the group IP interface.</li> <li>Will not accept packets destined to the SRRP gateway MAC received on the group IP interface.</li> </ul>

**Table 13: SRRP State Effect on Subscriber Hosts Associated with Group IP Interface (Continued)**

<b>SRRP State</b>	<b>ARP</b>	<b>Local Proxy ARP Enabled</b>	<b>Remote Proxy ARP Enabled</b>	<b>Subscriber Host Routing</b>
Master	<ul style="list-style-type: none"> <li>• Will respond to ARP for all owned subscriber subnet IP addresses (hardware address and source MAC = group IP interface native MAC).</li> <li>• Will respond to ARP for subscriber subnet SRRP gateway IP addresses (hardware address = SRRP gateway IP address, source MAC = group IP interface native MAC).</li> </ul>	<ul style="list-style-type: none"> <li>• Will respond to ARP for all subscriber hosts on the subscriber subnet (hardware address = SRRP gateway MAC address, source MAC = group IP interface native MAC).</li> </ul>	<ul style="list-style-type: none"> <li>• Will respond to ARP for all reachable remote IP hosts (hardware address = SRRP gateway MAC address, source MAC = group IP interface native MAC).</li> </ul>	<ul style="list-style-type: none"> <li>• All routing out the group IP interface will use the SRRP gateway MAC address.</li> <li>• Subscriber hosts mapped to the redundant IP interface are remapped to the group IP interface.</li> <li>• Will accept packets destined to the SRRP gateway MAC received on the group IP interface.</li> </ul>
Becoming Backup (redundant IP interface operational)	<ul style="list-style-type: none"> <li>• Will respond to ARP for all owned subscriber subnet IP addresses (hardware address and source MAC = group IP interface native MAC).</li> <li>• Will not respond to ARP for subscriber subnet SRRP gateway IP addresses</li> </ul>	<ul style="list-style-type: none"> <li>• Will not respond to ARP for any subscriber hosts on the subscriber subnet.</li> </ul>	<ul style="list-style-type: none"> <li>• Will not respond to ARP for any remote IP hosts.</li> </ul>	<ul style="list-style-type: none"> <li>• Will not route out the group IP interface for subscriber hosts associated with the subscriber subnet.</li> <li>• Subscriber hosts mapped to the group IP interface are remapped to the redundant IP interface.</li> <li>• Will accept packets destined to the SRRP gateway MAC received on the group IP interface.</li> </ul>

**Table 13: SRRP State Effect on Subscriber Hosts Associated with Group IP Interface (Continued)**

<b>SRRP State</b>	<b>ARP</b>	<b>Local Proxy ARP Enabled</b>	<b>Remote Proxy ARP Enabled</b>	<b>Subscriber Host Routing</b>
Becoming Backup (redundant IP interface not available)	<ul style="list-style-type: none"> <li>• Will respond to ARP for all owned subscriber subnet IP addresses (hardware address and source MAC = group IP interface native MAC).</li> <li>• Will not respond to ARP for subscriber subnet SRRP gateway IP addresses.</li> </ul>	<ul style="list-style-type: none"> <li>• Will not respond to ARP for any subscriber hosts on the subscriber subnet.</li> </ul>	<ul style="list-style-type: none"> <li>• Will not respond to ARP for any remote IP hosts.</li> </ul>	<ul style="list-style-type: none"> <li>• Will route out the group IP interface for subscriber hosts associated with the subscriber subnet using the group IP interface native MAC address.</li> <li>• Subscriber hosts mapped to the redundant IP interface are remapped to the group IP interface.</li> <li>• Will accept packets destined to the SRRP gateway MAC received on the group IP interface</li> </ul>
Backup (redundant IP interface operational)	<ul style="list-style-type: none"> <li>• Will respond to ARP for all owned subscriber subnet IP addresses (hardware address and source MAC = group IP interface native MAC).</li> <li>• Will not respond to ARP for subscriber subnet SRRP gateway IP addresses.</li> </ul>	<ul style="list-style-type: none"> <li>• Will not respond to ARP for any subscriber hosts on the subscriber subnet</li> </ul>	<ul style="list-style-type: none"> <li>• Will not respond to ARP for any remote IP hosts</li> </ul>	<ul style="list-style-type: none"> <li>• Will not route out the group IP interface for subscriber hosts associated with the subscriber subnet.</li> <li>• Subscriber hosts mapped to the group IP interface are remapped to the redundant IP interface.</li> <li>• Will not accept packets destined to the SRRP gateway MAC received on the group IP interface.</li> </ul>

**Table 13: SRRP State Effect on Subscriber Hosts Associated with Group IP Interface (Continued)**

<b>SRRP State</b>	<b>ARP</b>	<b>Local Proxy ARP Enabled</b>	<b>Remote Proxy ARP Enabled</b>	<b>Subscriber Host Routing</b>
Backup (redundant IP interface not available)	<ul style="list-style-type: none"> <li>• Will respond to ARP for all owned subscriber subnet IP addresses (hardware address and source MAC = group IP interface native MAC).</li> <li>• Will not respond to ARP for subscriber subnet SRRP gateway IP addresses.</li> </ul>	<ul style="list-style-type: none"> <li>• Will not respond to ARP for any subscriber hosts on the subscriber subnet.</li> </ul>	<ul style="list-style-type: none"> <li>• Will not respond to ARP for any remote IP hosts.</li> </ul>	<ul style="list-style-type: none"> <li>• Will route out the group IP interface for subscriber hosts associated with the subscriber subnet using the group IP interface native MAC address.</li> <li>• Subscriber hosts mapped to the redundant IP interface are remapped to the group IP interface.</li> <li>• Will not accept packets destined to the SRRP gateway MAC received on the group IP interface.</li> </ul>

### SRRP Messaging

SRRP uses the same messaging format as VRRP with slight modifications. The source IP address is derived from the system IP address assigned to the local router. The destination IP address and IP protocol are the same as VRRP (224.0.0.18 and 112, respectively).

The message type field is set to 1 (advertisement) and the protocol version is set to 8 to differentiate SRRP message processing from VRRP message processing.

The vr-id field has been expanded to support an SRRP instance ID of 32 bits.

Due to the large number of subnets backed up by SRRP, only one message every minute carries the gateway IP addresses associated with the SRRP instance. These gateway addresses are stored by the local SRRP instance and are compared with the gateway addresses associated with the local subscriber IP interface.

Unlike VRRP, only two nodes may participate in an SRRP instance due to the explicit association between the SRRP instance group IP interface, the associated redundant IP interface and the multi-chassis synchronization (MCS) peering. Since only two nodes are participating, the VRRP skew timer is not utilized when waiting to enter the master state. Also, SRRP always preempts when the local priority is better than the current master and the backup SRRP instance always inherits the master's advertisement interval from the SRRP advertisement messaging.

SRRP advertisement messages carry a *becoming-master* indicator flag. The *becoming-master* flag is set by a node that is attempting to usurp the master state from an existing SRRP master router. When receiving an SRRP advertisement message with a better priority and with the *becoming-master* flag set, the local master initiates its *becoming-backup* state, stops routing with the SRRP gateway MAC and sends an SRRP advertisement message with a priority set to zero. The new master continues to send SRRP advertisement messages with the *becoming-master* flag set until it either receives a return priority zero SRRP advertisement message from the previous master or its *becoming-master* state timer expires. The new backup node continues to send zero priority SRRP advertisement messages every time it receives an SRRP advertisement message with the *becoming-master* flag set. After the new master either receives the old master's priority zero SRRP advertisement message or the *become-master* state timer expires, it enters the *master* state. The *become-master* state timer is set to 10 seconds upon entering the *become-master* state.

The SRRP advertisement message is always evaluated to see if it has higher priority than the SRRP advertisement that would be sent by the local node. If the advertised priority is equal to the current local priority, the source IP address of the received SRRP advertisement is used as a tie breaker. The node with the lowest IP address is considered to have the highest priority.

The SRRP instance maintains the source IP address of the current master. If an advertisement is received with the current masters source IP address and the local priority is higher priority than the masters advertised priority, the local node immediately enters the *becoming-master* state unless the advertised priority is zero. If the advertised priority is zero, the local node bypasses the *becoming-master* state and immediately enters the *master* state. Priority zero is a special case and is sent when an SRRP instance is relinquishing the master state.

## SRRP and Multi-Chassis Synchronization

In order to take full advantage of SRRP resiliency and diagnostic capabilities, the SRRP instance should be tied to a MCS peering that terminates on the redundant node. The SRRP instance is tied to the peering using the **srrp srrp-id** command within the appropriate MCS peering configuration. Once the peering is associated with the SRRP instance, MCS will synchronize the local information about the SRRP instance with the neighbor router. MCS automatically derives the MCS key for the SRRP instance based on the SRRP instance ID. For example, an SRRP instance ID of 1 would appear in the MCS peering database with a MCS-key srrp-0000000001.

The SRRP instance information stored and sent to the neighbor router consists of:

- The SRRP instance MCS key
- Containing service type and ID
- Containing subscriber IP interface name
- Subscriber subnet information
- Containing group IP interface information
- The SRRP group IP interface redundant IP interface name, IP address and mask
- The SRRP advertisement message SAP
- The local system IP address (SRRP advertisement message source IP address)
- The Group IP interface MAC address
- The SRRP gateway MAC address
- The SRRP instance administration state (up / down)
- The SRRP instance operational state (disabled / becoming-backup / becoming-master / master)
- The current SRRP priority
- Remote redundant IP interface availability (available / unavailable)
- Local receive SRRP advertisement SAP availability (available / unavailable)

### SRRP Instance

The SRRP instance uses the received information to verify provisioning and obtain operational status of the SRRP instance on the neighboring router.

- [SRRP Instance MCS Key on page 924](#)
  - [Containing Service Type and ID on page 924](#)
  - [Containing Subscriber IP Interface Name on page 924](#)
  - [Subscriber Subnet Information on page 925](#)
- 

### SRRP Instance MCS Key

The SRRP instance MCS key ties the received MCS information to the local SRRP instance with the same MCS key. If the received key does not match an existing SRRP instance, the MCS information associated with the key is ignored. Once an SRRP instance is created and mapped to an MCS peering, the SRRP instance evaluates received information with the same MCS key to verify it corresponds to the same peering. If the received MCS key is on a different peering than the local MCS key an SRRP peering mismatch event is generated detailing the SRRP instance ID, the IP address of the peering the MCS key is received on and the IP address to which the local MCS key is mapped. If the peering association mismatch is corrected, an SRRP peering mismatch clear event is generated.

---

### Containing Service Type and ID

The Containing Service Type is the service type (IES or VPRN) that contains the local SRRP instance. The Containing Service ID is the service ID of that service. This information is supplied for troubleshooting purposes only and is not required to be the same on both nodes.

---

### Containing Subscriber IP Interface Name

The containing subscriber IP interface name is the subscriber IP interface name that contains the SRRP instance and its group IP interface. This information is supplied for troubleshooting purposes only and is not required to be the same on both nodes.



## Subscriber Subnet Information

The subscriber subnet information includes all subscriber subnets backed up by the SRRP instance. The information for each subnet includes the Owned IP address, the mask and the gateway IP address. If the received subscriber subnet information does not match the local subscriber subnet information, an SRRP Subscriber Subnet Mismatch event is generated describing the SRRP instance ID and the local and remote node IP addresses. Once the subscriber subnet information matches, an SRRP Subscriber Subnet Mismatch Clear event is generated.

---

## Containing Group IP Interface Information

The containing group IP interface information is the information about the group IP interface that contains the SRRP instance. The information includes the name of the group IP interface, the list of all SAPs created on the group IP interface, the administrative and operational state of each SAP and the MCS key and the peering destination IP address associated with each SAP. To obtain the MCS information, the SRRP instance queries MCS to determine the peering association of the SRRP instance and then queries MCS for each SAP on the group IP interface. If the local SRRP instance is associated with a different MCS peering than any of the SAPs or if one or more SAPs are not tied to an MCS peering, an SRRP group interface SAP peering mismatch event is generated detailing the SRRP instance ID, and the group IP interface name.

When receiving the remote containing group IP interface information, the local node compares the received SAP information with the local group IP interface SAP information. If a local SAP is not included in the SAP information or a remote SAP is not included in the local group IP interface, an SRRP Remote SAP mismatch event is generated detailing the SRRP instance ID and the local and remote group IP interface names. If a received SAP's MCS key does not match a local SAP's MCS Key, an SRRP SAP MCS key mismatch event is generated detailing the SRRP instance ID, the local and remote group IP interface names, the SAP-ID and the local and remote MCS keys.

---

## Remote Redundant IP Interface Mismatch

If the group IP remote redundant IP interface address space does not exist, is not within the local routing context for the SRRP instances group IP interface or is not on a redundant IP interface, the local node sends redundant IP interface unavailable to prevent the remote neighbor from using its redundant IP interface. An SRRP redundant IP interface mismatch event is generated for the SRRP instance detailing the SRRP instance, the local and remote system IP addresses, the local and remote group IP interface names and the local and remote redundant IP interface names and IP addresses and masks. The local redundant IP interface may still be used if the remote node is not sending redundant IP interface unavailable.

### **Remote Sending Redundant IP Interface Unavailable**

If the remote node is sending redundant IP interface unavailable, the local node will treat the local redundant IP interface associated with the SRRP instances group IP interface as down. A Local Redundant IP Interface Unavailable event is generated detailing the SRRP instance ID, the local and remote system IP addresses, the local group IP interface name, the local redundant IP interface name and the redundant IP interface IP address and mask.

---

### **Remote SRRP Advertisement SAP Non-existent**

If the remote node's SRRP advertisement SAP does not exist on the local SRRP instances group IP interface, the local node sends local receive SRRP advertisement SAP unavailable to the remote node. An SRRP receive advertisement SAP non-existent event is generated detailing the SRRP instance ID, the local and remote system IP addresses, the local group IP interface name and the received remote SRRP advertisement SAP. Since SRRP advertisement messages cannot be received, the local node will immediately become master if it has the lower system IP address.

---

### **Remote Sending Local Receive SRRP Advertisement SAP Unavailable**

If the local node is receiving local receive SRRP advertisements stating that the SAP is unavailable from the remote node, an SRRP Remote Receive advertisement SAP Unavailable event will be generated. This details the SRRP instance ID, the local and remote system IP addresses, the remote group IP interface name and the local SRRP advertisement SAP. Since the remote node cannot receive SRRP advertisement messages, the local node will immediately become master if it has the lower system IP address.

---

### **Local and Remote Dual Master Detected**

If the local SRRP state is master and the remote SRRP state is master, an SRRP dual master event is generated detailing the SRRP instance ID and the local, remote system IP addresses and the local and remote group IP interface names and port numbers.

## Subscriber Subnet Owned IP Address Connectivity

In order for the network to reliably reach the owned IP addresses on a subscriber subnet, the owning node must advertise the IP addresses as /32 host routes into the core. This is important since the subscriber subnet is advertised into the core by multiple routers and the network will follow the shortest path to the closest available router which may not own the IP address if the /32 is not advertised within the IGP.

---

## Subscriber Subnet SRRP Gateway IP Address Connectivity

The SRRP gateway IP addresses on the subscriber subnets cannot be advertised as /32 host routes since they may be active (master) on multiple group IP interfaces on multiple SRRP routers. Without a /32 host route path, the network will forward any packet destined to an SRRP gateway IP address to the closest router advertising the subscriber subnet. While a case may be made that only a node that is currently forwarding for the gateway IP address in a master state should respond to ping or other diagnostic messages, the distribution of the subnet and the case of multiple masters make any resulting response or non-response inconclusive at best. To provide some ability to ping the SRRP gateway address from the network side reliably, any node receiving the ICMP ping request will respond if the gateway IP address is defined on its subscriber subnet.

---

## Receive SRRP Advertisement SAP and Anti-Spoof

The group IP interface SAPs are designed to support subscriber hosts and perform an ingress anti-spoof function that ensures that any IP packet received on the group IP interface is coming in the correct SAP with the correct MAC address. If the IP and MAC are not registered as valid subscriber hosts on the SAP, the packet is silently discarded. Since the SRRP advertisement source IP addresses are not subscriber hosts, an anti-spoof entry will not exist and SRRP advertisement messages would normally be silently discarded. To avoid this issue, when a group IP interface SAP is configured to send and receive SRRP advertisement messages, anti-spoof processing on the SAP is disabled. This precludes subscriber host management on the SRRP messaging SAP.

## PPPoE MC Redundancy

This feature minimizes the downtime for PPPoE clients in an ESM environment when a single node fails.

But it is not necessary that an entire BNG fails before it triggers the corrective action. The solution outlined in this document will natively include protection against interfaces and line card failures within the BNG. The redundant (protective) entity, however, does not reside within the same BNG on which the failure occurs but instead it is on a separate BNG node.

The PPPoE MC Redundancy is based on SRRP and MC-LAG because SRRP is already established in ESM providing IPoE MC Redundancy. With some modifications, SRRP approach is adopted to PPPoE deployments.

---

## Hardware Support

This feature is supported on the following platforms:

- 7750 SR-7/12
- 7750-c4/12
- 7450 in mixed chassis mode.

MCS across different platform types (7750 SR-7/12, 7750-c4/12, 7450) is not supported. For example, MCS between 7750 SR-7/12 and 7750-c12 is not supported.

This feature is supported in the following chassis modes: B, C, D and SR Mixed Mode (IPv6 on IOM3 while IOM1 cards can be present in the same system).

IPv4 functionality is supported on IOM2 cards and IPv4/IPv6 on IOM3 cards.

**Note:** ESM v6 is supported only on IOM3 cards. IPv6 forwarding in ESM between IOM3 and IOM2 cards is not supported – for example, if the access side is IOM3 and the network side is IOM2. However, plain routing (non-ESM related) is supported between these two cards.

## SRRP Considerations for PPPoE

SRRP is based on VRRP whose purpose is to provide a default gateway redundancy for clients sharing the transport medium such as Ethernet. IPoE would be a typical example of this where IPoE clients use a virtual IP and MAC address that is shared between two default gateway nodes in the Master/Backup configuration. SRRP supports only two nodes in a cluster but VRRP allows multiple nodes to be configured in a cluster with a priority that will determine which node will assume Mastership. Although it is mandatory for the proper operation of IPoE clients that the same SRRP IP address is shared between the two BNG nodes providing redundancy, having the same SRRP IP address is not necessary for the operation of SRRP itself. In other words, SRRP itself (Master/Backup states) will work with different SRRP IP addresses on each node. Same is valid for MAC addressing. It is possible by configuration that the redundant BNG nodes use different IP/MAC addresses on a pair of SRRP instances.

Upon a switchover, a gratuitous ARP is sent from a newly selected active node so that each IPoE client can update the ARP table, if the MAC address has indeed changed (it does not have to). More importantly, if an Layer 2 aggregation network is in place between the BNG and the IPoE client, all intermediate Layer 2 devices will have to update their port-to-mac mappings (Layer 2 FIB). The above described process will ensure proper packet addressing on the IPoE client side as well as the proper forwarding path through Layer 2 aggregation network to the newly activated BNG.

When considering PPPoE in conjunction with SRRP, keep in mind that PPP protocol (point-to-point protocol) is adopted for the Ethernet (shared medium) by enabling an extra Ethernet related layer in PPP that allows sharing of point-to-point sessions over Ethernet (shared medium). The result is a PPPoE protocol designed to ‘tunnel’ each PPP session over Ethernet.

PPPoE is not aware of ARP (Address Resolution Protocol) and it will not react to gratuitous ARP packets sent by a newly active BNG. The destination MAC address that PPPoE clients will use when sending traffic is determined not by ARP but by the PPPoE Discovery phase at the beginning of the session establishment. This originally discovered destination MAC is used throughout the lifetime of the session. This has a couple of consequences:

1. If SRRP is used for PPPoE then the ‘SRRP’ MAC address between the redundant BNG nodes must be shared. It is not allowed to use a unique ‘SRRP’ MAC address per BNG in the redundant pair of BNG nodes (as it is allowed today for IPoE). Every PADx conversation is based on the SRRP shared MAC address, that is, the PADO reply must have the shared SRRP MAC address as the source MAC. This has a significant impact on the operation of MSAP in conjunction with this feature.
2. Since PPPoE sessions are not ARP aware, the only purpose of the gratuitous ARP would be to update the Layer 2 FIB in the aggregation network (and not the PPPoE client destination MAC address). For IPoE, the gratuitous ARP is sent for ALL subnet gateway IP addresses found under the subscriber interface over either all SAPs (default) or top-tags only. For PPPoE, the gratuitous ARP is sent only for the system IP address. The purpose of the gratuitous ARP in PPPoE scenario is only to update Layer 2 network path which is otherwise IP unaware. It is not necessary to send the gratuitous ARP for every default-gateway address found under the subscriber-interface. Since this feature is only applicable to PPPoE deploy-

ments, therefore, only PPPoE is present under the group interface. This is indicated by the following command under the SRRP node:

```
group-interface <name>
srrp <id>
    one-garp-per-sap
```

---

### SRRP Fact-Checks

1. Once Multi-chassis Synchronization (MCS) for subscriber management and SRRP is enabled, both BNG nodes, Master and Backup will in general forward packets (for subscribers) in both directions.
  2. Traffic flows through an SRRP enabled node according to the entries in the SRRP sync database and the SRRP state of the node:
    - Backup SRRP directs downstream traffic over the redundant-interface towards the Master SRRP node. If the redundant interface is unavailable, traffic is sent directly to the subscriber.
    - Master SRRP always directly forwards the downstream traffic towards the subscriber.
    - In the upstream direction, the active SRRP node accepts subscriber traffic addressed either to the MAC address of the SRRP active group OR the native interface MAC address.
    - The standby node accepts in the upstream direction only packets addressed to its native interface MAC address.
  3. If both SRRP nodes become Masters then both forward traffic to/from subscribers unaware of the link failure somewhere in the Layer 2 network. As a result, downstream traffic can be blackholed. Whether downstream traffic will be lost depends on the native routing on the network side, which is unaware of the failures in the aggregation network.
- 

### State Synchronization

PPPoE sessions are synchronized between the redundant BNG nodes. The subscriber synchronization is achieved through Multichassis Synchronization (MCS) protocol in a similar way it is performed for IPoE.

```
multi-chassis
    peer <IP@>create
        sync
            local-dhcp-server
            SRRP
            sub-mgmt [ipoe | pppoe]
        :
        :
        no shutdown
    exit
    no shutdown
exit
```

A two keywords, **ipoe** and **pppoe** enable a more granular control over which type of subscribers the MCS should be enabled.

Subscriber synchronization is important for following reasons:

1. Forwarding of downstream traffic between the redundant BNG nodes through a redundant interface is an artifact of how natural routing steers traffic through the network.
2. Subscriber instantiation on the node which did not originally create subscriber session. This drastically reduces downtime during the SRRP switchover.
3. Monitors operational aspects of the subscriber management through show commands.

---

### PPPoE Multi-chassis Synchronization (MCS) Model

PPPoE MCS model is based on SRRP synchronization and can be used in a centralized or distributed environment with or without Layer 2 aggregation network in-between DSLAMs and BNG nodes. The failure detection speed is dependent on SRRP timers. Traffic load can be balanced per SRRP group over the two links. In this model ([Figure 81](#)), PPPoE states are synchronized between the redundant BNG nodes. If one BNG fails, the newly activated BNG sends out a 'MAC update' (gratuitous ARP) message prompting the intermediate Layer 2 nodes to update their forwarding tables so that forwarding can resume. The SRRP timers can be configured in the sub-second range. In reality, the limiting factor for timer values is the scale of the deployment, in particular the number of SRRP groups per node.

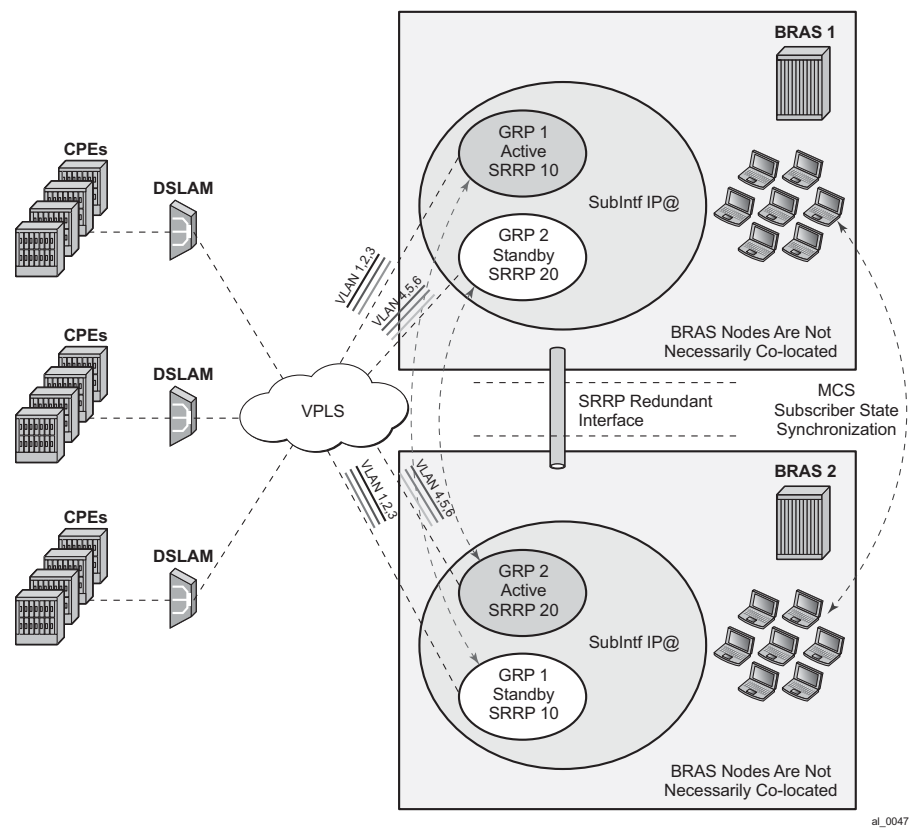


Figure 81: Fully Redundant “Statefull 1:1” Model



## Traffic Control and Redundant Interface

To preserve QoS and Accounting, subscriber's traffic must flow in both directions through the Master BNG node.

In the upstream direction, this is always true as traffic is steered to the master SRRP node just by the virtue of SRRP operation.

In the downstream direction which represents bulk of traffic, SRRP can not be relied up on to steer traffic through the Master node. This poses a problem in a very common environment where IP subnets are shared over multiple group-interfaces with SRRP enabled. A particular subnet will be advertised to the network side from both BNG nodes, Master and Backup. Natural routing on the network side will determine which BNG node will receive subscriber's traffic in the downstream direction. If the Backup SRRP node receives the traffic, it cannot simply send the traffic directly to the access network where the subscriber resides by just inserting the source MAC address of the SRRP instance in the outgoing packet. This would break the operation of SRRP. Instead, the Backup BNG node must send the traffic to the Master BNG node via a redundant-interface. The Master SRRP node would then forward traffic directly to the subscriber. Source MAC address of this traffic would then be the MAC address of SRRP instance. This traffic shunting over the redundant interface can result in a substantial load on the link between the two BNG nodes.

The increase in shunted traffic can quickly become an issue if the redundant BNG nodes that are not collocated. To minimize the shunt traffic, more granular routing information must be presented to the network core. This would lead to more optimal routing where downstream subscriber traffic would be directed towards the Master BNG node, without the need to cross the redundant interface. The downside of this approach is that this would further fragment the IP address space within the network core. In the extreme case where /32 (subscriber) IP addresses are advertised, the churn that /32s can cause in the core routing would most likely be unsustainable. In this case, routing updates in the core would be triggered by subscribers coming on/off-line.

Optimal operation would call for the shunt traffic to be eliminated and at the same time, a high IP route aggregation on the network side is achieved. The existence of the shunt traffic stems from the fact that routing protocols advertise subscriber subnets into the network with no awareness of the SRRP activity state (Master/Standby). To address this problem along with better aggregation of advertised subnets, two SRRP enhancements are introduced:


- SRRP fate-sharing
- SRRP aware routing

Both of this concepts are described under the 'SRRP Enhancements' section.

Traffic destined to/from the subscriber is forwarded under the condition that the subscriber-interface is operationally UP. This applies also to shunting of downstream subscriber traffic from the STANDBY to MASTER node. It is always necessary to keep the subscriber-interface operationally UP by configuring a dummy *group-interface* with a command *oper-up-while-empty* under it. This is especially true for the MC-LAG which causes the messaging SAP on the STANDBY node always to be in the INIT state. In case that MSAPs are used on such group-interfaces, the group-interfaces would be also operationally DOWN, causing the subscriber-interface to be operationally DOWN.

## Subnet Assignment and Advertisement - Option 'A'

A single IP subnet is used for all subscribers terminated within the redundant BNG nodes. The upside of the Option 'A' is that it offers aggregated IP addressing in the network core per pair of redundant BNG nodes. The downside is that the subscriber termination point (active BNG for the SRRP group) is hidden from the network core. Since both BNG nodes share the same IP subnet for the subscribers, the natural routing can cause downstream traffic to be sent to the standby BNG which in turn will have to shunt the traffic to the active BNG. It is likely that half of the traffic will be shunted over the redundant-interface with this approach. This scenario is shown in [Figure 82](#).



al\_0048

Figure 82: Shared Subscriber IP Space

## Subnet Assignment and Advertisement - Option 'B'

With the option 'B', an IP address pool (or subnet) can be allocated per group of SRRP instances that are in the Master state. The routing decision on the network side is further influenced by the static increase of the metric of the advertised route on the BNG node hosting the active SRRP groups (Figure 83).

This approach would cause greater IP space segmentation in the network core, but at the same time, it would indirectly provide more information about the subscriber whereabouts and thus minimize or eliminate the shunt traffic during the normal operation. However, in the case of a SRRP switchover, the shunt traffic would ensue. The amount of the shunted traffic would depend on the scale of the failure. From the Figure 83, it can be concluded that:

- In the depicted scenario on Figure 83, there is no shunted traffic.
- If any of the SRRP instances transitions out of the Master state, traffic for an entire IP network associated with this failed SRRP instance would be shunted. The reason for this is that the advertised route metric is static and it does not follow changes in SRRP state.

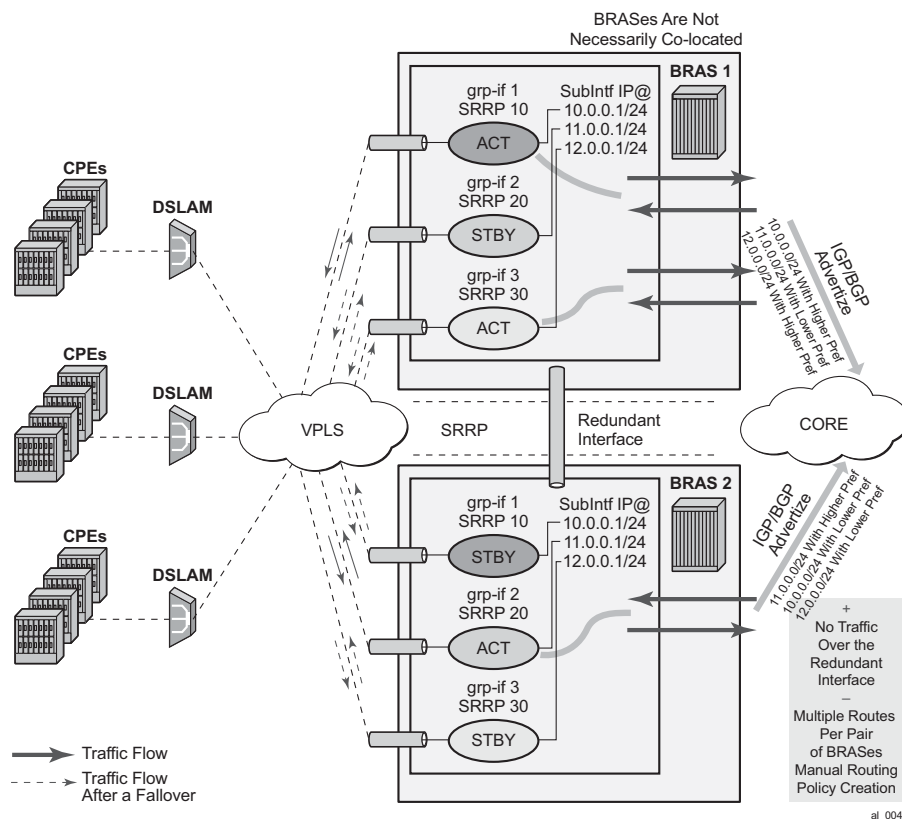


Figure 83: Option 'B' – IP Subnet per Active SRRP Group

## MSAP Considerations

As per RFC 2516 (PPPoE), this has the implications on the operation of the capture SAP. In IPoE environment, the initial DHCP traffic related to host establishment will use its native MAC of the physical port on 7x50. Once the group-interface is learnt (later in the process, via RADIUS or msap-policy), the MAC address is switched to SRRP MAC address (virtual MAC). The IPoE client will adapt easily to this change. On the contrary, for the proper operation of PPPoE with SRRP, the initial destination MAC address learned by the PPPoE client does not change during the lifetime of the session.

This is ensured by indirectly referencing the grp-if under the capture SAP:

```
configure>service>vpls
  sap 1/1/1:1.* capture-sap
track-srrp 10
  sap 1/1/1:2.* capture-sap
track-srrp 20

configure>service>vprn>
  subscriber-interface <if-name>
    group-interface <grp-if-name>
      sap 1/1/1:1.1
      srrp 10
      message-path 1/1/1:1.1

    group-interface <grp-if-name>
      sap 1/1/1:2.1
      srrp 20
      message-path 1/1/1:2.1
```

With this approach the grp-if is nailed during the session initiation phase by referencing the SRRP instance in track-srrp statement (srrp is a grp-if wide concept). RADIUS returned grp-if name must match the one on which referenced SRRP instance runs.

The capture SAP of the form

```
sap port-id:*. * capture-sap
  track-srrp X
```

assumes that there is only one grp-if associated with all msaps under this capture SAP.

A check is put in place to make sure that the MAC addresses associated with the SRRP instance is the same as the MAC address of the associated capture SAP. A log is raised if there is a discrepancy between the MAC addresses while the grp-if is operationally UP. If there is a MAC address change (user misconfiguration) then the existing PPPoE sessions will time out and the new sessions will fail to establish until the condition is corrected.

## Unnumbered Interface Support

For unnumbered subscriber-interface support in PPPoE, the gw IP address that is used to send gratuitous ARP is not available. For this reason, the system IP address is used to send gratuitous ARPs. Gratuitous ARP is used to update the Layer 2 network forwarding path towards the BNG node in the upstream direction.

The system IP address is used automatically if the subscriber interface is unnumbered.

---

## Compatibility with MC-LAG

SRRP for PPPoE works in an environment where MC-LAG is enabled. For example, the standby LAG link automatically puts the SRRP node in a Backup state and the SRRP becomes master on the active MC-LAG link. It is important that the SRRP on the standby leg of the MC-LAG is forced into a Backup state, or any new state that will force the downstream traffic to use the redundant interface.

Traffic destined to/from the subscriber is forwarded under the condition that the subscriber-interface is operationally UP. This applies also to shunting of downstream subscriber traffic from the STANDBY to MASTER node. It is always necessary to keep the subscriber-interface operationally UP by configuring a dummy group-interface with a command oper-up-while-empty under it. This is especially true for the MC-LAG which causes the messaging SAP on the STANDBY node always to be in the INIT state. In case that MSAPs are used on such group-interfaces, the group-interfaces would be also operationally DOWN, causing the subscriber-interface to be operationally DOWN.

---

## IPv6 Support

Prerequisite for MC IPv6 Redundancy is to synchronize PPPoEv6 and IPoEv6 subscribers between the nodes via MCS.

In PPPoE environment, SRRP is used to refresh the forwarding path (MAC addresses) in the access aggregation network (via gratuitous ARP). SRRP ensures that the upstream traffic is steered to the Master BNG node. In the downstream direction, the Backup BNG directs traffic over to the Master BNG node via redundant-interface.

The IPv6 functionality currently relies on IPv4 based SRRP and IPv4 based redundant-interface. In other words, IPv4 is required to run on the access side as well as on the redundant-interface.

The redundant-interface is used in the downstream direction. Traffic arriving on the network links on the Standby node is shunted over to the Master node over the redundant-interface. This is required to ensure consistent QoS and accounting functionality across the nodes (upstream and downstream traffic on the access links for a subscriber must traverse the same BNG node). There is no IPv6 related CLI associated with the redundant-interface.

All IPv6 subscriber traffic that arrives on the Standby node in the downstream direction is automatically shunted over the IPv4 redundant-interface to the Master node. When IPv6 traffic arrives over the redundant-interface on the Master node, it is either PPPoEv6 encapsulated or left as plain IPoEv6 before it is forwarded to the subscriber.

In the upstream direction (AN->BNG) the behavior is the following:

- PPPoEv6

On the switchover, gratuitous ARPs is sent from the new Master on each vlan. The IP address in gARP is the IPv4 gw-ip address or the system IP in the case of unnumbered interfaces. This updates the Layer 2 network path with the proper SRRP MAC address.

- IPoEv6

IPv4 based SRRP is used to update the Layer 2 forwarding path in the case of a switchover. A gratuitous ARP is sent in the same fashion as it is used for IPoE v4 hosts. Router Advertisements (RA) are not sent out in the case of the switchover.

However, the two BNG nodes share the same virtual Link Local (LL) IPv6 address. This address is used by the clients as a default-gw and only the Master BNG advertises this LL address in RAs. RAs are suppressed on the Standby node. As already mentioned, RAs are not sent during the switchover. RAs are sent:

- When the client first gets established – this is how the client learns its default-gw (in PPPoE case RA can also be used for SLAAC – stateless address configuration).
- As a reply to Router Solicitations messages sent by the clients.
- Periodically to each client.

Note that RAs are unicasted to each client.

Neighbor Advertisements (NA) used for address resolution are sent only from the Master. NA has the SRRP MAC address in the target link layer option on SRRP enabled group interfaces (on non-SRRP enabled group-interfaces, NAs contains the group interface MAC address).

The syntax to configure the LL address on the subscriber interface is the following:

```
configure>service>ies | vprn>
    subscriber-interface <if-name>
        ipv6
[no] link-local-address <ipv6-address>

<ipv6-address> : ipv6-address - x:x:x:x:x:x:x:x:x:x:x:x.d.d.d.d
x      [0..FFFF]H
d      [0..255]D
```

The LL IPv6 address must be the same on both nodes. In addition, the gw-mac address must be the same on both nodes. The IPv6 clients will not be aware of the switchover and therefore they will not send NS to solicit the update of its neighbor cache with the possibly different gw-mac address.

Note that the current version of SRRP relies only on IPv4 routes. The connection between SRRP and IPv4 routes is done via the subnets with gw IP addresses defined under the subscriber-

interfaces in the ESM context. This connection is needed so that SRRP can send Gratuitous ARP properly.

These are the cases for PPPoEv6 MC Redundancy that are supported:

- unnumbered subscriber-interfaces (config>service>subscriber-interface hierarchy)
- numbered IPv4 subscriber-interfaces (config>service>subscriber-interface hierarchy)
- numbered IPv4 AND IPv6 subscriber-interfaces (config>service>subscriber-interface and config>service>subscriber-interface>ipv6 hierarchy)

numbered IPv6 only subscriber-interfaces (config>service>subscriber-interface>ipv6 hierarchy) is not supported

---

## Considerations with Local DHCP Server

When local DHCP Server redundancy/synchronization is used, using address-range failover local | remote, in conjunction with PPPoE in multi-chassis environment, both DHCP servers must be referenced under the corresponding group-interface on each node. For address-range failover access-driven configurations only one DHCP server must be referenced.

```
subscriber-interface <sub-if>
  group-interface <grp-if>
    dhcp
      server <local-dhcp-ip-address> <remote-dhcp-ip-address>
```

Otherwise, the PPPoE clients will not be synchronized via MCS.

Note that this is not the requirement in IPoE environment. In IPoE environment, it is enough that the DHCP server points to the IP address of the local DHCP server. If the IP lease is originally assigned by the peer DHCP server, the request for renewal is automatically forwarded to the remote DHCP server by the virtue of the IP address of the original DHCP server that is included in the renewal request.

It is necessary for the successful renewal of the IP address on the remote DHCP server, that the remote DHCP server has a valid return path back to the gi-address of the forwarder of the renewal request.

## Redundant Interface Considerations

In PPPoE dual-chassis environment without the redundant-interface in place, SRRP aware routing should always be used. Otherwise, if the downstream traffic arrives on the backup node, it will get forwarded directly to the client over the access network (assuming that the access network is operational) with the source MAC address of the group-interface (instead of gw-mac). This grp-if MAC address is different from the MAC address (gw-mac) negotiated during the initial PPPoE phase, and therefore, this traffic will be dropped by the client. It must be ensured that the downstream traffic is always attracted to the Master node in the absence of redundant.




## Routed Central Office (CO)

This section describes the Alcatel-Lucent routers acting as a Broadband Service Router (BSR), with Enhanced Subscriber Management enabled.

In the so called Routed Central Office (Routed CO) model, a router is positioned directly behind a DSLAM. This design removes the need for a Layer 2 aggregation network between the router and the DSLAM, however it does involve more routing entities in the network.

Figure 84 shows a DSLAM connected to a router using a Layer 3 interface within an IES or VPRN service. Operators that do not require an aggregation network can implement this topology. Typical DSLAM connection models include:

- One SAP for all subscribers with all services.
  - Subscriber management will be used for subscriber separation with DSCP/Dot1p service separation.
- One SAP per service.
  - Subscriber management will be used for subscriber separation with the SAP being the service differentiator.
- One SAP per subscriber.
  - Model with SAP level subscriber separation with DSCP/Dot1p service separation.



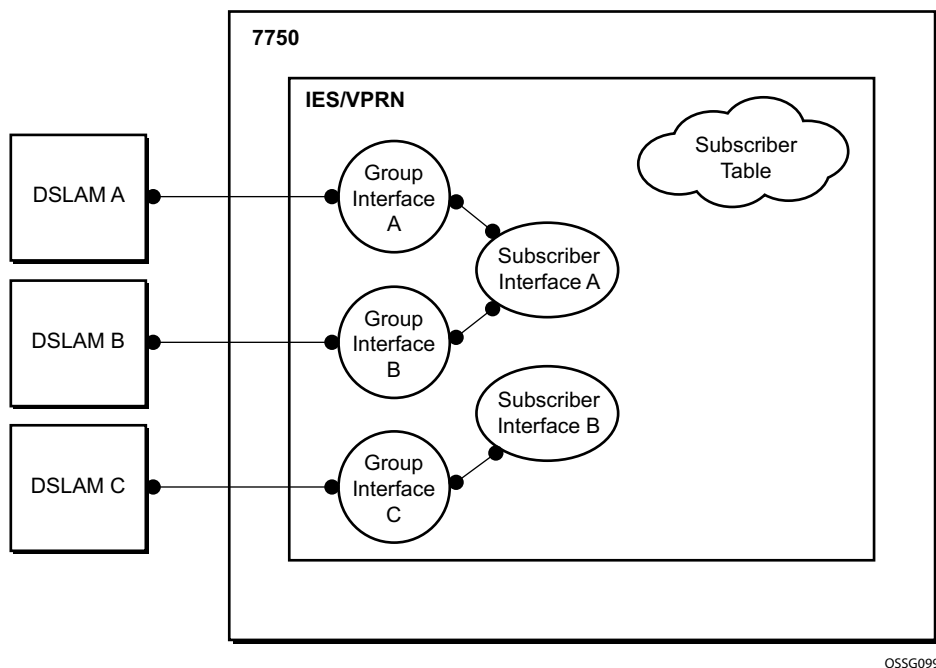
**Figure 84: DSLAM Connection**

## Layer 3 Subscriber Interfaces


On regular interfaces in an IES or VPRN service, only one SAP can be associated. A group-interface allows multiple SAPs to be configured as part of a single interface. All SAPs in a single group-interface must be within the same port. Since broadcast is not allowed in this mode, forwarding to the subscriber is based on IP/MAC addresses information gathered by the subscriber management module and stored in the subscriber management table. These entries are based on both static and dynamic DHCP hosts. Routed CO must be used with standard subscriber management or enhanced subscriber management. DSLAMs are typically deployed with Ethernet interfaces.

This model is a combination of two key technologies, subscriber interfaces and group interfaces. While the subscriber interface define the subscriber subnets, the group interfaces are responsible for aggregating the SAPs.

As depicted in [Figure 85](#), an operator can create a new subscriber interface in the IES or VPRN service. A subscriber interface allows for the creation of multiple group interfaces. The IP space is defined by the subnets of the subscriber interface's addresses. [Figure 85](#) shows the details of group interface A.




**Figure 85: Subscriber Interface in an IES/VRN Service**



OSSG301

**Figure 86: Details of a Group Interface**

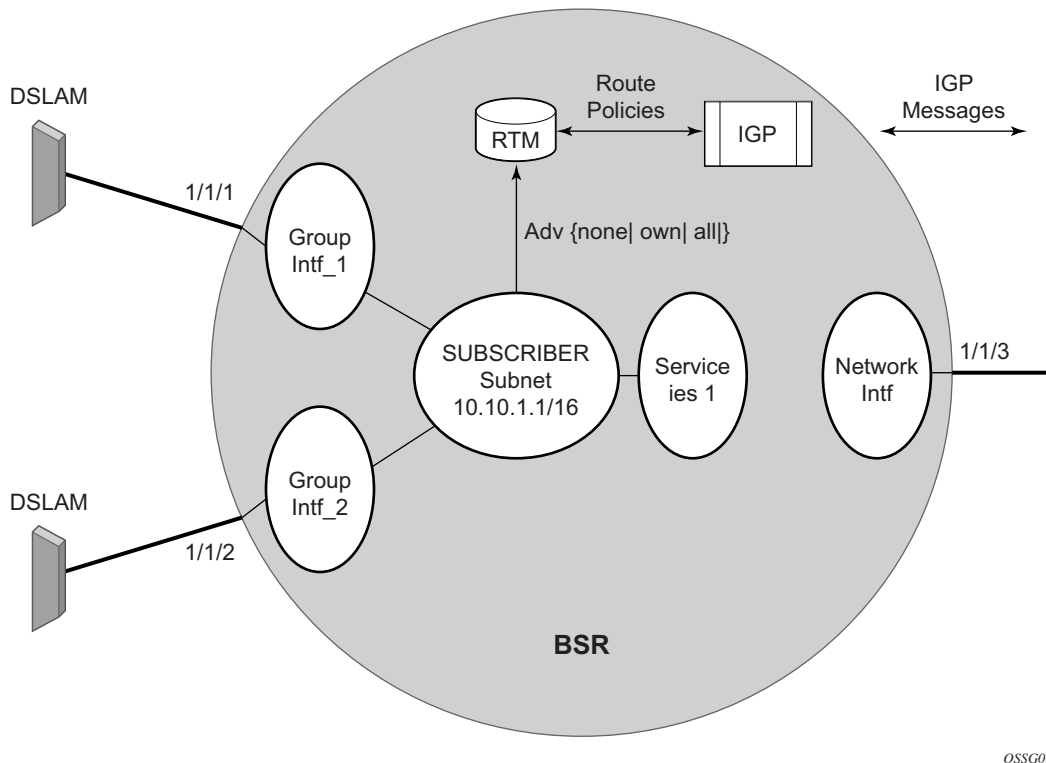
Figure 87 shows a network diagram where the DSLAM are connected directly to a Broadband Service Router (BSR) providing access to an IP subnet. Subscribers from multiple DSLAMs can be part of the same subnet. Note that BSR is also referred to as Broadband Network Gateway (BNG).



OSSG089

**Figure 87: Aggregation Network with Direct DSLAM-BSR Connection**

The BSR can be configured with multiple subnets, allowing subscribers to be part of a single subnet as well as providing mechanisms for re-addressing or expanding existing services without affecting existing users.



OSSG090

**Figure 88: Detailed View of Configurable Objects Related to Layer 3 Subscriber Interfaces**

Figure 88 shows a detailed view of a router and the configuration objects implemented to support Layer 3 subscriber interfaces.

- A subscriber service is defined by an IES Service. One or more IES services can be created.
- Each IES service concentrates a number of subscriber-interfaces. The operator can create multiple subscriber interfaces (represented as a subscriber subnet). A subscriber interface will define at least one subnet.
- A group-interface will be provisioned within the subscriber interface for each DSLAM connected. All group interfaces created under the subscriber interface will share the same subnet (or subnets). Group interfaces (shown as intf\_1 and intf\_2 in Figure 88) are configured as unnumbered and are associated with the subscriber-interface under which they are configured.
- SAPs can be configured under the group-interface. In a VLAN-per-DSLAM model only, one SAP per group-interface is needed, while in the VLAN-per-subscriber model, a subscriber of

the DSLAM will require its own SAP. All SAPs on a group-interface must be on the same physical port or LAG.

The individual features related to subscribers, such as DHCP relay, DHCP snooping and anti-spoofing filters, are enabled at group-interface level. For a Routed CO model of subscriber management, and when enhanced subscriber management (if sub-sla-mgmt is configured). Then, hashing will be based on an internally assigned subscriber-ID. Having a unique subscriber ID configured in CLI will ensure that each subscriber is assigned a unique internal subscriber ID.

It is assumed that individual end-user devices (further referred to as subscriber hosts) get their IP address assigned through either DHCP or static configuration. The management of individual subscriber hosts (such as creation, queue allocation, etc.) is performed by Enhanced Subscriber Management.

The operator can provision how the system advertises routes. While most deployments will advertise the full subnet it is possible to have the system advertise only the active, discovered or static host routes.

The distribution of this information into routing protocols will be driven by import/export route policies configured by the operator.

## DHCP Interactions

The DHCP relay process has been enhanced to record incoming DHCP discover and request messages. Since forwarding to the SAPs is done by the information in the subscriber management table and multiple SAPs are allowed in one interface it was impossible to know which SAP will be used to forward the DHCP replies. The node maintains a cache of the DHCP requests. The cache can be viewed using the **tools>dump>router>dhcp>group-if-mapping** command. The cache holds an entry for 30 seconds. If an ACK/NAK packet was not received from the server within the timeout the node discards the cache entry. The node can use the Option 82 circuit-id field as part of the temporary host entry. If used, the ACK must contain the same circuit-id field in Option 82 to be found in the cache only if the match-circuit-id is specified at the DHCP level of the group-interface. When the match-circuit-id command is enabled a check is performed for option 82 circuit-id.

---

## Routed CO for IES Service

The routed CO model depends on subscriber management to maintain the subscriber host information. To create a group-interface the operator must first create a subscriber interface within the service (**config>service>ies>subscriber-interface** *ip-int-name*). The subscriber interface maintains up to 256 subscriber subnets and is configured with a host address for each subnet.

When a DHCP ACK is received the IP address provided to the client will be verified to be in one of the subscriber subnets associated with the egress SAP. It will be noted that when DHCP snooping is enabled for regular IES interfaces the same rule will apply.

The subscriber interface is an internal loopback interface. The operational state is driven from the child's group-interface states and the configuration of an address in the RTM.

The group interface is an unnumbered interface. The interface will be operationally up if it is in the no shutdown state and if at least one SAP has been defined and is up and the parent subscriber interface is administratively up. The first SAP defined will determine the port for the group-interface. If the user attempts to define a subsequent SAP that is on a different port will result in an error. When the subscriber-interface or the group-interface is in shutdown state no packets will be delivered/received to/from the subscriber hosts but the subscriber hosts, both dynamic and static, will be maintained based on the lease time.

In the routed CO model, the router acts as a DHCP relay agent and also serves as the subscriber-identification agent. The DHCP actions are defined in the group-interface. All SAPs in that interface inherit these definitions. The group-interface DHCP definition are a template for all SAPs.

Lease-populate is enabled by default with the number-of-entries set to 1. This enables DHCP lease state for each SAP in the group-interface.

Since the group-interface can aggregate subscribers in different subnets a GI address must be defined for the DHCP relay process. The address must be in one of the host addresses defined for the subscriber interface. The GI address can be defined at the subscriber interface level which will

cause all child group interface to inherit that route. The GI address can then be overridden at the group interface level. A GI address must be defined in order for DHCP relay to function.

Because of the nature of the group-interface, local-proxy-arp, as well as arp-populate, should be enabled. This would allow the system to respond to subscriber ARP requests if the ARP request contains an IP address which is in the same subnet as one of the subscriber interface subnets.

When an authentication policy is specified for a SAP under a group interface, DHCP will intercept DHCP discover messages for RADIUS authentication. If the system is a DHCP-relay defined in a group-interface and the GI address was not configured the operational state of DHCP will be down.

## Routed CO for VPRN Service

Much like in Routed CO for IES service, the Routed CO model for VPRN depends on subscriber management to maintain the subscriber host information. To create a group-interface, the operator must first create a subscriber interface in the config>service>vprn context. The subscriber interface can maintain up to 256 subscriber subnets and can be configured with a host address for each subnet. The host IP address can be installed as a result of both relaying to a DHCP server and proxy to a RADIUS server. In both cases the host IP address must be in the subnet defined by the VPRN's subscriber interface.

Basic subscriber management is allowed only in a subscriber/SAP model and can be used only in a dedicated VPRN architecture. A RADIUS service selection (using Managed SAPs) will require Enhanced Subscriber Management. The subscriber interface's subnets are allowed to be advertised to both IGP and BGP within a VPRN.

When an authentication policy is specified for a group-interface, DHCP snooping must be enabled to intercept DHCP discover and renew messages for RADIUS authentication. Subscriber management RADIUS extensions are allowed if the operator chooses to have the RADIUS server return the subscriber identification, subscriber profile and sla-profile strings using RADIUS.

The node can be defined with both a DHCP relay or proxy function. If the user configures a DHCP relay, the local-proxy-server command will enable DHCP split leases. In that configuration the node will provide the configured DHCP lease to the client using either RADIUS or the real DHCP server as the source of the IP address to be provided.

The RADIUS server can send a Change of Authorization (CoA) message containing the DHCP forcerenew VSA which prompts the local-proxy-server to send a forcerenew message to the client. The node ACKs when the Force-Renew has been sent, regardless of whether the subscriber responds. If the client fails to respond or if a new session cannot be established due to resource management issues or otherwise the node must respond with a NACK to the RADIUS server.

If the CoA message contains an IP address that is different than the configured IP address (when RADIUS was providing IP addresses) the node must send a forcerenew message to the client and NAK the request and provide a new IP address. If the node fails to receive a request the CoA is ACK'd when the force-renew has been sent

The operational state of group and subscriber interfaces are dependent on the state of active SAPs. A group interface can become operationally up only if at least one SAP is configured and is in an operationally up state. A subscriber interface becomes operationally up if at least one group interface is operationally up or the associated wholesale forwarding interface is operationally up. This ensures that, in a failure scenario that affects all group interfaces in a given subscriber subnet, the node will stop advertising the subnet to the network. The SRRP state will affect this behavior as well and can cause the subnet to be removed if all group interfaces (and SRRP instances) are in backup state.



## Wholesale Retail Routed CO

VPRN Routed CO allows a provider to resell wholesaler services (from a carrier) while providing direct DSLAM connectivity. An operator can create a VPRN service for the retailer and configure the access from subscribers as well as to the retailer network. Any further action will be as if the VPRN is a standalone router running the Routed CO model. All forwarding to these servers must be done within the VPRN service. The operator can leak routes from the base routing instance. In this model, the operator can use RADIUS for subscriber host authentication, DHCP relay and DHCP proxy. This provides maximum flexibility to the retailer while minimizing the involvement of the wholesaler. Access cannot be shared among retailers unless a subscriber SAP is used. This requires that the wholesaler maintain a different access node (DSLAM) for each retailer that does not scale well. The wholesale retail model described in this section overcomes these limitations.

## Wholesale Retail Model

In the wholesale retail model (Figure 89), the wholesaler instance connections that are common to the access nodes are distributed to many retail instances. A subscriber host attached to an access node connected in the wholesaler service can be instantiated in a retail service and obtain IP addresses from the retailers address space. The service context of the retailer is determined during the subscriber host authentication phase (for example via the Alc-Retail-Serv-Id attribute in RADIUS or the retail-service-id CLI in the local user database).

Upstream subscriber traffic ingresses into the wholesaler instance and after identification is then forwarded into the retail instance. The reverse will occur for traffic in the downstream direction.

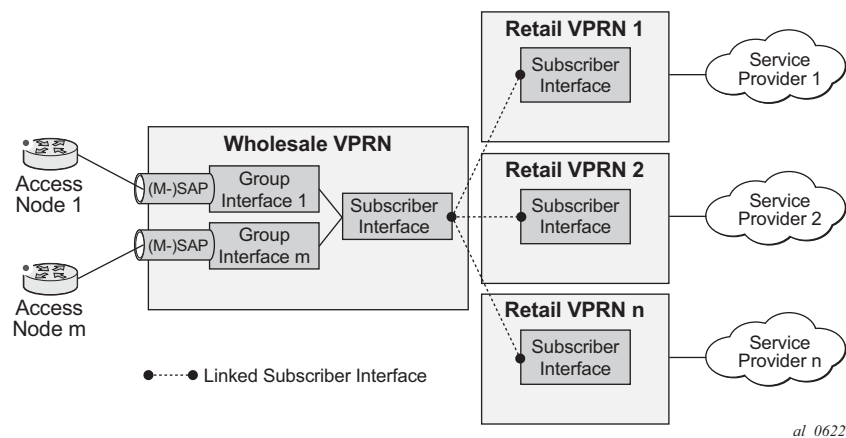


Figure 89: Wholesale Retail Model

In a wholesale retail model, two subscriber interfaces must be configured and linked together: one in the wholesale VPRN and one in the retail service.

The wholesale subscriber interface defines the IP subnets and host specific configuration parameters for subscriber hosts belonging to the wholesaler. There are associated group interfaces that contain the SAPs which connect to the access nodes.

The retail subscriber interface defines the IP subnets and host specific configuration parameters for subscriber hosts belonging to the retailer. The retail subscriber interface is linked to a wholesale subscriber interface for forwarding via explicit configuration. There are no associated group interfaces.

For example:

```
config>service
    vprn 1000 customer 1 create
        subscriber-interface "sub-int-ws-1" create
        # wholesale subscriber interface
        --- snip ---
        group-interface "group-int-1-1" create
        --- snip ---
        sap 1/1/1:1 create
        --- snip ---
        exit
    exit
exit

vprn 1001 customer 1 create
    subscriber-interface "sub-int-rt-1" fwd-service 1000  \\
        fwd-subscriber-interface "sub-int-ws-1" create
    # linked retail subscriber interface
    --- snip ---
    exit
exit
```

A retail subscriber interface can be linked to a single wholesale subscriber interface and context only. Subscriber interface chaining (linking a retail subscriber interface to another retail subscriber interface) is not supported. Multiple retail subscriber interfaces belonging to different retail contexts can be associated with a single wholesale subscriber interface. When a retail subscriber interface is linked to a wholesale context, all other retail subscriber interfaces from the same retailer must be linked to the same wholesale context.

## Configuration and Applicability

As explained in the previous section, the wholesale retail model is provisioned with the linking of a subscriber interface in a retail service to a subscriber interface in the wholesale VPRN service.

Because a retail subscriber interface does not have a group-interface context, some group-interface specific CLI parameters such as to configure dhcp relay are made available at the retail subscriber interface level. Other CLI parameters such as to provision Radius or local user database authentication are configured at the wholesale subscriber or group interface and apply to both wholesale and retail subscriber hosts.

The DHCP lease-populate configuration is special in wholesale retail as it is configured in both wholesale and retail context. The lease-populate value in the wholesale group-interface dhcp context controls the per SAP limits while the lease-populate value configured in the retail subscriber interface dhcp context controls the limits for the retailer subscriber interface. Both limits must be satisfied before a new subscriber host can be instantiated.

The sample configurations below enable dual stack IPoE devices to connect to wholesale service VPRN 4000 and retail service VPRN 4001. Hosts connected in VPRN 4000 get their IP address assigned from Radius, hence the proxy server configuration. Hosts connected in VPRN 4001 get their IP address from a DHCP server, hence the DHCP relay configuration.

Only the service configurations are shown. They have to be completed with authentication policies and subscriber management configuration such as radius-server-policies, sub- and sla-profiles, etc.

### Sample configuration – Wholesale VPRN Service:

```
config>service
    vprn 4000 customer 1 create
        autonomous-system 64500
        route-distinguisher 64500:4000
        auto-bind-tunnel
            resolution-filter
                ldp
                rsvp
            exit
        resolution filter
    exit
    vrf-target target:64500:4000
    subscriber-interface "sub-int-1" create
        address 10.10.1.254/24
        address 10.10.2.254/24
        ipv6
            delegated-prefix-len variable
            subscriber-prefixes
                prefix 2001:db8:a:100::/56 wan-host
                prefix 2001:db8:a001::/48 pd
            exit
        exit
    group-interface "group-int-1" create
        ipv6
            router-advertisements
                no shutdown
```

## Routed Central Office (CO)

```
        exit
        dhcp6
            proxy-server
                no shutdown
            exit
        exit
    exit
    arp-populate
    dhcp
        proxy-server
            emulated-server 10.10.1.254
            no shutdown
        exit
        lease-populate 100
        no shutdown
    exit
    authentication-policy "auth-policy-1"
    sap 1/1/4:1201.27 create
        sub-sla-mgmt
            sub-ident-policy "sub-ident-1"
            multi-sub-sap 100
            no shutdown
        exit
    exit
    exit
    exit
    no shutdown
exit
```

### Sample configuration – Retail VPRN Service:

```
config>service>
    vprn 4001 customer 1 create
        autonomous-system 64501
        route-distinguisher 64500:4001
        auto-bind-tunnel
            resolution-filter
                ldp
                rsvp
            exit
            resolution filter
        exit
        vrf-target target:64500:4001
        interface "int-loopback-1" create
            address 192.0.2.5/32
            ipv6
                address 2001:db8::5/128
            exit
            loopback
        exit
        subscriber-interface "sub-int-rt-4000-1" fwd-service 4000 fwd-subscriber-inter-
face "sub-int-1" create
            address 10.10.11.254/24
            address 10.10.12.254/24
            dhcp
                server 192.0.2.4
                lease-populate 100
                gi-address 10.10.11.254
```

```

        no shutdown
    exit
    ipv6
        subscriber-prefixes
            prefix 2001:db8:b:100::/56 wan-host
            prefix 2001:db8:b001::/48 pd
        exit
        dhcp6
            relay
                source-address 2001:db8::5
                server 2001:db8::4
                no shutdown
            exit
        exit
        router-advertisements
            no shutdown
        exit
    exit
exit
no shutdown
exit

```

The wholesale retail model applies to all IPoE, PPPoE PTA, IPv4 and IPv6 host types.

The wholesale service type must be VPRN. For IPEv4 hosts, the retail service type must be a VPRN. For all other host types, the retail service type can be IES or VPRN.

Multicast per host replication can be enabled without support for multi-chassis redundancy.

The wholesale retail model can be deployed in combination with managed SAPs

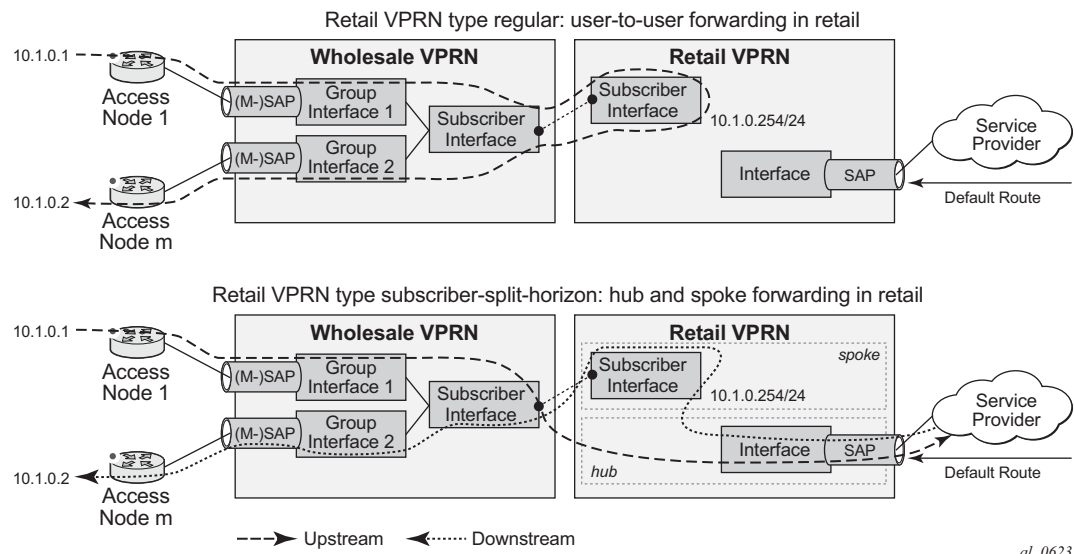
## Hub-and-Spoke Forwarding

In some cases, hub-and-spoke-type forwarding is needed for the retailer's VPRN. When the retailer expects all subscriber traffic to reach its router (for accounting, monitoring, wiretapping, etc.) normal best-hop behavior within the retailer VPRN is not desired. Any subscriber-to-subscriber traffic will be forwarded within the VPRN preventing the retailer from receiving these packets. To force all subscriber packets to the retailer network a new type of hub-and-spoke topology is defined: "type subscriber-split-horizon". It can be used to force all subscriber traffic (upstream) to the retailers network. The system requires that the operator will shutdown the VPRN service to enable this flag.

With retail VPRN type configured to subscriber-split-horizon, routes learned from MBGP, IGP through a regular interface, static routes through regular interfaces and locally attached regular interface routes are considered hub routes and are used for upstream traffic forwarding. Subscriber subnets cannot be used for upstream traffic forwarding. Downstream traffic will use routes in both hub and spoke routing instances.

[Figure 90](#) shows user to user traffic forwarding for both retail VPRN type regular and subscriber-split-horizon.


Routed Central Office (CO)



al\_0623

Figure 90: Wholesale Retail – Hub and Spoke Forwarding

## Routed Subscriber Hosts



**Figure 91: Router Subscriber Hosts**

A routed subscriber host associated route is a global routable subnet/prefix behind a routed CPE or Home Gateway. The routed CPE is identified in the BNG as an ESM subscriber host: QoS, accounting and anti-spoofing is enforced per CPE. The associated routes are installed in the BNG route table with next-hop pointing to the routed subscriber host's WAN address.

Routed subscriber host associated routes are supported on IES/VPRN subscriber interfaces in a routed CO configuration. To put a SAP or MSAP in routed subscriber mode, the anti-spoof type for the SAP or MSAP must be configured to nh-mac:

```
configure
  service ies/vprn <service-id>
    subscriber-interface <ip-int-name>
    group-interface <ip-int-name>
    sap <sap-id>
      anti-spoof nh-mac

configure
  subscriber-mgmt
    msap-policy <msap-policy-name>
    ies-vprn-only-sap-parameters
      anti-spoof nh-mac
```

There are three ways to learn about a routed subscriber host associated IPv4 route:

1. Configuration for a static host
2. A dynamic BGP peer
3. The RADIUS [22] Framed-Route attribute

A routed subscriber host associated IPv6 route can only be learned with the RADIUS [99] Framed-IPv6-Route attribute.

---

### Static Configured IPv4 Managed Route

The routes associated with a static host are populated in the routing table as “Remote Managed” routes. Up to sixteen managed routes can be configured for a static host.

```
config>service>ies>sub-if>grp-if>sap
config>service>vprn>sub-if>grp-if>sap

static-host ip 10.1.1.20 create
sla-profile "sla-profile-1"
sub-profile "sub-profile-1"
subscriber "static-host-1"
managed-routes
route 172.20.1.0/24
. . .
route 172.20.16.0/24
exit
no shutdown
exit
```

To display the managed routes associated with a routed subscriber host, use following commands:

**show service id *service-id* static-host detail**

---

### Static Configured IPv6 Managed Route

The routes associated with a static host are populated in the routing table as “Remote Managed” routes. Up to sixteen managed routes can be configured for a static host.

```
config>service>ies>sub-if>grp-if>sap
config>service>vprn>sub-if>grp-if>sap
anti-spoof nh-mac
static-host ip 2001::1/128 create
sla-profile "sla-profile-1"
sub-profile "sub-profile-1"
subscriber "static-host-1"
managed-routes
route 2000::/56
. . .
route 3000::/56
exit
no shutdown
exit

static-host ip 2001:1::/64 create
sla-profile "sla-profile-1"
sub-profile "sub-profile-1"
subscriber "static-host-1"
managed-routes
route 4000::/56
```



```

        . . .
        route 5000::/56
    exit
    no shutdown
exit

```

To display the managed routes associated with a routed subscriber host, use following commands:

**show service id *service-id* static-host detail**

---

## Dynamic BGP Peering

Routed subscriber host associated IPv4 routes can be learned over a dynamic BGP IPv4 peer that is automatically set up when a subscriber host is instantiated. The parameters for the BGP peer are configured in a BGP peering policy or obtained in Radius VSA attributes. The subscriber-host IPv4 address is used as the BGP peer IP address. The BGP peering is torn down and the associated routes removed from the routing table as soon as the subscriber-host is removed.

Dynamic BGP peering is supported for routed subscriber hosts terminated in a VPRN service and is not supported for routed subscriber hosts terminated in an IES service. The BGP learned routes scaling is limited by the BGP scaling limits. The routes learned via a dynamic BGP peer are populated in the routing table as “Remote BGP” routes.

To display the BGP learned routes associated with a routed subscriber host, use the regular BGP commands. For example:

**show router *service-id* bgp neighbor *ip-address* received-routes**

A dynamic BGP group must be configured in the BGP cli context of the VPRN service where the subscriber host is started:

**config>service>vprn>bgp**

```

    group "dynamic-peer-1" dynamic-peer
exit

```

The BGP peering policy to be used must be configured in the subscriber-mgmt CLI context:

**config>subscr-mgmt**

```

    bgp-peering-policy "bgp-policy-1" create
exit

```

A dynamic BGP peer is established for a subscriber host if the RADIUS attribute [26-6527-55] “Alc-BGP-Policy” returned in the Access-Accept contains the name of a local configured bgp-peering-policy and if a dynamic peer group is configured in the VPRN BGP context.

BGP peering parameters can be specified from multiple sources:

- Use BGP peering parameters returned in Radius VSA attributes
- If not available from RADIUS, use BGP peering parameters configured in the bgp-peering-policy

- If not configured in the bgp-peering-policy, use BGP peering parameters configured for the dynamic-peer group
- If not configured in the dynamic-peer group, use the BGP peering parameters configured in the VPRN service BGP CLI context.
- If not configured in the VPRN service BGP CLI context, use the defaults

The import and export policies to be used for the dynamic bgp peer are determined in following priority order:

1. Use import/export policies returned in RADIUS VSA attributes and append policies configured in the bgp-peering-policy.
2. If not available from RADIUS AND not configured in the bgp-peering-policy, use the policies configured in the dynamic-peer group.
3. If not configured in the dynamic-peer group, use the policies configured in the VPRN service BGP CLI context.

[Table 14](#) details the RADIUS VSA attributes that can be used to setup dynamic BGP peering.

**Table 14: RADIUS VSA Attributes to Setup Dynamic BGP Peering**

Attribute-ID	Attribute Name	Description
26-6527-55	Alc-BGP-Policy	Mandatory attribute to setup a dynamic BGP peer. References a bgp peering policy configured in the “configure subscriber-mgmt bgp-peering-policy <policy-name>” CLI context.
26-6527-56	Alc-BGP-Auth-Keychain	Optional. References a keychain configured in the “configure system security keychain <keychain-name>” CLI context.
26-6527-57	Alc-BGP-Auth-Key	Optional. The MD5 authentication key used between BGP peers for BGP session establishment.
26-6527-58	Alc-BGP-Export-Policy	Optional. References a pre-configured BGP export routing policy.
26-6527-59	Alc-BGP-Import-Policy	Optional. References a pre-configured BGP import routing policy.
26-6527-60	Alc-BGP-PeerAS	Optional. Specifies the Autonomous System number for the remote peer

## RIP Listener

If a routed subscriber host is associated with a RIP policy, the host's IPv4 routes can be learned over RIP. The BNG only supports RIP listener and does not support sending RIP routes to subscribers. To enable RIP for a subscriber, the subscriber must first be associated with a **rip-policy**. The group interface of the subscriber must also be configured as a RIP neighbor. The rip-policy can be associated to the subscriber during authentication from LUDB or via Radius. It can also be configured directly for static hosts. The RIP routes learned from a subscriber is removed as a subscriber is purged or shutdown from the system. RIP listening for ESM host is supported on both IES and VPRN.

To display the RIP learned routes associated with a routed subscriber host, use the regular RIP commands. For example:

```
show router service-id rip neighbor interface advertised-routes
```

The group interface must be configured in the RIP CLI context of the routed instance where the subscriber host is created:

```
config>router/service vprn>rip
      group "rip-listener"
      neighbor "group-interface-01"
```

The RIP policy is configured in the subscriber-mgmt CLI context:

```
config>sub-mgmt
      rip-policy "rip-policy-01" create
```

A RIP neighbor is established for a subscriber host if the RADIUS attribute [26-6527-207] "Alc-RIP-Policy" is returned in the Access-Accept or in LUDB. RIP parameters such as **authentication key** and **type** can be specified in the RIP policy.

For more information about RIP, refer to the 7x50 SR OS Routing Protocols Guide.

## RADIUS: Framed-Route and Framed-IPv6-Route

RADIUS attribute [22] Framed-Route can be specified in a Radius Access-Accept message to associate an IPv4 route with an IPv4 routed subscriber host and Radius attribute [99] Framed-IPv6-Route can be used to associate an IPv6 route with an IPv6 routed subscriber wan host (DHCPv6 IA-NA or SLAAC). These routes are populated in the routing table as “Remote Managed” routes. Up to sixteen managed routes can be installed for a routed subscriber host; this corresponds with up to sixteen Framed-Routes and sixteen Framed-IPv6-Routes for a dual stack routed subscriber. Framed-IPv6-Routes cannot be associated with a Prefix Delegation host (DHCP IA-PD).

The Framed-Route and Framed-IPv6-Route attributes should be formatted as:

```
"<ip-prefix>[/<prefix-length>] <space> <gateway-address> [<space> <metric>] [<space> tag  
<space> <tag-value>] [<space> pref <space> <preference-value>]"
```

where:

<space> — is a white space or blank character.

<ip-prefix>[/prefix-length] — is the managed route to be associated with the routed subscriber host. The prefix-length is optional for an IPv4 managed route. When not specified, a class-full class A,B or C subnet is assumed. The prefix-length is mandatory for an IPv6 managed route.

<gateway-address> — must be the routed subscriber host IP address.

“0.0.0.0” is automatically interpreted as the host IPv4 address for managed IPv4 routes.

“::” and “0:0:0:0:0:0:0:0” are automatically interpreted as the wan-host IPv6 address for managed IPv6 routes.

[<metric>] — Optional. Installed in the routing table as the metric of the managed route. If not specified, metric zero is used. Value = [0.. 65535].

[tag <tag-value>] — Optional. The managed route will be tagged for use in routing policies. If not specified, or tag-value = 0, then the route is not tagged. Value = [0..4294967295].

[pref <preference-value>] — Optional. Installed in the routing table as protocol preference for this managed route. If not specified, preference zero is used. Value = [0..255].

If the optional metrics (metric, tag and/or preference) are specified in a wrong format or with out of range values, then the defaults are used for all metrics: metric=0, no tag and preference=0. No event is logged.

If the Framed-Route or Framed-IPv6-Route is invalid (for example because the gateway address specified does not match the host wan IP address or because the host bits are not zero) then the routed subscriber host is instantiated without the ill defined managed route. An event is logged in this case.

Equal Cost Multi-Path (ECMP) is supported for IPv4 Framed-Route:

The maximum number of equal cost paths in a routing instance is configured with:

```
config>router>  
config>service>vprn>
```

```
ecmp <max-ecmp-routes>
```

If an identical managed route is associated with different routed subscriber hosts in the context of the same IES/VP RN service, up to *<max-ecmp-routes>* managed routes are installed in the routing table. Candidate ECMP Framed-Routes have:

- Identical prefix
- Equal lowest preference
- Equal lowest metric

A tie breaker determines if more candidate ECMP Framed-Routes are available than the configured *<max-ecmp-routes>* is: Lowest ip next-hop.

Other identical managed routes are shadowed and an event is logged.

Note that Candidate ECMP Framed-Routes must not belong to the same subscriber.

Equal Cost Multi-Path (ECMP) is not supported for Framed-IPv6-Route.

If an identical managed IPv6 route is associated with different routed subscriber hosts in the context of the same IES/VP RN service only one managed IPv6 route is installed in the routing table. The selection criteria are (in order of priority):

1. Lowest preference
2. Lowest metric
3. Lowest ip next-hop

Other identical managed IPv6 routes are shadowed and an event is logged. Valid Framed-Routes and Framed-IPv6-Routes are persistent (stored in the persistency file for recovery after reboot) and synchronized in a Multi-Chassis Redundancy configuration.

RADIUS-learned Framed-Route/Framed-IPv6-Route and static host associated managed routes that are installed in the routing table can be identified in routing policies for redistribution as protocol “managed”.

To display the managed routes associated with a routed subscriber host, use following commands:

```
show service id service-id dhcp lease-state detail
```

```
show service id service-id dhcp6 lease-state detail
```

```
show service id service-id slaac host detail
```

```
show service id service-id ppp session detail
```

```
show service id service-id pppoe session detail
```

```
show service id service-id arp-host detail
```

Valid RADIUS-learned managed routes can be included in Radius accounting messages with following configuration:

```
configure
  subscriber-mgmt
```

## Routed Central Office (CO)

```
radius-accounting-policy <name>  
    include-radius-attribute  
        framed-route  
        framed-ipv6-route
```


Associated managed routes for an instantiated routed subscriber host are included in RADIUS accounting messages independent of the state of the managed route (Installed, Shadowed, HostInactive, etc.)

In case of a PPP session, when a Framed-Route or Framed-IPv6-Route is available while the corresponding routed subscriber host is not yet instantiated, the managed route is in the state “notYetInstalled” and will not be included in Radius accounting messages.

## Dual Homing

All residential networks are based on two models: Layer 2 CO and Layer 3 CO. Dual homing methods for Layer2 CO include MC-LAG and MC-Ring. Dual homing for Layer 3 CO is based on SRRP and can be done in ring-topologies (I3-mc-ring or with directly attached nodes). All methods use multi-chassis synchronization protocol to sync subscriber state.

### Dual Homing to Two PEs (Redundant-Pair Nodes) in Triple Play Aggregation



Fig\_40

**Figure 92: Dual-Homing to Two PEs**

Figure 92 depicts dual-homing to two different PE nodes. The actual architecture can be based on a single DSLAM having two connections to two different PEs (using MC-LAG) or ring of DSLAMs dual-connected to redundant pair of PEs.


Similarly to previous configuration, both aggregation models (VLAN-per-subscriber or VLAN-per-service) are applicable.

Configurations include:

- Loop resolution and failure recovery — Can be based on MC-LAG or mVPLS.
- DHCP-lease-state persistency — Stores all required information to recover from node failure.
- DHCP-lease-state synchronization — A mechanism to synchronize the DHCP lease-state between two PE nodes in the scope of redundancy groups (a group of SAPs used for dual homing).
- IGMP snooping state synchronization — Similarly to DHCP lease-state synchronization, IGMP snooping state is synchronized to ensure fast switchover between PE nodes. In a VPLS network, a BTV stream is typically available in all PE nodes (the ring interconnecting all PEs with Mcast routers is typically used) so the switch over can be purely driven by RSTP or MC-LAG.
- ARP reply agent responses — The ARP reply agent can response to ARP requests addressing a host behind the given SAP if the SAP is in a forwarding state. This prevents the FDB table in the VPLS from being “poisoned” by ARP responses generated by the node with a SAP in a blocking state (see [Figure 93](#)).

[Figure 93](#) shows a typical configuration of network model based on Layer 2 CO model. Individual rings of access nodes are aggregated at BSA level in one (or multiple) VPLS services. At higher aggregation levels (the BSR), individual BSAs are connected to Layer 3 interfaces (IES or VPRN) by spoke SDP termination. Every Layer 3 interface at BSR level aggregates all subscribers in one subnet.





Fig\_39

**Figure 93: Layer 2 CO Dual Homing - Network Diagram**

Typically, BTV service distribution is implemented in a separate VPLS service with a separate SAP per access-node. This extra VPLS is not explicitly indicated in [Figure 93](#) (and subsequent figures) but the descriptions refer to its presence.

From a configuration point of view in this model, it is assumed that all subscriber management features are enabled at the BSA level and that synchronization of the information (using multi-chassis synchronization) is configured between redundant pair nodes (BSA-1 and BSA-2 shown in [Figure 93 on page 965](#)). The multi-chassis synchronization connection is used only for synchronizing active subscriber host database and will operate independently from dual-homing connectivity control. At the BSR level, there are no subscriber management features enabled.

The operation of redundancy at the BSR level through VRRP is the same as dual homing based on MC-LAG. The operation of dual homing at BSA level is based on two mechanisms. Ring control connection between two BSAs have two components, in-band and out-of-band communication. With in-band communication, BFD session between BSA-1 and BSA-2 running through the access ring and using dedicated IES/VP RN interface configured on both nodes. This connection uses a separate VLAN throughout the ring. The access nodes provides transparent bridging for this

VLAN. The BFD session is used to continuously verify the integrity of the ring and to detect a failure somewhere in the ring.

With out-of-band communication, the communication channel is used by BSA nodes to exchange information about the reachability of individual access nodes as well as basic configurations in order to verify the consistency of the ring. The configuration information is synchronized through multi-chassis synchronization and therefore it is mandatory to enable multi-chassis synchronization between two nodes using the multi-chassis-ring concept.


In addition, the communication channel used by MC-LAG or MC-APS control protocol is used to exchange some event information. The use of this channel is transparent to the user.

Ring node connectivity check continuously checks the reachability of individual access nodes in the ring. The session carrying the connection is conducted on separate VLAN, typically common for all access nodes. SHCV causes no interoperability problems.

### Steady-State Operation of Dual-homed Ring

Figure 94 illustrates the operation of the dual-homed ring. The steady state is achieved when both nodes are configured in a consistent way and the peering relation is up. The multi-chassis ring must be provisioned consistently between two nodes.

In-Band Ring Control Connection (IB-RCC) is in operational UP state. Note that this connection is set up using a bi-directional forwarding session between IP interfaces on BSA-1 and BSA-2.



**Figure 94: Dual Homing Ring Under Steady-State Condition**

In [Figure 94](#), the ring is fully closed and every access node has two possible paths towards the VPLS core. [Figure 94](#) refers to these as **path-a** and **path-b**. In order to avoid the loop created by the ring, only one of the paths can 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.

The selection of the BSA master for both paths will be based on the IP address of the interface used for IB-RCC communication (bi-directional forwarding session). The BSA with the lower IP address of the interface used as IB-RCC channel will become master for ring nodes and their respective VLANs assigned to path-a. The master of path-b will be other BSA.

In this example, each path in the ring has a master and standby BSA. The functionality of both devices in steady state are as follows:

In the master BSA:

- All SAPs that belong to the path where the given BSA is a master, are operationally UP and all FDB entries of subscriber hosts associated with these SAPs point to their respective SAPs.
- The master of a path performs periodical Ring Node Connectivity Verification (RNCV) check to all ring nodes.
- In case of a RNCV failure, the respective alarm will be raised. Note that the loss of RNCV to the given ring node does not trigger any switchover action even if the other BSA appears to have the connection to that ring node. As long as the BFD session is up, the ring is considered closed and the master/standby behavior is driven solely by provisioning of the individual paths.
- The ARP reply agent replies to ARP requests addressing subscriber hosts where the BSA master.

In the standby BSA:

All SAPs that belong to a BSA's path, the standby will be operationally down and all FDB entries of subscriber hosts associated with those SAPs will be pointing towards SDP connecting to master BSA (also called a shunt SDP).

In both BSAs:

- The information on individual paths assignment is exchanged between both BSAs through multi-chassis synchronization communication channel and conflicting SAPs (being assigned to different paths on both BSA nodes) will be forced to path-a (the default behavior).
- For IGMP snooping, the corresponding multi-chassis IDs are targeting all subscriber-facing SAPs on both nodes. On the standby BSA node, the corresponding SAPs are in an operationally down state to prevent the MC traffic be injected on the ring twice.


## Broken-Ring Operation and the Transition to this State

Figure 95 illustrates the model with a broken ring (link failure or ring node failure). This state is reached in following conditions:

- Both nodes are configured similarly.
- Peering is up.
- The multi-chassis ring is provisioned similarly between two nodes
- IB-RCC is operationally down.

In this scenario, every ring node has only one access path towards the VPLS core and hence, the Path-a and Path-b notion has no meaning in this situation.

Functionally, both BSAs are now the master for the reachable ring nodes and will take action as described in [Steady-State Operation of Dual-homed Ring on page 966](#). For all hosts behind the unreachable ring nodes, the corresponding subscriber host FDB entries point to the shunt SDP.



Fig\_38

**Figure 95: Broken Ring State**

The mapping of individual subscriber hosts into the individual ring nodes is complicated, especially in the VLAN-per-service model where a single SAP can represent all nodes on the ring. In this case, a given BSA can have subscriber hosts associated with the given SAP that are behind

reachable ring nodes as well as subscriber hosts behind un-reachable ring nodes. This means that the given SAP cannot be placed in an operationally down state (as in a closed ring state), but rather, selectively re-direct unreachable subscriber states to the shunt SDP.

All SAPs remain in an operationally up state as long as the ring remains broken. This mainly applies for BTV SAPs that do not have any subscriber hosts associated with and do not belong to any particular ring node.

In order to make the mapping of the subscriber-hosts on the given ring node automatically provisioned, the ring node identity will be extracted during subscriber authentication process from RADIUS or from a Python script. The subscriber hosts which are mapped to non-existing ring node will remain attached to the SAP.

At the time both BSA detect the break in IB-RCC communication (if BFD session goes down) following actions are taken:

- Both nodes trigger a RNCV check towards all ring nodes. The node, which receives the reply first, will assume a master functionality and will inform the other BSA through an out-of-band channel. This way, the other node can immediately take actions related to the standby functionality without waiting for an RNCV timeout. Even if the other node receives an RNCV response from the given ring node later, the master functionality remains with the node the received the response first.
- After assuming the master functionality for hosts associated with the given SAP(s), the node will send out FIB population messages to ensure that new path towards the VPLS core is established. The FIB population messages are sourced from the MAC address of the default gateway used by all subscriber hosts (such as the VRRP MAC address) which is provisioned at the service level.

## Transition from Broken to Closed Ring State

By its definition, the multi-chassis ring operates in a revertive mode. This means that whenever the ring connectivity is restored, the BSA with lower IP address in the IB-RCC communication channel will become master of the path-a and vice versa for path-b.

After restoration of BFD session, the master functionality, as described in [Steady-State Operation of Dual-homed Ring on page 966](#), is assumed by respective BSAs. The FDB tables are updated according to the master/standby role of the given BSA and FDB population messages is sent accordingly.

---

## Provisioning Aspects and Error Cases

The multi-chassis ring can operate only if both nodes similarly configured. The peering relation must be configured and both nodes must be reachable at IP level. The multi-chassis ring with a corresponding sync-tag as a ring-name identifying a local port ID must be provision on both nodes. And, BFD session and corresponding interfaces needs to be configured in a consistent way.


In case the multi-chassis rings are not provisioned consistently, the ring will not become operational and the SAP managed by it will be in operationally up state on both nodes.

The assignment of individual SAPs to path-a and path-b is controlled by configuration of VLAN ranges according to the following rules:

- By default, all SAPs (and hence all VLANs on the given port) are assigned to path-a.
- An explicit statement defining the given VLAN range assigns all SAPs falling into this range to the path-b.
- An explicit statement defining the given VLAN range defines all SAPs that are excluded from the multi-chassis ring control.
- In case of a conflict in the configuration of VLAN ranges between two redundant nodes is detected, all SAPs falling into the “conflict-range” will be assigned to path-a, on both nodes regardless the local configuration.
- For QinQ-encapsulated ports the VLAN range refers to the outer VLAN.

## Dual Homing to Two BSR Nodes

Figure 96 depicts a single DSLAM dual-homed to two BSRs.



**Figure 96: Low**


In order to provide dual-homing in the context of subscriber interfaces, the following items must be configured on both BSRs:

- Group interface (dslam-1) with corresponding SAPs (vlan 1-100)
- SRRP instance controlling given group interface
- Redundant interface between BSRs to provide “shunt” connectivity
- MCS connection to provide synchronization of dynamic subscriber-host entries

During the operation, BSR-1 and BSR-2 will resolve master-backup relation and populate respective FIBs in such a way that at master side, subscriber-host entries point to corresponding group-interface while at the back-up side, subscriber-host entries point to the redundant interface. Note that the logical operation of the ring in the Layer 3 CO model is driven by SRRP. For more details on SRRP operation, refer to [Subscriber Routed Redundancy Protocol \(SRRP\) on page 916](#).

## MC Services

The typical implementation of MC services at the network level is shown in [Figure 97](#).



**Figure 97: MC Services in a Layer 3-Ring Topology (a)**

The IGMP is used to register joins and leaves of the user. IGMP messaging between BSRs is used to determine which router performs the querier role (BSR2 in [Figure 97](#)). PIM is used to determine which router will be the designated router and the router that sends MC streams on the ring.

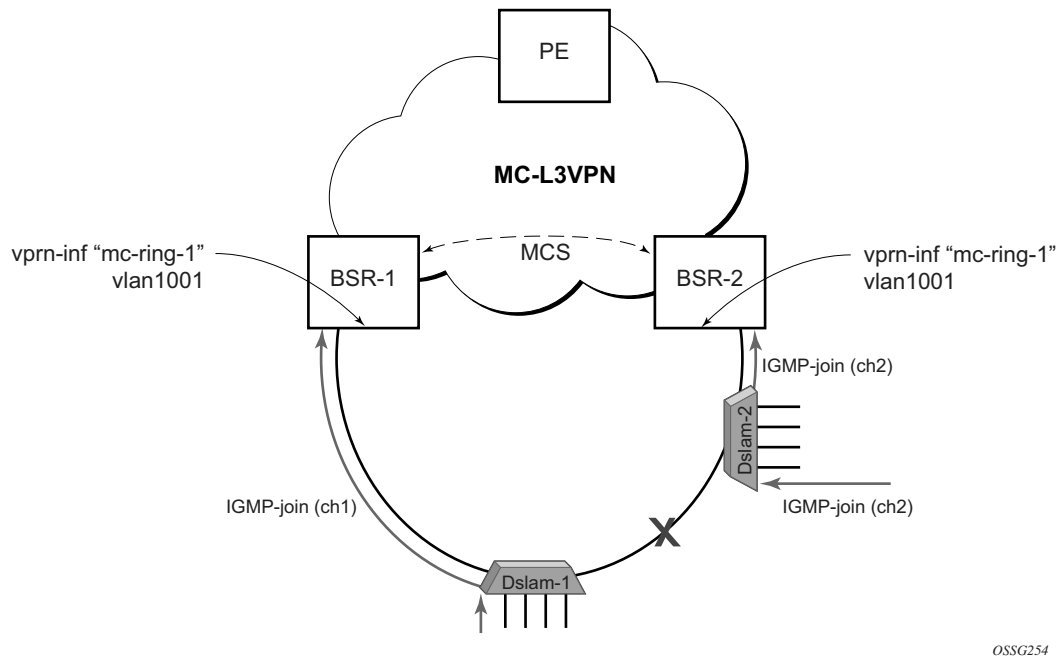
The access nodes have IGMP snooping enabled and from IGMP messaging between BSR, they are aware which router is the querier. In the most generic case, IGMP snooping agents (in access nodes) send the IGMP-joins messages only to IGMP-querier. The synchronization of the IGMP entries can be then be performed through MCS. In some cases, access nodes can be configured in such a way that both ring ports are considered as m-router ports and IGMP joins are sent in both directions.

All of the above is a steady state operation which is transparent to the topology used in a Layer 2 domain.

A ring-broken state is shown in [Figure 98](#).

In this case, IGMP and PIM messaging between BSRs is broken and both router assume role of querier and role of designated router. By the virtue of ring topology, both routers will see only IGMP joins/leaves generated by host attached to a particular “half” of the ring. This means that both routers will have “different” view on dynamic IGMP state.





**Figure 98: MC Services on a Layer 3-Ring Topology (b)**

In principle, MCS could be used to synchronize both routers, but in case of a Layer 2 ring, the implementation sends all IGMP messages to a “ring-master” which then performs IGMP processing and consequently, MCS sync. As a result, any race conditions are avoided.

Another ring-specific aspect is related to ring healing. The ring continuity check is driven by BFD which then drives SRRP and PIM messaging. BFD is optimized for fast detection of ring-down events while ring-up events are announced more slowly. There is a time window when routers are not aware that the ring is recovered. In the case of MC, this means traffic will be duplicated on the ring.

To avoid this, the implementation of BFD provides a “raw mode” which provides visibility on “ring-up” events. The protocols, such as SRRP and PIM, use this raw mode rather than the BFD API.

### Routed CO Dual Homing

Routed CO dual homing is a solution that allows seamless failover between nodes for all models of routed CO. In the dual homed environment, only one node will forward downstream traffic to a given subscriber at a time. Dual homing involves several components:

- Redundant Interface — This is used to shunt traffic to the active node for a given subscriber for downstream traffic.
- SRRP — This is used to monitor the state of connectivity to the DSLAM. Refer to the SRRP section for more detail.
- MCS — This is used to exchange subscriber host and SRRP information between the dual homed nodes.

Routed CO dual homing can be configured for both wholesaling models. Dual homing is configured by creating a redundant interface that is associated with the protected group interfaces. The failure detection mechanism can be VRRP. If VRRP is used, each node monitors the VRRP state to determine the priority of its own interface.

Dual homing is used to aggregate a large number of subscribers in order to support a redundancy mechanism that will allow a seamless failover between nodes. Because of the Layer 3 nature of the model, forwarding is performed for the full subscriber subnet.

---

### Redundant Interfaces

In dual homing, a redundant interface must be created. A redundant interface is a Layer 3 spoke SDP-based interface that allows delivery of packets between the two nodes. The redundant interface is required to allow a node with a failed link to deliver packets destined to subscribers behind that link to the redundant node. Since subscriber subnets can span multiple ports it is not possible to stop advertising the subnet, thus, without this interface the node would black hole.

The redundant interface is associated with one or more group interfaces. An interface in backup state will use the redundant interface to send traffic to the active interface (in the active node). The SAP structure under the group interface must be the same on both nodes as the synchronization of subscriber information is enabled on a group interface basis. Traffic can be forwarded through the redundant interface during normal operation even when there are no failed paths. See [Figure 99](#).

---

### SRRP in Dual Homing

Subscriber Router Redundancy Protocol (SRRP) allows two separate connections to a DSLAM to operate in an active/standby fashion similar to how VRRP interfaces operate. Since the SRRP state is associated with the group-interface, multiple group-interfaces may be created for a given port such that some of the SAPs will be active in one node and others active on the other node. While each SRRP pair is still allowed to be active/backup, the described configuration is allowed for load balancing between the nodes. Note that in a failure scenario subscriber bandwidth will be affected. For more information about SRRP, refer to [Subscriber Routed Redundancy Protocol](#)

(SRRP) on page 916.

If SRRP is configured before the redundant interface is up, and in backup state the router will forward packets to the access node via the backup interface but will not use the gateway MAC address. This applies to failures in the redundant interface as well. If the redundant interface exists and up the router will send downstream packets to the redundant interface and will not use the backup group interface.


In a dual homing architecture the nodes must be configured with SRRP to support redundant paths to the access node. The nodes must also be configured to synchronize subscriber data and IGMP state. To facilitate data forwarding between the nodes in case some of the ports in a given subscriber subnet are affected a redundant interface must be created and configured with a spoke. The redundant interface is associated with one or more group interfaces.

The service IDs for both the wholesale VPRN and the retailer VPRN must be the same in both nodes.

An interface in backup state will use the redundant interface to send traffic to the active interface (in the active node). The SAP structure under the group interface must be the same on both nodes as the synchronization of subscriber information is enabled on a group interface basis.

SRRP is associated a group-interface. Multiple group-interfaces can be created for a given port such that some of the SAPs will be active in one node and others active on the other node. While every SRRP pair is still allowed to be active/backup the described configuration will allow for load balancing between the nodes. Note that in a failure scenario subscriber bandwidth will be affected.

## Page 976






Figure 99: Dual Homing Example

### Synchronization

To establish subscriber state the nodes must synchronize subscriber information. Refer to the 7750 SR Basic Configuration Guide for multi-chassis synchronization configuration information. The operator must complete the configuration and the system must have data synchronized before the backup node may deliver downstream packets to the subscriber.

If dual homing is used with regular interfaces that run IGMP the nodes must be configured to synchronize the Layer 3 IGMP state.

The service IDs for both the wholesale VPRN and the retailer VPRN must be the same in both nodes.

---

### Wholesale-Retail Multi-Chassis Redundancy

Multi-Chassis Redundancy for a retail service is enabled with the SRRP and redundant interface configuration on the wholesale group-interface parented by the forwarding subscriber interface. The Multi-Chassis state (active or standby) of the retail subscriber host is determined from the SRRP state (master / non-master) of the group-interface that parents the SAP of the retail subscriber host. The retail service id must be equal on both nodes.

Sample wholesale service configuration:

```
vprn 3000 customer 1 create
  description "Wholesale service"
  route-distinguisher 64500:3000
  auto-bind mpls
  vrf-target import target:64500:3000
  redundant-interface "red-int-1" create
    address 192.168.100.0/31
    ip-mtu 1500
    spoke-sdp 12:3000 create
      no shutdown
    exit
  exit
subscriber-interface "sub-int-1" create
  address 10.1.1.253/24 gw-ip-address 10.1.1.254
  address 10.1.2.253/24 gw-ip-address 10.1.2.254
  group-interface "group-int-1-1" create
    dhcp
    --- snip ---
  exit
  redundant-interface "red-int-1"
  sap 1/1/6:1.4001 create
    description "SRRP 1 message path"
  exit
  srrp 1 create
    message-path 1/1/6:1.4001
    send-fib-population-packets outer-tag-only
    no shutdown
  exit
  pppoe
  --- snip ---
```

```

        exit
    exit
    group-interface "group-int-1-2" create
        dhcp
        --- snip ---
    exit
    redundant-interface "red-int-1"
    sap 1/1/6:2.4001 create
        description "SRRP 2 message path"
    exit
    srrp 2 create
        message-path 1/1/6:2.4001
        priority 50
        send-fib-population-packets outer-tag-only
        no shutdown
    exit
    pppoe
        --- snip ---
    exit
exit
no shutdown
exit

```

#### Sample retail service configuration:

```

vprn 3001 customer 1 create
    description "Retail service"
    route-distinguisher 64500:3001
    auto-bind mpls
    vrf-target target:64500:3001
    subscriber-interface "sub-int-rt-3000-1" fwd-service 3000 fwd-subscriber-inter-
face "sub-int-1" create
    address 10.1.11.253/24 gw-ip-address 10.1.11.254
    address 10.1.12.253/24 gw-ip-address 10.1.12.254
    dhcp
        --- snip ---
    exit
    pppoe
        --- snip ---
    exit
exit
no shutdown
exit

```

Overlapping addresses in retail subscriber interfaces (enabled with private-retail-subnets) cannot be used in combination with Multi-Chassis Redundancy. Retail subscriber subnets leaked in the wholesale service are needed to forward shunted traffic. The address of an unnumbered subscriber host (enabled with **unnumbered** *ip-int-name|ip-address* or **allow-unmatching-subnets** on the retail subscriber interface) is not contained in the subnets configured on the retail subscriber interface. The export of the retail subscriber host routes to the wholesale service must be explicitly enabled with the **export-host-routes** command:

## Dual Homing

```
vrpn 3001 customer 1 create
  subscriber-interface "sub-int-rt-3000-1" fwd-service 1000 fwd-subscriber-inter-
face "sub-int-1" create
  allow-unmatching-subnets
  address 10.1.11.253/24 gw-ip-address 10.1.11.254
  address 10.1.12.253/24 gw-ip-address 10.1.12.254
  export-host-routes
  --- snip ---
```

Multi-Chassis Redundancy is supported for IPoEv4 and PPPoEv4 retail subscriber hosts.



## SRRP and Multi-Chassis Synchronization

In order to take full advantage of SRRP resiliency and diagnostic capabilities, the SRRP instance will be tied to a MCS peering that terminates on the redundant node. Once the peering is associated with the SRRP instance, MCS will synchronize the local information about the SRRP instance with the neighbor router. MCS automatically derives the MCS key for the SRRP instance based on the SRRP instance ID. An SRRP instance ID of 1 would appear in the MCS peering database with a MCS-key srrp-0000000001.

The SRRP instance information stored and sent to the neighbor router consists of:

- The SRRP instance MCS key
- Containing service type and ID
- Containing subscriber IP interface name
- Subscriber subnet information
- Containing group IP interface information
- The SRRP group IP interface redundant IP interface name, IP address and mask
- The SRRP advertisement message SAP
- The local system IP address (SRRP advertisement message source IP address)
- The group IP interface MAC address
- The SRRP gateway MAC address
- The SRRP instance administration state (up/down)
- The SRRP instance operational state (disabled/becoming-backup/backup/becoming-master/master)
- The current SRRP priority
- Remote redundant IP interface availability (available/unavailable)
- Local receive SRRP advertisement SAP availability (available/unavailable)

---

## Dual Homing and ANCP

Alcatel-Lucent provides a feature related to exchange of control information between DSLAM and BRAS (BSA is described in this model). This exchange of information is implemented by in-band control connection between DSLAM and BSA, also referred to as ANCP connection.

In case of dual homing, two separate connections will be set. As a consequence, there is no need to provide synchronization of ANCP state. Instead every node of the redundant-pair obtains this information from the DSLAM and creates corresponding an ANCP state independently.

## SRRP Enhancement

The SRRP enhancements addressed in this section is to reduce the need for redundant-interface between the pair of redundant nodes without sacrificing the subnet aggregation on the back-end.

Redundant BNG nodes are not always collocated. This means that the logical link associated with the redundant (shunt) interfaces is taking the uplink path thus wasting valuable bandwidth (downstream traffic that arrives to the Standby node is routed via uplinks for the second time over to the Master node).

To meet the requirement to reduce the existence of shunted traffic only to the short transitioning period between SRRP switchovers while the routing on the network side is converging, the following was required (referring to [Figure 100](#)):

1. Share IP subnets over multiple SRRP instances. This is not mandatory, but it would help to load balance traffic over the two nodes. For example, IP subnets 10 and 11 can be shared over SRRP instances 10 and 20 on node 1, and the IP subnet 12 can be associated with the SRRP instance 30 on node 2.
2. *SRRP aware routing* – this allows to dynamically increase routing metric on the IP subnets advertised from the Master SRRP node in comparison to the Standby SRRP node. It also allows to advertise/withdraw routes from a routing protocol based on the SRRP state. In this fashion, downstream traffic is routed in a predictable manner towards the Master SRRP node.
3. *SRRP Fate Sharing* for SRRP instances 10 and 11. This ensures congruency of SRRP states on the same node. This is a necessary step towards *SRRP aware routing*.

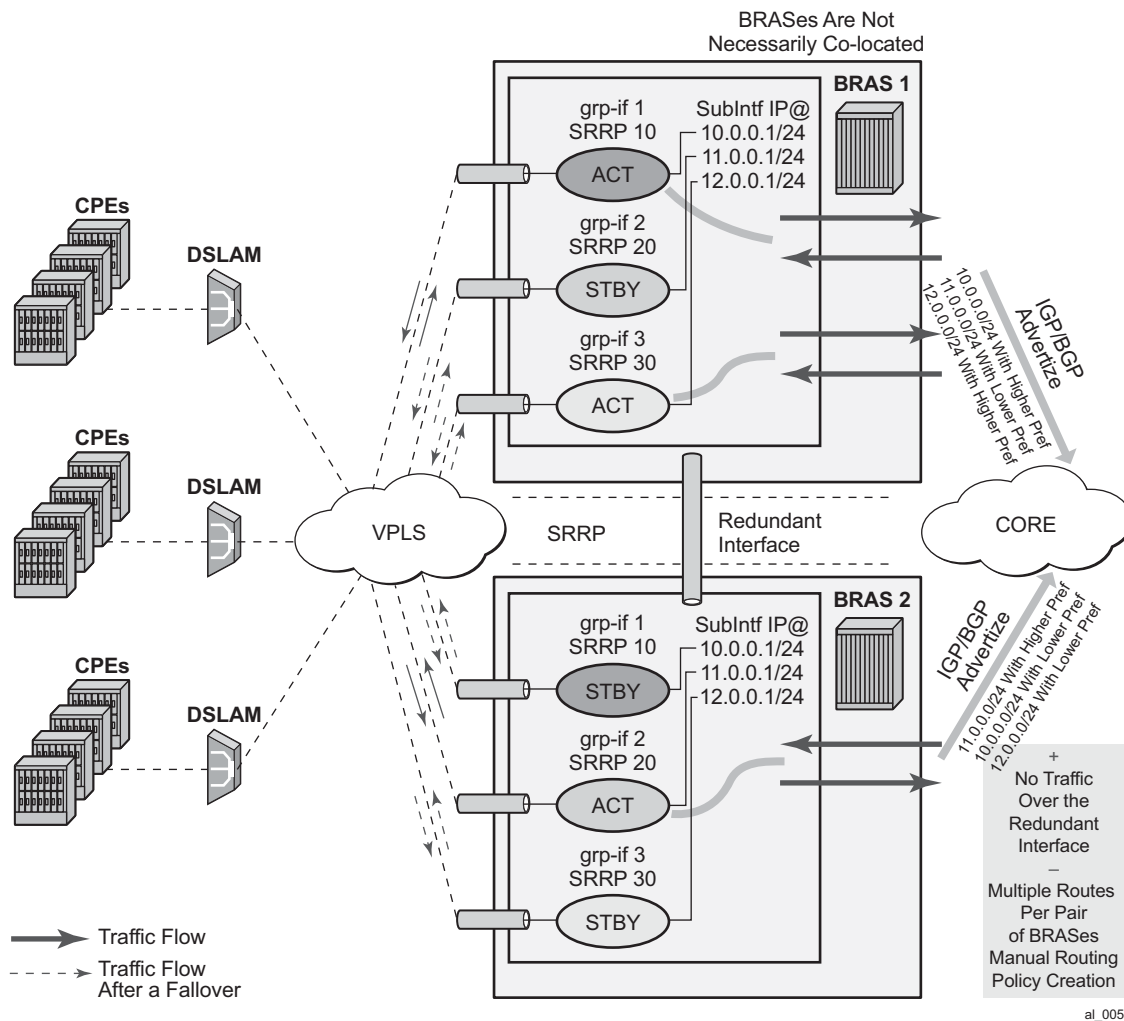


Figure 100: IP Subnet Per SRRP Master Group

## SRRP Fate Sharing

SRRP Fate Sharing is a concept in which a group of SRRP instances track a single operational-object comprised of SRRP messaging SAPs. The SRRP instances behave as one (in the single failure case) with regards to SRRP mastership. The group of SRRP instances that are sharing fate on a paired node are referred as a Fate Sharing Group (FSG).

Transition of a single messaging SAP within the FSG into a DOWN state forces the SRRP instance on top of it into the INIT state. Consequently, all other SRRP instances within the same FSG transitions into a Backup state. In other words, SRRP instances within the FSG all share the same fate as the failed SRRP instance as shown in Figure 101. SRRP Fate Sharing provides

optimal protection in the context of a single failure in the network.

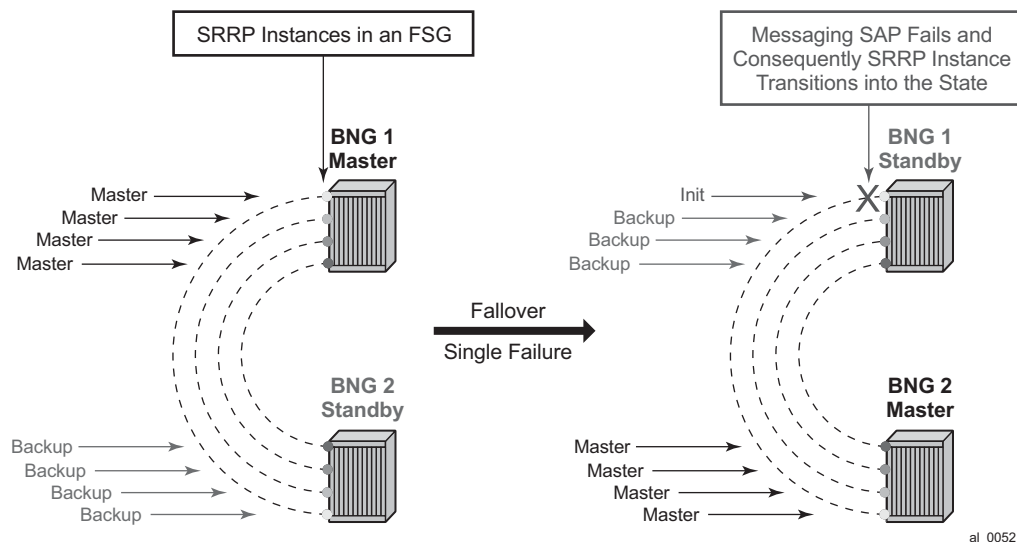


Figure 101: FSG — Single Network Failure

In the case of multiple network failures, the concept of the FSG breaks as there is a possibility that a 'FSG' contains SRRP instances that are in any of the three possible SRRP states: Master, Backup, or Init. This Fate Sharing feature may not provide optimal protection when there are multiple network failures distributed over both redundant nodes.




Figure 102: Multiple Network Failures

The whereabouts of the failure in the network path that SRRP is designed to monitor are not always clearly reflected through SRRP states. For example, if the network failure is somewhere in the aggregation network beyond the direct reach of our BNG, SRRP assumes Mastership on both BNG nodes. This is a faulty condition and the reason why solely monitoring of the SRRP states is not enough to protect against failures. On the other hand, the SRRP messaging SAP states are more indicative of the network failure since they can be tied into Eth-OAM.

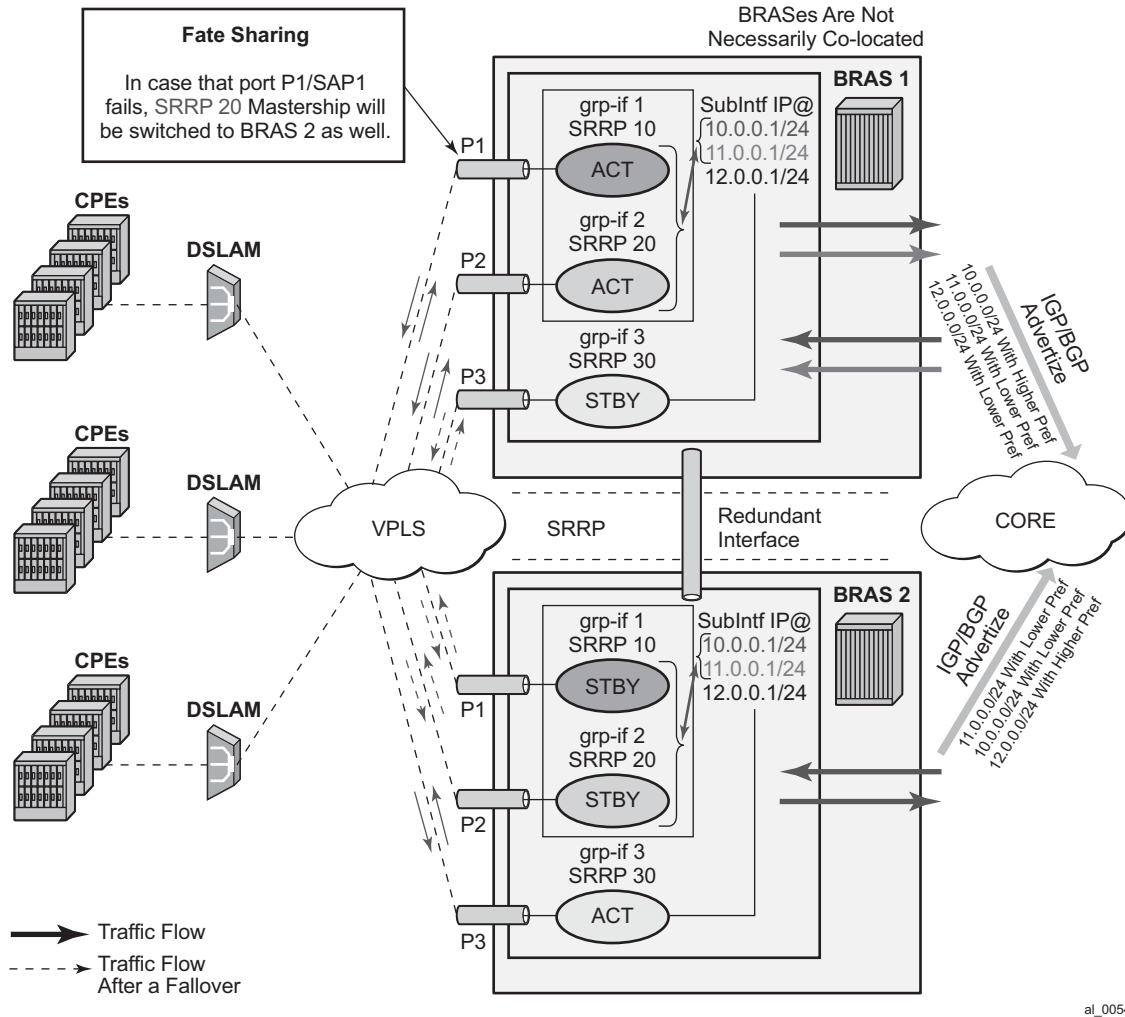
Once a single network failure is detected and as a result an SRRP instance transitions into a non-Master state, the remaining SRRP instances in the FSG are forced into a Backup state. This is achieved by changing the priority of each individual SRRP instance in the FSG.

In the case of simultaneous multiple failures (multiple ports fail at the same time), it is possible that the SRRP instances within the FSG settle in any of the three possible SRRP states: Master, Backup, or Init. In such scenario shunted traffic will ensue.

In the premise of SRRP Fate Sharing, the network failure will be reflected in the operational state of the messaging SAP over which SRRP runs. This will certainly be the case if the failure is localized to the BNG (somewhere on the directly connected link). In the case of non-localized failure (beyond the direct reach of the BNG node), Eth-OAM might be needed in to detect the remote end failure and consequently bring the SAP operationally into a DOWN state.

Once the single network failure is detected, all instance within the FSG transitions into a non-Master state.


If there are no failures in the network, all SAPs are UP and SRRP instances within the FSG are in a homogeneous and deterministic state based on their configured priorities.



**Figure 103: SRRP Fate Sharing**


## Failure Detection in a Fate Sharing Group

1. Dual homing over directly connected ports.  
No Eth-OAM is needed, AN is directly connected to the BNG.



**Figure 104: Scenario 1**

2. Dual homing with aggregation network - aggregation network has no redundancy between Layer 2 switches (STP). To determine whereabouts of failure at point 1 in the figure below, Eth-OAM is needed.



**Figure 105: Scenario 2**

3. Dual homing with aggregation network - aggregation network with redundancy between Layer 2 switches (STP).  
No Eth-OAM is needed in this case for successful operation. However, the failure detection is based on the failure of the directly attached ports.




Figure 106: Scenario 3

4. Single homing with aggregation network.  
In this case, SRRP can protect only against direct failures. Any remote failure leaves a part of the network isolated from the subscriber point of view.




Figure 107: Scenario 4



## Fate Sharing Algorithm

Fate Sharing Algorithm (FSG) is relying on tracking the state of messaging SAPs over which SRRP instances run. An SRRP instance with the messaging SAP operationally DOWN will transition into the Init state.

The transitioning of any messaging SAP in a FSG into an UP/DOWN state will trigger SRRP priority adjustment within the FSG. The SRRP priorities should be chosen carefully to achieve the desired behavior. They are modified dynamically as the SAP states change. The range in which SRRP priorities can be modified is from 1 to the SRRP priority that is initially configured under the SRRP node. Here are some general guidelines for choosing SRRP priorities in a FSG:

- Initially configured SRRP priorities for all SRRP instance within the FSG within the node should be the same.
- Initially configured SRRP priorities should be different between pairing FSGs. For example, SRRP instances in the BNG node A within an FSG will all have the same SRRP priority 'X', while corresponding SRRP instances on the paired node within corresponding FSG will all have SRRP priority 'Y'. This ensures that SRRP mastership is clearly defined between the two BNG nodes. Note that this step is not mandatory as SRRP will naturally break the Master-ship tie in the case that all SRRP priorities are the same. However, following this step may provide a clearer view from an operational perspective.
- The priority-step used for dynamic SRRP priority adjustment must be greater than the difference in initially configured SRRP priorities between two BNG nodes. This ensures that a single failure event triggers the SRRP switchover. Otherwise, if the dynamically lowered SRRP priority is still greater than the one from the SRRP peer, the switchover would not be triggered. Therefore, the fate sharing concept would not function as intended.
- Initially configured SRRP priority of each SRRP instance should be greater than the (anticipated) number of SRRP instances in a FSG multiplied by the SRRP priority-step. This ensures that the dynamically priority never tries to go below 1. There is a code check that prevents SRRP priority going below 1. Nonetheless, it is recommended not to get into a situation where this needs to be enforced in the code.

**Note:** The priorities will never be less than 1 or greater than initially configured SRRP priority.

Example scenarios:

Assume 3 SRRP instances in a FSG. The SRRP instances in the FSG-1 on BNG 1 have the priority of 100, while the SRRP instances in the FSG-2 on BNG 2 have the priority of 95. The priority-step is 6. The SRRP instances and underlying messaging SAPs will be referred as SRRP 1, 2, 3 and SAP 1,2,3, respectively.

Initialization:

Scenario 1 – all SAPs are operationally UP.

BNG 1 boots up and all messaging SAPs transition into the UP state. When the first SRRP instance in FSG-1 comes up, it looks under the FSG to find out how many messaging SAPs are operationally UP. Since all messaging SAPs are operationally UP, this first SRRP instance

assumes its initially configured priority of 100. The other two SRRP instances in the same FSG follows the same sequence of events.

BNG 2 follows the same flow of events. As a result, BNG 1 assumes mastership over BNG 1 for all SRRP instances within the corresponding FSG.

Scenario 2 – messaging SAP 1 is operationally DOWN on BNG 1, the rest of the messaging SAPs are operationally UP.

SRRP 2 and 3, during the initialization, pick up SRRP priority of 94 ( $100 - 1 * \text{priority-step}$ ).

On BNG 2, all messaging SAPs are UP and consequently all SRRP instances within the FSG on BNG 2 have SRRP priority of 95. The SRRP instances on BNG 2 assumes Mastership.

Scenario 3 – Continuing from scenario 2, the SAP 1 on BNG 1 transitions into the UP state. SRRP priority of each SRRP instance in FSG-1 is increased by 6, bringing it to 100, enough to assume Mastership.

---

### Adding a New Instance into an FSG

To introduce minimal network disruption, first create messaging SAPs in both BNG nodes and ensure that both SAPs are operationally UP. Then a new SRRP 4 instance should be created on both BNG nodes. The next step would be to include this new messaging SAP into a SAP monitoring group. And finally, the SRRP-4 is added into the FSG (1 and 2). This triggers the recalculation of SRRP priorities for the existing FSG-1 and FSG-2. Since all SRRP priorities are at the max (initially configured priority), nothing changes.

There are more disruptive ways of adding an SRRP instance into a FSG. One such example would be in the case where SRRP priorities are not at their maximum (initially configured) priority. If an SRRP instance is first added into an FSG that is in a 'Backup' state, this would increase the FSG priority and potentially cause a switchover. If the SRRP instances is then added in a FSG on the peer BNG (previously Master), the priority of this FSG would be increased again and the switchover would unnecessarily occur for the second time. The new SRRP instances, once operational, should always be added in the Master FSG first.

SRRP priority re-calculation within the FSG is triggered by the following events:

- SRRP initialization
- addition of a SAP under the monitoring group
- messaging SAP failure

This priority calculation looks into how many SAPs are in the DOWN state within the monitored SAP group. Based on this number, the priority is calculated as follows:

$\text{SRRP priority} = \text{configured-priority} - \text{priority-step} * \text{num\_down\_SAPs}$ .

## SRRP Aware Routing - IPv4/IPv6 Route Advertisement Based on SRRP State

There are three cases that need to be covered, each case with its own specifics:

- Subscriber Interface Routes (IPv4/IPv6)
- Managed Routes
- Subscriber Management Routes (/32 IPv4 hosts routes and IPv6 PD wan-host routes)

Depending on the route type, the action is to either modify the route metric based on the SRRP state that the route is tracking, OR to advertise/withdraw the route based on the SRRP state that the route is tracking. The action is defined in the routing policy and it is based on the new attributes with which the routes are associated.

To achieve a better granularity of the routes that are advertized, an origin attribute is added to the subscriber management routes (/32 IPv4 routes and IPv6 PD wan-host) with three possible values:

### aaa

#### IPv4

*subscriber-management* /32 host routes that are originated through RADIUS framed-ip-address VSA other than 255.255.255.254. The 255.255.255.254 returned by the RADIUS indicates that the BNG (NAS) should assign an IP address from its own pool.

#### IPv6

*subscriber-management* routes that are originated through framed-ipv6-prefix (SLAAC), delegated-ipv6-prefix (IA\_PD) or alc-ipv6-address (IA\_NA) RADIUS attributes . This is valid for IPoE and PPPoE type host.

### dhcp

#### IPv4

*subscriber-management* /32 host routes that are originated via DHCP server (local or remote) and also RADIUS framed-ip-address=255.255.255.254 (RFC 2865).

#### IPv6

*subscriber-management* routes that are assigned via local DHCPv6 server pools whose name is obtained through Alc-Delegated-IPv6-Pool (PD pool) and Framed-IPv6-Pool (NA pool) RADIUS attributes. This is valid for IPoE and PPPoE type hosts.

In addition, for IPoEv6 only, the pool name can be also obtained via ipv6-delegated-prefix-pool (PD pool) and ipv6-wan-address-pool (NA pool) from LUDb.

### ludb

#### IPv4

*subscriber-management* /32 host routes that are originated via LUDb. This also covers RADIUS fallback category (RADIUS falls back to system-defaults or to LUDb).

#### IPv6

subscriber-management routes obtained from LUDB via ipv6-address (IA\_NA) or ipv6-prefix (IA\_PD). This is supported only for IPoE.

Overall, the following new route attribute is added:

*state: srrp-master, srrp-non-master*

The existing origin attribute is expanded to contain the following values:

*origin: aaa, dhcp, ludb*

These two attribute types are applied in the following fashion:

The state attribute is applied to all three route types: *subscriber interface routes*, *managed routes* and *subscriber management routes*. Each route listens to the SRRP state.

If an attribute is defined in the routing policy as a match condition (from statement) but the route itself does not have this attribute, the route is evaluated into a non-match condition.

The origin attribute is always applied only to subscriber management routes. No additional statement is needed to explicitly apply this attribute as it may be the case for the state attribute.

Every time there is a change in the attribute associated with the route, the route is re-evaluated in the RTM by the routing policy and corresponding action is taken.

---

## Subscriber Interface Routes (IPv4 and IPv6)

Optimized routing and elimination of downstream shunt traffic during normal operation can be achieved by **statically** favoring the routes on the network side that are advertized with an increased metric by Master SRRP nodes.

The downside of this static approach is that during the port/card failure and consequently a SRRP switchover, the node with the failed port/card will continue to advertise routes with the same high metric as long as the subscriber interface is in the 'UP' state (or a single SAP under it). That is, the network side will not be aware of the switchover. It will continue to forward traffic to the standby node, and as a result, heavy shunt traffic will ensue. To effectively deal with this, the network side must be aware of the routing change that occurred in the access layer.

When failure is detected, the metric for the route is changed automatically based on the following configuration:

```
configure
  service <type> <id>
    subscriber-interface <intf-name>
address <ip-address> gw-ip-address <gw-address> track-srrp <srrp-inst> holdup-time <msec>
ipv6
subscriber-prefixes
  prefix <ipv6-prefix> pd track-srrp <srrp-id> holdup-time <msec>
  prefix <ipv6-prefix> wan-host track-srrp <srrp-id> holdup-time <msec>

policy-options
  begin
  policy-statement <name>
```

```

        entry 1
        from
protocol direct
        state 'srrp-master'
        exit
        action accept
        metric set 100
        exit
    exit
    entry 2
    from
protocol direct
        state 'srrp-non-master'
        exit
        action accept
        metric subtract 10
        exit
    exit
    entry 3
    from
        protocol direct
    exit
    action accept
    exit
exit

```

This configuration ensures that the route metric is changed for the subscriber interface routes based on the SRRP state while the other, non-subscriber directly attached routes are unaffected by SRRP.

*Route Advertisement based on SRRP State* requirement is applicable to BGP (IPv4, IPv4-IPVPN) and IGP.

*Routing policy* also provides the flexibility to prevent route advertisement (*action reject*) instead of changing the route metric.

Although this feature is designed to minimize or eliminate the use of the redundant-interface, it is important to note that the redundant-interfaces would still be used in the case of transient conditions. An example of such condition would be:

1. Messaging SAP Fails
2. SRRP switches over
3. Stale routing in the core is still in the effect while the metric is being propagated (or the route is being advertised/withdrawn). During this time, traffic is flowing over the redundant interface.
4. Network convergence is complete
5. Traffic in the network core is redirected to the new Master SRRP node

## Managed Routes

Only the state attribute is applicable to managed routes, and only to the ones that are synchronized (static and RADIUS obtained – framed-route and framed-ipv6-route). The managed routes obtained via BGP are not synchronized and this feature is not applicable to them.

Based on the SRRP state, the managed route can be either advertized with a modified metric or be withdrawn altogether.

For example:

*Managed routes* that are tracking SRRP state are only advertised from the Master node and denied from Backup node. All other managed routes that are not tracking SRRP state are advertised regardless of the SRRP state.

```

policy-options
  begin
    policy-statement <name>
      entry 1
        from
protocol managed
      state 'srrp-master'
      exit
      action accept
      exit
    exit
    entry 2
      from
protocol managed
      state 'srrp-non-master'
      exit
      action reject
      exit
    exit
    entry 3
      from
        protocol managed
      exit
      action accept
    exit
  exit

```

## Subscriber Management Routes (/32 IPv4 Host Routes, IPv6 PD WAN-Host Routes)

Both attributes (state and origin) are applicable to the subscriber management routes.

For Example:

A Service Provider wants to advertise only subscriber-management routes with the origin DHCP and AAA from the Master node. Routes with the LUDB origin are not advertised. Standby node is not advertising any /32 subscriber management routes.

```
policy-options
  begin
    policy-statement <name>
      entry 1
        from
          origin dhcp
          origin aaa
            state 'srrp-master'
        exit
        action accept
        exit
      exit
    exit
```

Default action is reject.

## Activating SRRP State Tracking

The SRRP state tracking by routes is turned on only when desired.

For subscriber-interface routes (IPv4 and IPv6), this is performed explicitly.

```
subscriber-interface <intf-name>
  address <ip-address> gw-ip-address <gw-address> track-srrp <inst-name> holdup-time <msec>
  ipv6
  subscriber-prefixes
    prefix <ipv6-prefix> pd track-srrp <srrp-id> holdup-time <msec>
    prefix <ipv6-prefix> wan-host track-srrp <srrp-id> holdup-time <msec>
```

For managed and subscriber management routes, this is explicitly enabled under the group interface:

```
group-interface <name>
  srrp-enabled-routing holdup-time <msec>
```

## SRRP in Conjunction with a PW in ESM Environment – Use Case

In certain cases, subscriber traffic is terminated on the BNG via an EPIPE. In this case, the subscriber traffic can be offloaded onto a plain Ethernet port via a VSM module (a ‘loop’) so that it can be terminated in ESM. EPIPEs can be configured in A/S configuration and terminated on two BNG nodes in multihomed environment.

In such multi-homed environment with EPIPEs and ‘loops’, the ESM itself would be detached from the EPIPE, which brings the subscriber traffic to the BNG. Because of that, the ESM would not know if the PW’s state is Active or Standby. As a result, in the downstream direction, traffic could end up being forwarded towards the Standby PW, effectively being black-holed.

To overcome this, SRRP can be used in conjunction with an additional mechanism to help monitor the activity of the PWs. This monitoring mechanism is very similar to Fate-sharing. The difference in this case is that the messaging SAP (instead of SRRP instance) is monitoring the activity of the PW. As a result, the SRRP messaging SAP reflects the state of the PW. For example, the PW in a Standby mode would cause the messaging SAP to be in the DOWN state while the PW Active state would cause the messaging SAP to be in the UP state. That is, the SRRP instance reflects the operational state of the messaging SAP. SRRP is indirectly tied into PW state.

Modifying the priority of SRRP instance based on PW’s state as a mean of mapping the Master SRRP into the Active PW would not help here as SRRP messages are not flowing over standby PWs. This is why SRRP state must be enforced via the messaging SAP.

Fate-sharing for PW termination in conjunction with SRRP is not supported.

Metric adjustment for the subscriber routes is supported. Once the tracked SRRP instance transitions into a non-Master state, the state attribute of the route changes and the appropriate action defined in the routing-policy is taken.

## Group-monitor

The failure detection mechanism to trigger an action within FSG relies on the operational state of the messaging SAP. Such failure detection mechanism is referred as a group monitor.

Group monitor can also be used to detect the state change of the PW. PW state change is reflected in the messaging SAP which in turn triggers the state change of an SRRP instance.

All this is implemented through an oper-group object which is described in the ‘Services Guide’. All entities that needs to be monitored (messaging SAPs and PWs) are associated with this oper-group object. Finally, an SRRP instance (in case of FSG) or a messaging SAP (in case of PW) is instructed to monitor the entities in the oper-group object. State transitions of objects in a oper-group object trigger state transitions of entities that are monitoring them (messaging SAPs and SRRP instances). State transitions of monitored objects in a oper-group will cause the following actions:

- In the case of an FSG, priorities of SRRP instances are recalculated
- In the case of PW termination on BNG, the operational state of the messaging SAP is changed.



This is an overview of the CLI syntax showing the principles of how should this work (for exact description of commands and full syntax, please see the command reference):

```
configure>service
oper-group <name> //oper-group creation

configure>service(IES | VPRN)>sub-if>grp-if>sap
    oper-group <name> //adds the SAP to the oper-group
    monitor-oper-group <name> // links the status of the oper-group to the SAP. In this
fashion a messaging SAP can monitor the state of a PW.

configure>service(IES | VPRN)>sub-if>grp-if>srrp x
    monitor-oper-group <name> priority-step [0-253] //with this, a state transition of the
objects in the oper-group should trigger SRRP priority recalculation. The state of the
oper-group is not important but in the state of the objects within. If an object within the
oper-group goes down, the SRRP priority is lowered by a priority-step. The SRRP priority
will be adjusted on every state transition of member objects.

configure>service>epipe>spoke-sdp
    oper-group <name> // this will add a PW to the oper-group. A messaging SAP monitors
this PW and it assumes the state according to the state of the PW in the oper-group. A
standby or a DOWN PW state causes the messaging SAP to assume a DOWN state. Otherwise the
messaging SAP would be in the UP state. In order for the SAP to assume the DOWN state, both
RX and TX side of the PW must be shut. In other words, a PW in standby mode also must have
the local TX disabled by the virtue of the 'slave' flag (standby-signaling-slave command
under the spoke-sdp hierarchy). Without the TX disabled, the SAP monitoring the PW would
not transition in the down state.
```

Hold timer is provided within the oper-group command to suppress flapping of the monitored object (SAP or pseudowire).

Example with ESM over pseudowire through a VSM 'loop'.

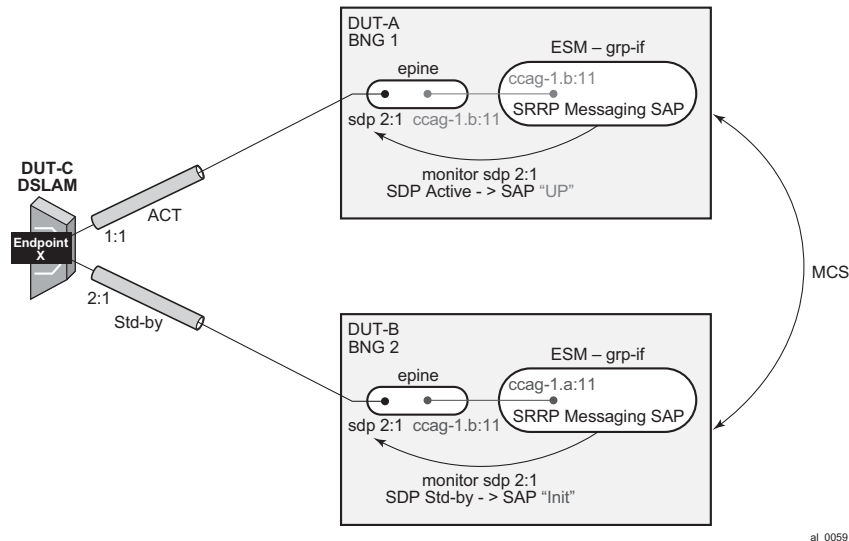


Figure 108: Pseudowire Example

```

*A:Dut-C>config>service>epipe# info
-----
    endpoint "x" create
        standby-signaling-master
    exit
    sap 1/1/7:1 create
    exit
    spoke-sdp 1:1 endpoint "x" create
        precedence primary
        no shutdown
    exit
    spoke-sdp 2:1 endpoint "x" create
        no shutdown
    exit
    no shutdown
-----

*A:Dut-A>config>service>epipe# info
-----
    sap ccag-1.b:11 create
    exit
    spoke-sdp 2:1 create
        standby-signaling-slave
        oper-group "1"
        no shutdown
    exit
    no shutdown
-----

*A:Dut-B>config>service>epipe# info
-----
    sap ccag-1.b:11 create
    exit
    spoke-sdp 2:1 create
        standby-signaling-slave
        oper-group "1"
        no shutdown
    exit
    no shutdown
-----

*A:Dut-A>config>service>ies# info
-----
    redundant-interface "redif11" create
        address 101.1.1.2/24 remote-ip 101.1.1.4
    spoke-sdp 1:1 create
        no shutdown
    exit
    exit
    subscriber-interface "subif_1" create
        shutdown
        address 1.1.1.2/24 gw-ip-address 1.1.1.100
    group-interface "grpif_1_2" create
        shutdown
        redundant-interface "redif11"
    exit
    exit
    subscriber-interface "subTest" create
        address 80.1.1.2/24 gw-ip-address 80.1.1.254
    group-interface "grpTest" create
        redundant-interface "redif11"

```

```

        sap ccag-1.a:1 create
        exit
        sap ccag-1.a:11 create
            monitor-oper-group "1"
        exit
        srrp 11 create
            message-path ccag-1.a:11
            no shutdown
        exit
    exit
exit
no shutdown
-----
*A:Dut-B>config>service>ies# info
-----
        redundant-interface "redif11" create
            address 101.1.1.4/24 remote-ip 101.1.1.2
            spoke-sdp 1:1 create
                no shutdown
            exit
        exit
        subscriber-interface "subif_1" create
            shutdown
            address 1.1.1.4/24 gw-ip-address 1.1.1.100
        exit
        subscriber-interface "subTest" create
            address 80.1.1.4/24 gw-ip-address 80.1.1.254
            group-interface "grpTest" create
                redundant-interface "redif11"
                sap ccag-1.a:1 create
                exit
                sap ccag-1.a:11 create
                    monitor-oper-group "1"
                exit
                srrp 11 create
                    message-path ccag-1.a:11
                    no shutdown
                exit
            exit
        exit
        exit
        no shutdown
-----
*A:Dut-B>config>service>ies# show srrp
=====
SRRP Table
=====
ID          Service      Group Interface      Admin      Oper
-----
11          1              grpTest              Up          initialize
-----
No. of SRRP Entries: 1
=====
*A:Dut-A>config>service>ies# show srrp
*A:Dut-A>config>service>ies#
=====
SRRP Table
=====
ID          Service      Group Interface      Admin      Oper
-----

```

SRRP Enhancement

11	1	grpTest	Up	master
-----				
No. of SRRP Entries: 1				
=====				

## Subscriber Override

This feature provides the ability to override queue and policer parameters (CIR, PIR, CBS, MBS) as well as HQoS parameters (egress aggregate-rate and root-arbiter rate) configured at sla-profile and sub-profile level in order to provide per-subscriber-(host) customizations. The goal is to avoid an explosion of the sub-profiles and sla-profiles to cover all service level combinations. This customization of QoS related parameters can in principle occur during authentication (auth-response message) or during sub-host life time by RADIUS CoA messages.

The QoS parameter customizations are communicated by RADIUS server in form of RADIUS VSAs which can be included in RADIUS-authentication response message or in CoA message.

The Alc-Subscriber-QoS-Override VSA (126) is a string with following layout “direction:type:[key:]values” where:

- Object-type:
  - **direction** represents single character indicating **i** for ingress and **e** for egress.
  - **type** represents single character indicating **q** for queue, **p** for policer, **r** for aggregate-rate and **a** for arbiter.
  - **key** is indicated the queue or policer-id. It is not used in case of aggregate-rate and root-arbiter.
  - **values** indicates actual values preceded with keywords used in CLI (e.g., cir).

Aspects, such as parent, priority level, stats- mode are not accessible through this customization. Instead, a new policy should be created on the node.

The key identifying the subscriber-host in the RADIUS CoA message is accounting-session-id. This is different in previous releases, where the *service-id* and *ip-address* are mandatory fields in RADIUS CoA message.

The operational value of the QoS objects (queues or schedulers) are derived from different inputs. As in queue/policer parameters, the following hierarchy of inputs are respected (highest priority is the first):

- On-line charging overrides
- RADIUS response/CoA overrides
- Queue overrides configured at sla-profile level
- Queue parameters set in QoS policy level

In the case of scheduler/arbiter-overrides, the following hierarchy of inputs apply:

- ANCP overrides
- RADIUS response overrides
- Scheduler/arbiter parameters as configured in scheduling/policer-control-policy.

The above rules are generic. If any given override mechanism is not yet applicable to policers (or arbiters) it will be skipped. The above rules indicate the priority if all mechanisms are supported.

## Subscriber Override

The QoS overrides received in RADIUS message (in the form of a VSA) are per definition related to a given subscriber-host the given message is referring to. Internally, the overrides are applied on per SLA instance level (queue/policer-overrides) and per-subscriber-level (scheduler/arbitrator overrides). The subscriber-overrides are applicable to PPPoE hosts only.

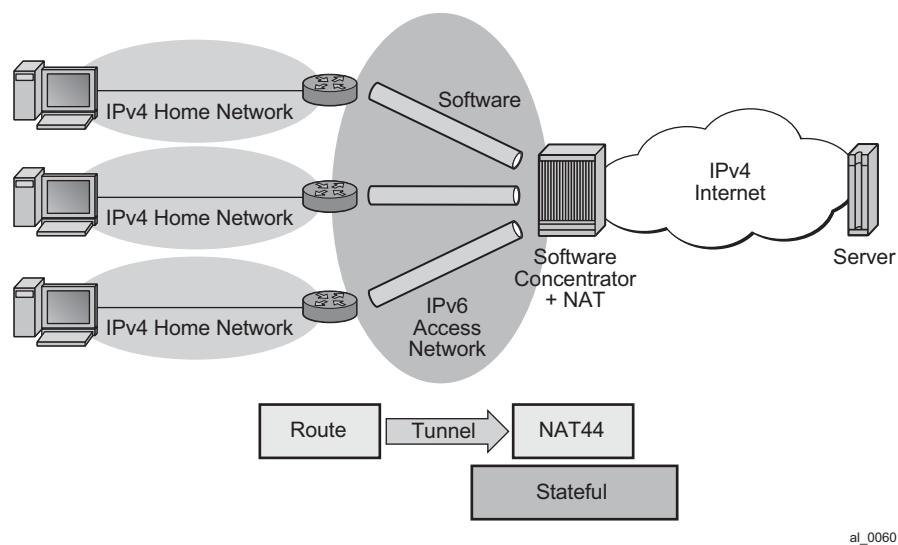
In a dual-homing environment, the adapted values are synchronized through MCS, but they do not need to be persistent.

## Dual Stack Lite

Dual Stack Lite feature is supported on the 7710 SR-Series in combination with the MS-ISA to function as a DS-Lite Address Family Transition Router (AFTR).

Dual Stack Lite is an IPv6 transition technique that allows tunnelling of IPv4 traffic across an IPv6-only network. Dual-stack IPv6 transition strategies allow service providers to offer IPv4 and IPv6 services and save on OPEX by allowing the use of a single IPv6 access network instead of running concurrent IPv6 and IPv4 access networks. Dual-Stack Lite has two components: the client in the customer network, known as the Basic Bridging BroadBand element (B4) and an Address Family Transition Router (AFTR) deployed in the service provider network.

Dual-Stack Lite leverages a network address and port translation (NAPT) function in the service-provider AFTR element to translate traffic tunneled from the private addresses in the home network into public addresses maintained by the service provider. On the 7750 SR, this is facilitated through the Carrier Grade NAT function.



al\_0060

**Figure 109: Dual-Stack Lite**

As shown in [Figure 109](#), Dual-Stack Lite has two components, a software initiator in the RG and a software concentrator, deployed in the service provider network, where control-less IP-in-IP (using protocol 4 - IPv4 in IPv6) is used for tunnelling. When using control-less protocol, packets are sent on the wire for the remote software endpoint without prior setup of a tunnel.

The software initiator in the home network is combined with a routing function, where the default route is passed to the software pseudo-interface. Note that there is no NAT function, therefore, the private IP addresses of the home network are encapsulated without source address modification,

and forwarded to the software concentrator where all NAT is performed. The software pseudo-interface unicasts all IPv4 traffic to the IPv6 address of the software concentrator, which was pre-configured.

When encapsulated traffic reaches the software concentrator, the device treats the source-IP of the tunnel to represent a unique subscriber. The software concentrator performs IPv4 network address and port translation on the embedded packet by re-using Large Scale NAT and L2-Aware NAT concepts.

IP-in-IP

As shown in Figure 110, IP-in-IP uses IP protocol 4 (IPv4) to encapsulate IPv4 traffic from the home network across an IPv6 access network. The IPv4 traffic tunnelling is treated as best-effort with no subscriber management or policy, and does not use ESM. The scale is dependant only on the internal structures of the MS-ISA and CPM, that is, the IP-in-IP model can support more subscribers than an ESM-based approach.

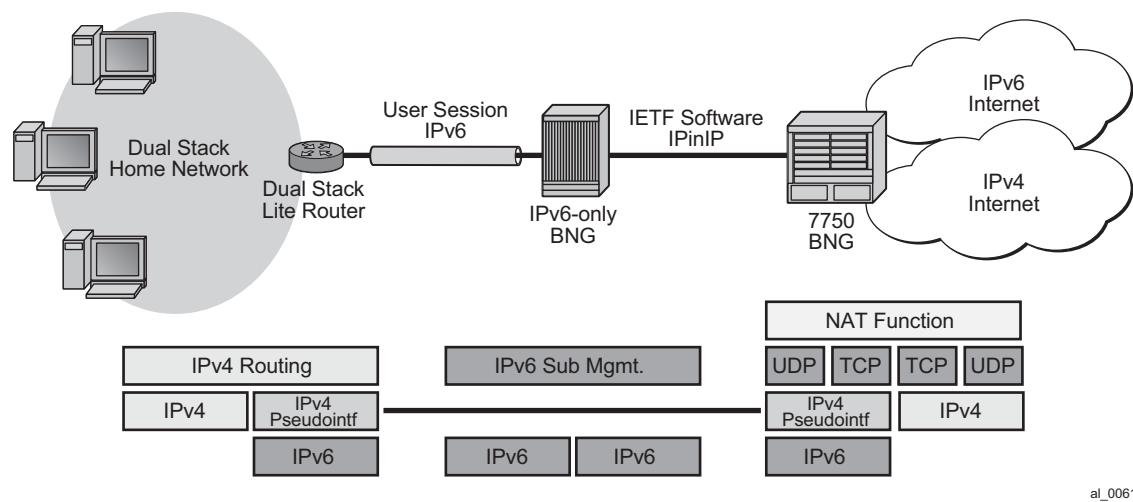


Figure 110: IP-in-IP

Dual-Stack Lite IP-in-IP is configured through the existing **nat** command that is inside the CLI statements that are within the base router or VPRN. A service performing large scale NAT supports Dual-Stack Lite.

Dual-Stack Lite expects a routing (non-NATing) gateway in the home, where many different IPv4 inside addresses exist for each subscriber. These inside addresses may overlap other subscriber's address, especially given the heavy use of RFC 1918 address space.



The lack of control of protocol for the IP-in-IP tunnels simplifies the functional model, since any received IPv4 packet to the ISA dual-stack-lite address can simply be:

- Checked for protocol 4 in the IPv6 header.
- Checked that the embedded IP packet is IPv4.
- Processed as if it were L2-Aware, where the source-IP of the tunnel (the source IPv6 address) is used as the subscriber identifier.

Note that the inside IP address in the NAT, tables must not be the IPv6 address of the tunnel, but the true IPv4 address of any host within the home. The subscriber-id must be the literal IPv6 address (appreciating this may be 34 characters in length).

---

## Configuring Dual Stack Lite

Dual Stack Lite is configured on an inside service and uses the existing Large Scale NAT nat-policies and outside pools. Dual-Stack Lite and NAT44 Large Scale NAT can operate concurrently on the same inside and outside services.

Dual Stack Lite is configured using the following CLI:


```
configure {router | service vprn service-id}
- [no] nat
  - inside
    - [no] dual-stack-lite
      - [no] *address ipv6-address
```

## L2TP over IPv6

In this mode, L2TP provides the transport for IPv4 that allows full ESM capabilities on the 7750 SR. From the 7750 perspective, the L2TP tunnel is no different in capability to those already supported. Only the underlying transport (IPv6 instead of IPv4) distinguishes this approach.

To support legacy IPv4 access, L2TP over IPv6 is combined with the existing L2-Aware NAT feature as shown in [Figure 111](#).

As ESM is used, scale is limited by the number of ESM hosts supported on a chassis and any associated resources like queues.



**Figure 111: L2TP over IPv6**

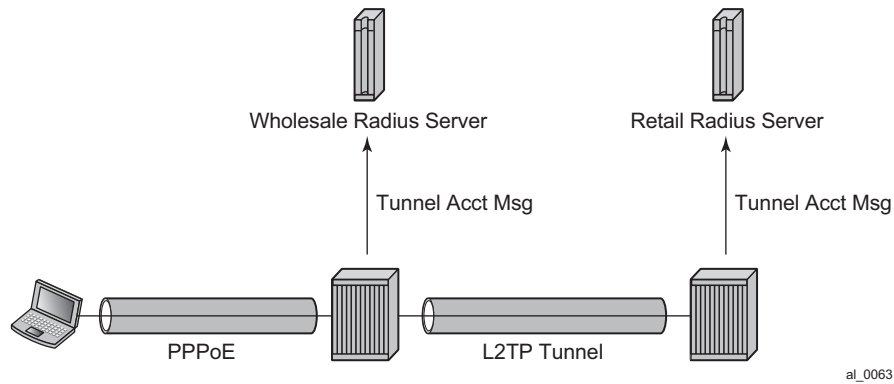
L2TP LNS over IPv6 is supported in both the base routing instance and VPRN that has 6VPE configured.

Like the SR-OS 8.0 LNS implementation, tunnels are terminated on any routing interface, including loopback, SAP, or network port. A single interface simultaneously supports IPv4 and IPv6 L2TP tunnel termination by having two different addresses configured.

For greater scalability, L2TP tunnel and session count per chassis are increased to allow 1 tunnel per session.

NAT capabilities are supported via existing L2-Aware NAT methods. Note that the L2TP LNS over IPv6 may be used without NAT as well and the L2TP sessions may be either IPv6-only or dual-stack.

## L2TP Tunnel RADIUS Accounting



**Figure 112: L2TP Tunnel Accounting**

When L2TP tunnel accounting is enabled, except for **host** or **sla-profile**-based accounting packets and attributes, the following are additional accounting packets and attributes:

- Accounting packets: tunnel-start/stop/reject; tunnel-link-start/stop/reject — There are no interim updates for L2TP tunnel/session accounting.
- RADIUS accounting attributes:
  - Tunnel-Assignment-Id (LAC only)
  - Acct-Tunnel-Connection
  - Acct-Tunnel-Packets-Lost

These attributes were added into current account-start/stop/interim-update packets (host accounting/sla-profile accounting)

Tunnel level accounting and session level accounting can be enabled or disabled independently.

New accounting packets and related RADIUS attribute list are described in [Table 15](#).

Some considerations of RADIUS attributes are described in [RADIUS Attributes Value Considerations on page 1012](#)

## Accounting Packets List

Table 15 describes L2TP tunnel accounting behavior along with some key RADIUS attributes (apply for both LAC and LNS):

**Table 15: L2TP Tunnel Accounting Behavior**

Act-Packet	When	Key Attributes	Remark
Tunnel-Start	A new L2TP tunnel is created	Acct-Session-ID	
		Event-Timestamp	
		Tunnel-Type:0	
		Tunnel-Medium-Type:0	
		Tunnel-Assignment-Id:0	
		Tunnel-Client-Endpoint:0	
		Tunnel-Client-Auth-Id:0	
		Tunnel-Server-Endpoint:0	
		Tunnel-Server-Auth-Id:0	
Tunnel-Reject	A new L2TP tunnel creation failed	Acct-Session-Id	
		Event-Timestamp	
		Tunnel-Type:0	
		Tunnel-Medium-Type:0	
		Tunnel-Assignment-Id:0	
		Tunnel-Client-Endpoint:0	
		Tunnel-Client-Auth-Id:0	
		Tunnel-Server-Endpoint:0	
		Acct-Terminate-Cause	
Tunnel-Stop	An established L2TP tunnel is removed	Acct-Session-Id	
		Event-Timestamp	
		Tunnel-Type:0	
		Tunnel-Medium-Type:0	

**Table 15: L2TP Tunnel Accounting Behavior (Continued)**

Act-Packet	When	Key Attributes	Remark
		Tunnel-Assignment-Id:0	
		Tunnel-Client-Endpoint:0	
		Tunnel-Client-Auth-Id:0	
		Tunnel-Server-Endpoint:0	
		Tunnel-Server-Auth-Id:0	
		Acct-Session-Time	
		Acct-Input-Gigawords	
		Acct-Input-Octets	
		Acct-Output-Gigawords	
		Acct-Output-Octets	
		Acct-Input-Packets	
		Acct-Output-Packets	
		Acct-Terminate-Cause	
Tunnel-Link-Start	An L2TP session is created	User-Name	
		Acct-Session-Id	This is the same as Acct-Session-id in access-request of host auth
		Event-Timestamp	
		Service-Type	Framed
		Class	
		Tunnel-Type:0	
		Tunnel-Medium-Type:0	
		Tunnel-Assignment-Id:0	
		Tunnel-Client-Endpoint:0	
		Tunnel-Client-Auth-Id:0	
		Tunnel-Server-Endpoint:0	
		Tunnel-Server-Auth-Id:0	

**Table 15: L2TP Tunnel Accounting Behavior (Continued)**

Act-Packet	When	Key Attributes	Remark
		Acct-Tunnel-Connection	See <a href="#">RADIUS Attributes Value Considerations on page 1012</a>
Tunnel-Link-Reject	A new L2TP session creation is failed	Acct-Session-Id	Should be as same as Acct-Session-id in access-request of host auth
		Event-Timestamp	
		Tunnel-Type:0	
		Tunnel-Medium-Type:0	
		Tunnel-Assignment-Id:0	
		Tunnel-Client-Endpoint:0	
		Tunnel-Client-Auth-Id:0	
		Tunnel-Server-Endpoint:0	
		Acct-Terminate-Cause	
		Acct-Tunnel-Connection	
Tunnel-Link-Stop	A established L2TP session is removed	User-Name	
		Acct-Session-Id	Should be as same as Acct-Session-id in access-request of host auth
		Event-Timestamp	
		Service-Type	Framed
		Class	
		Tunnel-Type:0	
		Tunnel-Medium-Type:0	
		Tunnel-Assignment-Id:0	
		Tunnel-Client-Endpoint:0	
		Tunnel-Client-Auth-Id:0	

**Table 15: L2TP Tunnel Accounting Behavior (Continued)**

<b>Act-Packet</b>	<b>When</b>	<b>Key Attributes</b>	<b>Remark</b>
		Tunnel-Server-Endpoint:0	
		Tunnel-Server-Auth-Id:0	
		Acct-Tunnel-Connection	
		Acct-Session-Time	
		Acct-Input-Gigawords	
		Acct-Input-Octets	
		Acct-Output-Gigawords	
		Acct-Output-Octets	
		Acct-Input-Packets	
		Acct-Output-Packets	
		Acct-Tunnel-Packets-Lost	
		Acct-Terminate-Cause	

Notes:

- Errors will occur if there are multiple hosts sharing the same sla-profile instance and then these hosts go to different tunnel.
- 7750 SRs have an internal limitation of 500 pps for accounting messages. This feature shares the same limitation

# RADIUS Attributes Value Considerations

- The value of Acct-Tunnel-Connection uniquely identify a L2TP session, and in order to match LAC and LNS accounting record, the value of Acct-Tunnel-Connection is determined by a method shared by LAC and LNS. This means for a given L2TP session, Acct-Tunnel-Connection from the LAC and LNS are the same.
- Current ESM stats are used in Tunnel-Link and tunnel level accounting. This applies for both standard attribute and the 7750's own VSA.
- Tunnel level accounting stats need to aggregate all sessions stats that belong to the tunnel. Note: there could be sessions come and go before tunnel is down, so system need to remember the stats of every session that has been created within the tunnel.  
  
This applies for both standard attribute and 7750's own VSA.
- The value of Acct-Tunnel-Packets-Lost is the aggregation of all discarded packets on both ingress and egress.

---

## Other Optional RADIUS Attributes

Table 16 lists the optional attributes that could be optionally included in tunnel accounting packet, some of them are applied for link level accounting only.

**Table 16: Optional RADIUS Attributes**

Attribute	Tunnel/Link
nas-identifier	Both
nas-port	Link level only
nas-port-id	Link level only
nas-port-type	Link level only



## RADIUS VSA to Enable L2TP Tunnel Accounting

In order to support pure RADIUS-enabled L2TP tunnel accounting on LAC side, the following RADIUS VSA are supported:

**Table 17: Supported RADIUS VSAs**

VSA	Type	Value
ALC-Tunnel-Accounting-Policy	String	Policy-name; if the name is <b>disable</b> then this means L2TP tunnel accounting is disabled for this tunnel

Note: ALC-Tunnel-Accounting-Policy takes precedence over what has been defined in CLI when Alc-Tunnel-Group is also returned.

## MLPPP on the LNS Side

With MLPPP, the counter on LNS side is only available for the bundle, not for each link, so the SR OS's behavior is:

- For each new link session system sends a tunnel-link-start.
- For each link session that is deleted system sends a tunnel-link-stop.
- For all link sessions except the last one system reports 0 for all counters.
- For the last link session, system reports the actual counters for the bundle.

## RADIUS Route Download

The RADIUS route download mechanism periodically polls a RADIUS server for routes to download. The main objective of this feature is to download, in advance, customer-assigned subnets so that they can be re-advertised to the corresponding routing protocols. In this way, subscriber bringup can potentially be done faster (as the routes are already in place and advertised) and, most importantly, reduce the routing protocol churn as subscribers connect and disconnect. The routes being learned through this mechanism could be both managed routes/delegated prefixes as well as the WAN IP assigned to the subscriber in the case PPPoE and un-numbered interfaces are being used.

The route download process requests the routes to a configured RADIUS server by triggering an access-request message. The key identifier for this message is the username, which is a combination of the system's name (or an optionally configured value), appended by a dash ( "-") and then a monotonically increasing integer. The download process sends an access request starting with 1 (such as "hostname-1") and the RADIUS server replies with an access-accept message and a number of routes embedded within the message. The system then increases the counter and sends another access request (this time being hostname-2) and receive a reply with the next batch of routes to download. The process continues, incrementing the counter by 1 each time until the system gets an access-reject or the maximum number of routes that can be downloaded is reached.

The routes to be accepted are in the following format:

```
[vrf {vprn-name | vprn-service-id}] prefix-mask {null0 | null 0 | black-hole} [metric] [tag
tag-value]
```

The prefix-mask could be in any form as 'prefix/length', 'prefix mask' or 'prefix' (in the latter case, for IPv4 routes, the mask shall be derived from the IP class of the prefix).

The route formats are supported:

- Framed-Route (RADIUS attribute 22)

```
Framed-Route = "192.168.3.0 255.255.255.0 null0"
```

```
Framed-Route = "vrf vrfboston 192.168.1.0/24 null 0 0 tag 6"
```

```
Framed-Route = "vrf 2001 192.168.10.0/24 black-hole 0 tag 8"
```

- Cisco-AVPair (Cisco VSA 26-1)

```
cisco-avpair = "ip:route = 192.168.3.0 255.255.255.0 null0"
```

```
cisco-avpair = "ip:route = vrf vrfboston 192.168.1.0/24 null 0 0 tag
6"
```

IPv6 routes are also supported. The format is based on using the IETF-defined IPv6 Framed-IPv6-Route (attribute 99). The following text shows the supported formats.

- Framed-IPv6-Route (RADIUS attribute 99)

```
Framed-Route = "2001:100:bad:cafe::/64 null0"
```

```
Framed-Route = "vrf vrfboston 2100:5aaa:dead:beaf::/96 null 0 0 tag 6"
```

```
Framed-Route = "vrf 3000 2200:1bbbb:dead::/48 black-hole 0 tag 6"
```

All the routes downloaded will be a new protocol type “**periodic**”. The downloader process restarts the AAA requests after a given interval (a configurable value but target refresh rate is 15 minutes) and routes shall be updated according to the following process:

- When the router initiates a new download process, the routes are kept in a temporary table until the download process completes (receives an access-reject from the AAA). The temporary download table is then checked for errors and finally, any changes reflected to the actual routing table.
- Routes no longer present in the download will be removed from the routing table.
- If the AAA server responds with an access-reject for the first username (that is, an implicit empty route-download table), all routes will be removed from the routing table.
- If there are any protocol errors (at the RADIUS level), such as time-out, no response, bad record format, too many records, etc., the download process is suspended and retried after a configurable timer. The minimum retry timer is at least 1 minute and given the light load this represents control-plane-wise (concurrent downloads are not supported) the retries can continue infinitely until the next refresh period occurs, where the download restarts from the beginning. An exponential backoff algorithm with a configured minimum and maximum delay will be used to determine the retry timer.
- In any case, the routes are only purged from the routing table after a complete download process was achieved (properly terminated with an access-reject message). Under any other failure condition, the routes shall remain active. Shutting down the download process should not remove the downloaded routes. A clear command will be provided to clear the periodic routes.
- All the imported routes (blackholes) will be imported into the line-card FIBs to avoid the routing loops caused by announcing the prefixes but not installing the actual blackholes.

## Managed SAP (MSAP)

Managed SAPs allow the use of policies and a SAP template for the creation of a SAP. Although the router supports automatic creation of subscriber hosts in a shared SAP, the most secure mode of operation and common mode is the subscriber per SAP model. In this model, each subscriber is defined with its own VLAN. This feature uses authentication mechanisms supported by the node to provide a SAP.

The reception of a trigger packet initiates a RADIUS or local user database authentication to provide the service context in which the MSAP should be created. The VLAN of the created MSAP is the same as the trigger packet. An MSAP is similar to a regular SAP but its configuration is not user editable and not maintained in the configuration file. The managed SAP remains active as long as the session is active.

The following trigger types are supported:

- DHCP discover (or requests if configured) for DHCP clients. The managed SAP lifetime is defined by the lease time.
- PPPoE PADI for the PPPoE client. The managed SAP lifetime is defined by the session time. The MSAP is installed after the IP address is provided. A short temporary state handles packets between the PADO and ACK.
- ARP defines the managed SAP lifetime. The ARP entry refresh behavior is maintained.
- DHCP6
- PPP
- rtr-solicit (SLAAC hosts)

Multiple trigger types can be enabled on a single capture-sap.

Trigger packets are received on a capture SAP that must be configured in a VPLS service. A capture SAP is defined in a similar way as a regular default SAP but it does not forward traffic. It's sole purpose is to capture trigger packets for authentication. Refer to the following configuration example:

```
vpls 10 customer 1 create
  sap 1/1/1:*. * capture-sap create
    description "capture sap"
    trigger-packet arp dhcp dhcp6 pppoe
    authentication-policy "auth-policy-1"
  exit
  no shutdown
exit
```

A capture SAP and default SAP cannot be configured at the same time for a single port with the dot1q encapsulation or for a single port:topq combination with qinq encap. The capture SAP is used if a more specific match for the Q or Q-in-Q tags is not found by the IOM. If a capturing SAP is defined, triggering packets are sent to the CPM. Non-triggering packets captured by the capturing SAP are dropped.

An ingress VLAN ID (VID) type mac filter can be configured on a capture-sap to have additional control on the vlans that are allowed to initiate a host setup. Other filter types are not supported on a capture-sap.

Supported modes:

- Port:\*: Provides a context for the trigger packet and SAP template.
- A capture SAP can be created in the format Port:Q.
- Port:Q: A specific Q-tag defined SAP for the port and already running managed SAPs.

Supported Q-in-Q modes:

- Port:\*,\*, or Q,\*: Both q-tags will always be sent to RADIUS. The MSAP created will bear both q-tags that arrived in the original packet if authenticated by RADIUS.
- Port:\*.Q: It is an inverse capture-sap that matches on a fixed inner tag with the outer tag identifying the user. The following restrictions apply when an inverse capture-sap is configured on a port:
  - Ethernet ports only
  - It is not possible to create y.\* saps when there is a \*.x capture sap present on the port (y=0,1..4094,\* and x=1..4094).
  - It is not possible to create a y.\* network interface when there is a \*.x capture SAP present on the port (y=0,1..4094,\* and x=1..4094).

A set of mandatory parameters should be provisioned for MSAP creation are as follows:

- **Service id:** service context in which the MSAP will be created.
- **Interface id:** name of the group-interface context in which the MSAP will be created. The group-interface must exist in the provided service for the MSAP to be installed (routed CO scenario only).
- **MSAP policy:** name of the policy that defines the MSAP parameters. The policy must exist in the subscriber-mgmt context.

These parameters can be obtained from the following order of preference:

1. Local user database lookup.
2. RADIUS attributes.
3. Defaults configured at the capture-sap context.

The MSAP parameters can be obtained from a local user database. The local user database should be configured at the capture sap and group-interface context. For example,

# IPoEv4 hosts

```
>config>service>vpls>sap# dhcp-user-db <local-user-db-name>
>config>service>ies>sub-if>grp-if>dhcp# user-db <local-user-db-name>
```

# IPoEv6 hosts

## Managed SAP (MSAP)

```
>config>service>vpls>sap# dhcp6-user-db <local-user-db-name>
>config>service>ies>sub-if>grp-if>ipv6>dhcp6# user-db <local-user-db-name>
```

### # PPP hosts

```
>config>service>vpls>sap# pppoe-user-db <local-user-db-name>
>config>service>ies>sub-if>grp-if>pppoe# user-db <local-user-db-name>
```

When RADIUS authentication is still required after local user database authentication, then the authentication policy must be specified in the local user database. In this case no authentication policy may be configured at the group-interface context. For example,

### # IPoE hosts

```
>config>subscr-mgmt>loc-user-db>dhcp>host# auth-policy <policy-name>
```

### # PPP hosts

```
>config>subscr-mgmt>loc-user-db>ppp>host# auth-policy <policy-name>
```

The MSAP parameters are configured at the local user database host context. For example,

```
>config>subscr-mgmt>loc-user-db>dhcp>host# msap-defaults
>config>subscr-mgmt>loc-user-db>ppp>host# msap-defaults
```

- msap-defaults

```
[no] group-interface - Configure the group interface
[no] policy          - Configure the MSAP policy
[no] service         - Configure the service
```

When RADIUS authentication is required to return the MSAP parameters without prior local user database authentication, then the authentication policy should be configured at the capture-sap context. In a Bridged CO model, the authentication policy specified in the capture-sap will also be used for the MSAP in the VPLS service. In a Routed CO model, the same authentication policy must also be configured at the group-interface context. For example,

```
>config>service>vpls>sap# authentication-policy <auth-policy-name>
>config>service>ies>sub-if>grp-if# authentication-policy <auth-policy-name>
```

The MSAP will not be created if the group-interface name returned from RADIUS has a different authentication policy than the authentication policy configured at the capture-sap.

The following table lists the RADIUS attributes (VSAs) to include in a RADIUS access accept message to obtain MSAP parameters in the RADIUS authentication phase.

Attribute name	Type	Purpose and Format
Alc-MSAP-Serv-Id [26-6527-31]	Integer	Service ID of the service context in which the MSAP will be created.
Alc-MSAP-Policy [26-6527-32]	String	Name of the policy that defines the MSAP parameters.
Alc-MSAP-Interface [26-6527-33]	String	Name of the group-interface context in which the MSAP will be created.

MSAP parameters that are not obtained from a local user database lookup, and that are not returned from RADIUS can be specified in the default-msap section of the capture-sap context (last resort):

```
>config>service>vpls>sap# msap-defaults ?
- msap-defaults

[no] group-interface - Configure the group interface
[no] policy          - Configure the MSAP policy
[no] service         - Configure the service
```

MSAPs can be created in IES or VPRN group interfaces (Routed CO model) and in a VPLS service (Bridged CO model).

The managed SAP configuration can be persistent. The template MSAP policy is stored with the subscriber host which in turn can be made persistent.

If local user database or RADIUS authentication did not provide all the required information to create the subscriber host (no IP address for example), then the MSAP is created with a short timer while waiting for the host to acquire the missing information. If no host is instantiated when the timer expires, the MSAP is deleted.

Multiple subscribers and/or subscriber hosts can share a single MSAP. The MSAP is created with the first instantiated subscriber host and deleted when the last associated subscriber host is removed from the system. Note that only a single MSAP policy is allowed to be specified for a given MSAP. An attempt to change the MSAP policy by a new subscriber host for an existing MSAP will result in a host setup failure.

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.

## ESM Identification Process

---

### SAP-ID ESM Identifier

Providers migrating from Basic Subscriber Management (BSM) can assign a subscriber to a SAP. The SAP ID ESM identifier makes the transition easier by allowing the operator to continue using the *sap-id* as a subscriber-ID.

An ESM SAP ID provides the system the ability to:

- Provide access to the SAP ID string in the Python script.
  - Allow the automatic assignment of the SAP-ID to a static subscriber or subscriber host.
- 

### DSLAM-ID

A DSLAM ID provides the system the ability to define a DSLAM-ID string provided through the Python script, RADIUS, or local user database. If the DSLAM-ID was provided, but the subscriber host is instantiated on a regular MDA (a non-HSMDA), the DSLAM-ID will be ignored.

The HSMDA and the ability to aggregate subscribers into DSLAMs for the purpose of QoS, can use the SAP ID to identify subscribers and associated DSLAMs.

---

### Default-Subscriber

This feature provides a default subscriber definition under the SAP. If the object was configured the operator may use ESM without enabling a processing script or a RADIUS authentication policy. In the event both have been disabled any host that was installed for the SAP will be installed with the configured default subscriber ID. If a RADIUS policy was used or if a script was enabled but a subscriber ID was not returned the default subscriber ID will be used.



## Multicast Management

The multicast-management CLI node contains the bandwidth-policy and multicast-info-policy definitions. The bandwidth-policy is used to manage the ingress multicast paths into the switch fabric. The multicast-info-policy is used to define how each multicast channel is handled by the system. The policy may be used by the ingress multicast bandwidth manager, the ECMP path manager and the egress multicast CAC manager.

---

## Subscriber Mirroring

This section describes mirroring based on a subscriber match. Enhanced subscriber management provides the mechanism to associate subscriber hosts with queuing and filtering resources in a shared SAP environment. Mirroring used in subscriber aggregation networks for lawful intercept and debugging is required. With this feature, the mirroring capability allows the match criteria to include a subscriber-id.

Subscriber mirroring provides the ability to create a mirror source with subscriber information as match criteria. Specific subscriber packets can be mirrored mirror when using ESM with a shared SAP without prior knowledge of their IP or MAC addresses and without concern that they may change. The subscriber mirroring decision is more specific than a SAP. If a SAP (or port) is placed in a mirror and a subscriber host of which a mirror was configured is mirrored on that SAP packets matching the subscriber host will be mirrored to the subscriber mirror destination.

The mirroring configuration can be limited to specific forwarding classes used by the subscriber. When a forwarding class (FC) map is placed on the mirror only packets that match the specified FCs are mirrored. A subscriber can be referenced in maximum 2 different mirror-destinations: 1 for ingress and 1 for egress.

## Volume and Time Based Accounting

Time and volume-based accounting includes the following components:

- Metering function performing stateful monitoring of the service delivery to the subscriber.
  - Communication with an external management system that gets and updates credit per subscriber, notifications of credit exhaustion, etc.
  - Action on credit exhaustion takes pre-defined action when the credit has been exhausted
- 

### Metering

Metering represents the core of time and volume-based accounting. Service usage is typically measured by performing an accounting of the traffic passing through corresponding subscriber-host queues (volume usage) or by keeping lease-state while the given subscriber-host is connected to the network (time usage).

- Statefulness — The accounting information is compared with pre-defined credit expressed in terms of time or volume to monitor service usage.
- Sensitivity — Defining so called activity-threshold allows distinction between subscriber-host being connected and subscriber-host effectively using the service. This is particularly of interest in cases of time based charging.
- Aggregated usage per-category per-subscriber-host — Accounting information can be reported on per-queue per-sla instance of the given subscriber. In many situations, a certain level of aggregation (such as a per-subscriber or HSI ingress and egress traffic) is required to perform meaningful mechanism for pre-paid services.

## Categories Map and Categories

This feature introduces a new object category-map which defines individual aggregates (such as data in and out, video and data, etc.) and their mapping to individual forwarding queues.

The following output depicts a category-map configured in the subscriber management context.

```
*A:ALA-48>config>subscr-mgmt# info
-----
...
    category-map "triple-play" create
        category "data" create
            queue 1 ingress-egress
        exit
        category "video" create
            queue 2 egress-only
        exit
        category "voice" create
            queue 3 ingress-egress
        exit
    exit
category-map "aggr-subscriber-service" create
    category "data-services" create
        queue 1 ingress-egress
        queue 3 egress-only
    exit
exit
...
-----
*A:ALA-48>config>subscr-mgmt#
```

Based on a category-map the system gathers usage information (volume/time) on a per-sla-instance-per-category basis. In order to do so, statistics of all queues forming the category of the given sla-instance are aggregated.

- Single subscriber host (routed CPE) — Single SLA instance.
- Multiple subscriber hosts on the same SAP (bridged CPE) — Single SLA instance. Note that several hosts use the same credit and the renewal of one will cause renewal for all.
- Multiple subscriber hosts on different SAP (bridged CPE) — SLA instance per host.

The per-category usage gathered as described above is compared with per-subscriber-host-per-category credit and when credit is exhausted several actions can be taken.

There are several category-maps pre-configured on the system. The category-map applicable to a given subscriber-host will be derived at the host creation from the RADIUS VSA in an authentication-response, Python script, or static configuration in the local-user-database. All subscriber-hosts belonging to the same subscriber and created on the same SAP (hence, sharing the same sla-instance) must use the same category-map. In case of conflict, (an existing subscriber host has a different category-map than the one derived for the new host) the category-map of the last host will be applied to a given sla-instance. As a consequence, all previous information related to the status of the credit will be lost.

There can be multiple queues aggregated into one category. There can be up to three categories in a category map.

## Quota Consumption

There are two types of quota (credit), volume and time. In volume usage monitoring, the system accumulates byte counters per category-sla-instance and compares it with the assigned quota. Once the credit is exhausted (or threshold for renewal is met) the system attempts to renew it with corresponding management system.

In time-based credit, the distinction between active-usage and active-connection is made by defining an activity-threshold, where an object defines an average data rate under which the subscriber-host is considered silent.

As long as the effective rate of the application usage does not exceed the rate defined by the activity-threshold, the given subscriber host will be considered silent and its corresponding credit will not be used. As long as the application usage exceeds the rate, the application-credit will be consumed (in terms of time).

---

## RADIUS VSA Credit-Control-Quota

The quota in the RADIUS VSA Credit-Control-Quota uses this fixed format:

Alc-Credit-Control-Quota = “<volume quota>|<time quota>|<category name>”

- Where Volume: in bytes (B), kilobytes (K or KB), megabytes (M or MB), gigabytes (G or GB)
- Where Time: in seconds (s), in minutes (m), in hours (h), in days (d) or a combination (5m30s) but there is a restriction; a lower unity may never exceed the higher unity (5m60s is not allowed)

For example, Alc-Credit-Control-Quota = “1G|1h30m|cat1”

Volume quota, as well as time quota, needs to be specified.

- The minimum volume quota is 100 megabytes.
- The minimum time quota is 15 minutes.

## Credit Negotiation Mechanisms


The per-subscriber per-category credit can be obtained by several ways:

- RADIUS during authentication process.
- Static configuration - configured in the **config>subscr-mgmt>category-map>category** context.

Credit can be expressed by either

- Volume
- Time

The renewal of the credit using RADIUS authentication is triggered by credit exhaustion or (if configured) by depletion of the credit to exhausted-credit-threshold level. If this occurs, the system will send a RADIUS authentication message indicating the corresponding category and usage. The following are several possibilities for the RADIUS server response (as shown in figure below):



al 0064

**Figure 113: Threshold Configured/Not Configured**

1. No authentication response — The system will install out-of-credit action after the original credit has been used.
2. Authentication response with reject — The corresponding host is removed after the original credit has been used.

## Action on Credit Exhaustion

3. Authentication response with accept and no credit VSA included — The system will install out-of-credit action.
4. Authentication response with accept and credit VSA included — The out-of-credit will be installed.  
Note that the new credit is always reduced by the amount of credit consumed in time between renewal has been initiated and authentication-respond has been received. In case of a negative result (the newly receive credit is smaller than the amount consumed in the mean-time) the test cr
5. is installed.

In order to identify that the given RADIUS-auth request is related to credit renewal rather than to plain authentication, the node will include empty credit VSAs, depending on categories which has been exhausted. The RADIUS server can identify which category has requested credit renewal.

---

## Action on Credit Exhaustion

System supports configurable actions once the credit for given subscriber is exhausted:

- Sends an SNMP trap and continue (the credit-usage counter is reset).
  - Disconnect.
  - Changes to a pre-defined service level (such as adjusting the queue rate).
  - Blocks the category.
- 

## Action on Error-Conditions

During credit negotiation, the number of errors can occur which can lead to a given subscriber-host category with no new credit renewed. This is different from credit exhaustion where a separate configurable action will be taken. The following occurs:

- Sends an SNMP trap and continues.
- Sends a trap and blocks the category.

## **Applicability of Volume and Time Based Accounting**

Volume and time based accounting is applicable to the ESM mode of operation only. Using credit control concept is not mutually exclusive with other accounting methods. In many network implementations the more traditional accounting methods such as XML file or RADIUS accounting will be still used in a combination with the credit concept but with larger intervals. This is helpful when providing overviews of the average usage and service utilization.

## Subscriber Host Idle Timeout

An idle timeout is the maximum time that a subscriber session can be idle before the session is terminated or a connectivity check is started. Idle timeout applies to PPPoE, PPPoEoA, PPPoA and IPoE hosts.

The time/volume based accounting model is used to configure an idle timeout:

- Create a category-map ([Categories Map and Categories on page 1023](#))
  - Define a category with queues and/or policers to be monitored for activity (packets being forwarded).
  - An activity threshold (in kbps) must be configured for idle timeout to take effect. The activity threshold suppresses background traffic (for example control flows) from activity monitoring.

Example:

```
config>subscr-mgmt
  category-map "idle-timeout" create
    activity-threshold 25
  category "cat-1" create
    queue 1 ingress-egress
  exit
exit
```

- In the sla-profile, associate the category-map and optionally define
  - An idle-timeout (60..15552000 seconds). The default is infinite (no idle-timeout).

The idle-timeout can also be specified from RADIUS in an access-accept or CoA message with the [28] Idle-Timeout attribute. A RADIUS specified idle-timeout overrides the CLI-configured value. The values outside the limits are accepted but rounded to these boundaries.

Attribute ID	Attribute name	Type	Limits	Purpose and Format
28	Idle-Timeout	integer	[60..15552000] seconds	0 = infinite (no idle-timeout) [60..15552000] in seconds For example: Idle-Timeout = 3600

- An idle-action:
  - **shcv-check** — Perform a subscriber host connectivity check (IPoE hosts only). Host connectivity verification should be enabled on the corresponding group-interface for the **idle-action shcv-check** to take effect:

```
configure service ies|vprn service-id subscriber-interface ip-int-name group-  
interface ip-int-name host-connectivity-verify
```



If the shcv check is successful, the subscriber host is not disconnected and the idle-timeout timer is reset to zero. If the shcv check fails, the subscriber host is disconnected (same as terminate).

For PPP hosts, the **idle-action shcv-check** is ignored and has the same effect as “idle-action terminate”

- Terminate (default): disconnect the subscriber hosts
  - IPoE:
    - Delete the subscriber host
    - Send a DHCP release message to the DHCP server
    - Send an Accounting Stop message to the RADIUS accounting server
  - PPP:
    - Delete the subscriber host
    - Send a terminate request message to the CPE
    - Send an Accounting Stop message to the RADIUS accounting server

#### Example

```
config>subscr-mgmt
  sla-profile "sla-profile-1" create
    category-map "idle-timeout"
      category "cat-1" create
        idle-timeout 3600
        idle-timeout-action terminate
      exit
    exit
  exit
exit
```

At host instantiation, a timer is initialized to the idle-timeout value (one timer per sla-profile instance). Each queue or policer in the category is monitored for activity over a fixed polling interval:

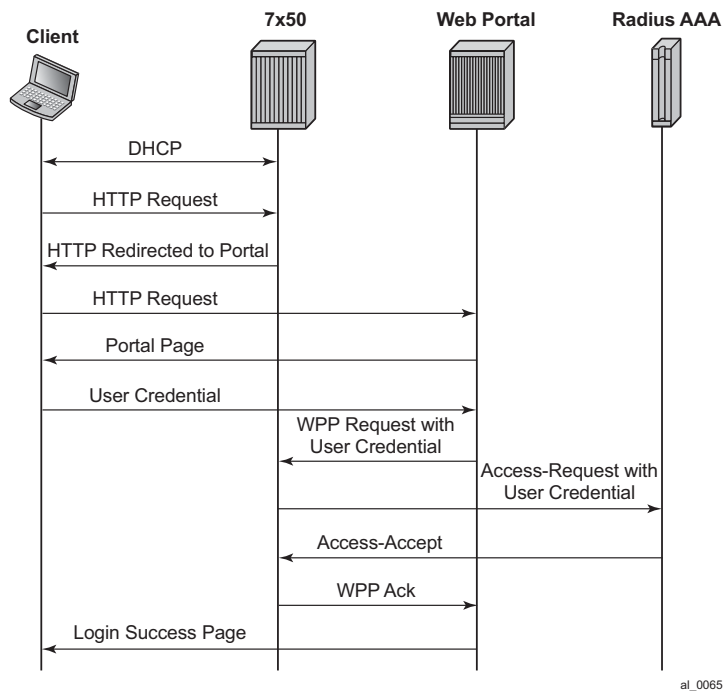
- During the polling interval:
  - if the forwarding rate falls below the configured activity threshold then the timer is deducted by the polling interval (time elapsed).
  - If the forwarding rate is above the configured activity threshold then the timer is initialized to the idle-timeout value.

When the timer becomes zero, the idle-timeout-action is performed for all hosts associated with the SLA-profile-instance (all hosts from a subscriber on a single sap and that share the same sla-profile).

## Web Authentication Protocol (WPP)

The Web Authentication Protocol (WPP) is a protocol running between BNG and Web portal server. WPP is used for web portal authentication of WLAN users (DHCP Host). It can function like a web portal that can trigger BNG to do RADIUS authentication for WLAN users, or send user disconnection notification to BNG.

The [Figure 114](#) illustrates high level of call flow of WPP authentication.



**Figure 114: WPP Authentication**

The following describes WPP authentication call flow:

1. When the WLAN user starts a DHCP exchange with a 7750, the 7750 will create a DHCP host from following configurations:
  - Sub-id is the default sub-id configured in the **sap>sub-sla-mgmt** context.
  - sla-profile/sub-profile/aa-profile will take the configuration from CLI command **grp-if>wpp>initial-sla-profile/initial-sub-profile/initial-app-profile**.
  - IP address from local or external DHCP server will be assigned to the host.
2. When the user sends an HTTP request to visit a web site by browser, the 7750 redirects the HTTP request to the web portal.
3. The portal server sends an authentication page to the WLAN user.

4. WLAN user enters username and password in the authentication page and submit to the Portal Server.
5. The portal server sends a WPP request to router together with the user credentials.
6. The 7750 sends an access-request to RADIUS server with user credentials.
7. RADIUS returns an access-accept if authentication succeeds.
8. The 7750 returns a WPP ack to portal server.
9. If it was access-accept, then the 7750 can optionally override the following host properties:
  - Sub-id: sub-id from RADIUS. If there is no sub-id from RADIUS, then the host will keep using current sub-id.
  - Sla-profile/sub-profile/aa-profile: The system will use the RADIUS server returned values. If the RADIUS server did not return these then the system will try to use the LUDB (in local DHCP server) return values if they are available. If not, the system will try to use the default values configured under SAP.

---

## WPP Configurations

A minimal WPP configurations must include the following:

- WPP portal server — Specifies the name and IP address of the WPP portal server.
  - Enable WPP under the group-interface:
    - WPP portal server that system should listen to.
    - **authentication-policy** on **group-interface** that specifies address of RADIUS server.
    - **def-sub-id** under `sap>sub-sla-mgmt` that is used for DHCP host before user is authenticated by portal server.
    - **initial-sla-profile** and **initial-sub-profile** that are used for the DHCP host before user is authenticated by portal server.
- Note:** **initial-sla-profile** should include a ingress filter that has **http-redirection** entry.

The following is an example configuration:

```
#-----
echo "Web Portal Protocol Configuration"
#-----
wpp
  portals
    portal "portal-1" address 9.9.9.9 create
    no shutdown
  exit
  exit
  no shutdown
  exit
config>service>vprn# info
#-----
```

## Web Authentication Protocol (WPP)

```
subscriber-interface "sub-if" create
  address 192.168.10.1/24
  group-interface "grp-if" create
    dhcp
      server 1.1.1.1
      gi-address 192.168.10.1
      no shutdown
    exit
  authentication-policy "radius-auth"
  sap 1/1/9 create
    sub-sla-mgmt
      def-sub-id "WLAN-User-Unauth"
      no shutdown
    exit
  wpp
    initial-sla-profile "webportal"
    initial-sub-profile "webportal"
    portal router "Base" name "portal-1"
    no shutdown
  exit
exit
exit
...
```

-----

## WPP Triggered Host Creation

In some cases, a 7750 SR can sit behind a Layer 3 device (such as an CMTS), where the 7750 does not participate in client's DHCP process. Such a use case is different from a normal WPP use case where 7750s rely on getting client's DHCP request to create an initial ESM host.

This feature allows the system to create an ESM host upon successful WPP authentication without creating an initial host.

In the above use case (behind a Layer 3 device) the user also need to configure one or more default hosts on the SAP to allow HTTP redirection without an ESM host. The default-host subnet is the user's source subnet and the nexthop address is the Layer 3 device's interface address that connect to the SAP. Users also need to configure the **lease-populate l2-header** command in the **grp-if>dhcp** context to make HTTP redirection with default-host work. The **grp-if>dhcp** context could be shutdown in the meantime.

This feature does not work with wholesale/retail

---

## LUDB Support For WPP

Since R12.0R1, the SR OS supports LUDB lookup for WPP authentication. Users can configure LUDB to return WPP-related attributes such as a portal name, initial-sla-profile/initial-sub-profile, etc. The system could access LUDB when creating the initial host before WPP authentication and the LUDB returned attribute will override the corresponding configuration under the group-interface.

## WPP Multi-Chassis Redundancy Support

The SR OS supports multi-chassis redundancy to WPP. This can be achieved by doing following:

- Create a loopback interface on both 7750 with the same IP address X.
- Use the **track-srrp** parameter while configuring address X to track the corresponding SRRP instance.
- Configure a portal with the same name and same service-id on both nodes to send WPP packets to the destination address.
- Use an route-policy to export X to the routing protocol. The metric the route X can be set is based on the a specified SRRP state (master or non-master) so that master node can advertise route X with a better metric. Then the WPP packet from the portal will be attracted to the master.
- Only the master process WPP packet, however in case of standby node receives (such as routing is still re-converging) the WPP packet, then standby will shunt the WPP back to master.
- WPP hosts will be synced via MCS.

## One-time HTTP Redirection Overview

With this feature enabled, after an ESM host is created, only the FIRST HTTP request from the host will be redirected to a configured URL with specified parameters. Subsequent HTTP request will go through without being redirected.

This feature could be used by service providers to push a web-page to broadband users for purpose of advertisement, announcements, and such.

A **one-time-http-redirection** filter could be configured in **sla-profile**, this filter will be replaced by ingress filter in **sla-profile** after 1st HTTP request is redirected. There is also a RADIUS VSA ( ALC-Onetime-Http-Redirection-Filter-Id) that could be included in access-accept or CoA request to override CLI configuration. The format of ALC-Onetime-Http-Redirection-Filter-Id is **Ingr-v4:filter-id**; for example, **Ingr-v4:1000**. If the the filter-id is 0, then system will replace the current **one-time-http-redirection** filter with ingress filter.


**Note:** In case of CoA, if the host's **one-time-http-filter** has already been replaced then system will just ignore the ALC-Onetime-Http-Redirection-Filter-Id.

If a 7750 SR receives filter insertion via CoA or access-accept when **one-time-http-redirection** filter is still active then the received filter entries will only be applied to the ingress filter. And after 1st http redirection, the update ingress filter will replace the one-time-http-redirection filter.

This feature only supports IPv4 filter.

## ESM over MPLS Pseudowires

This feature allows IPoE and PPPoE (terminated or L2TP tunneled) subscriber sessions to be backhauled through an Ethernet aggregation network using MPLS pseudowires terminating directly on the BNG. The MPLS pseudowire originates from the first hop aggregation PE (referred to as access PE) upstream of the Access-Node (or directly from a multi-service AN), and terminates on the BNG. Multiple subscriber sessions from a given access-port on the Access-PE can be backhauled over a single P2P MPLS pseudowire towards the BNG. This capability allows the network to scale and does not require a MPLS pseudowire per subscriber between Access-PE and the BNG. The access-port on the Access-PE can be dot1q, q-in-q or NULL encapsulated. The BNG terminates the MPLS pseudowire, decapsulates the received frames, and provides ESM functions including HQoS, without requiring an internal or external loopback. Each MPLS pseudowire is represented on the BNG as a “PW-port” for which SAPs are created. A PW-port can be configured with capture SAP. Both static and managed SAPs are supported. The underlying Ethernet port is required to be in hybrid mode. The feature set is supported for IOM3-XP and HSMDAv2. This feature is supported on the 7750 SR and 7450 ESS in mixed mode.



al\_0066

Figure 115: ESM over MPLS Pseudowire Example



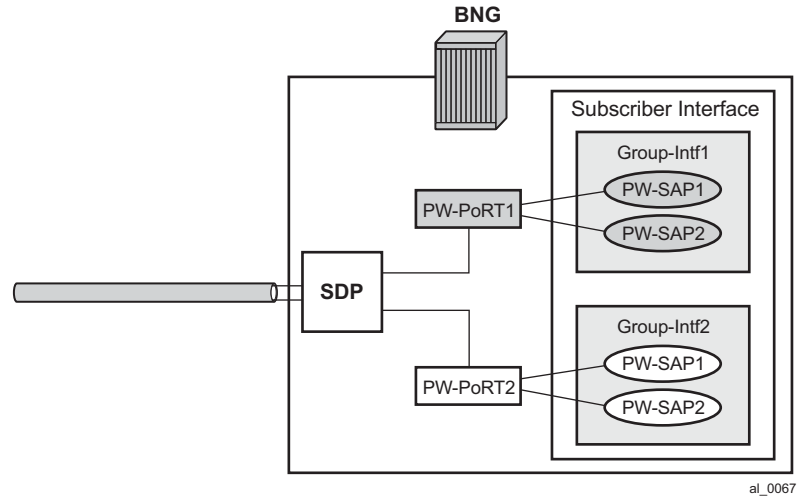
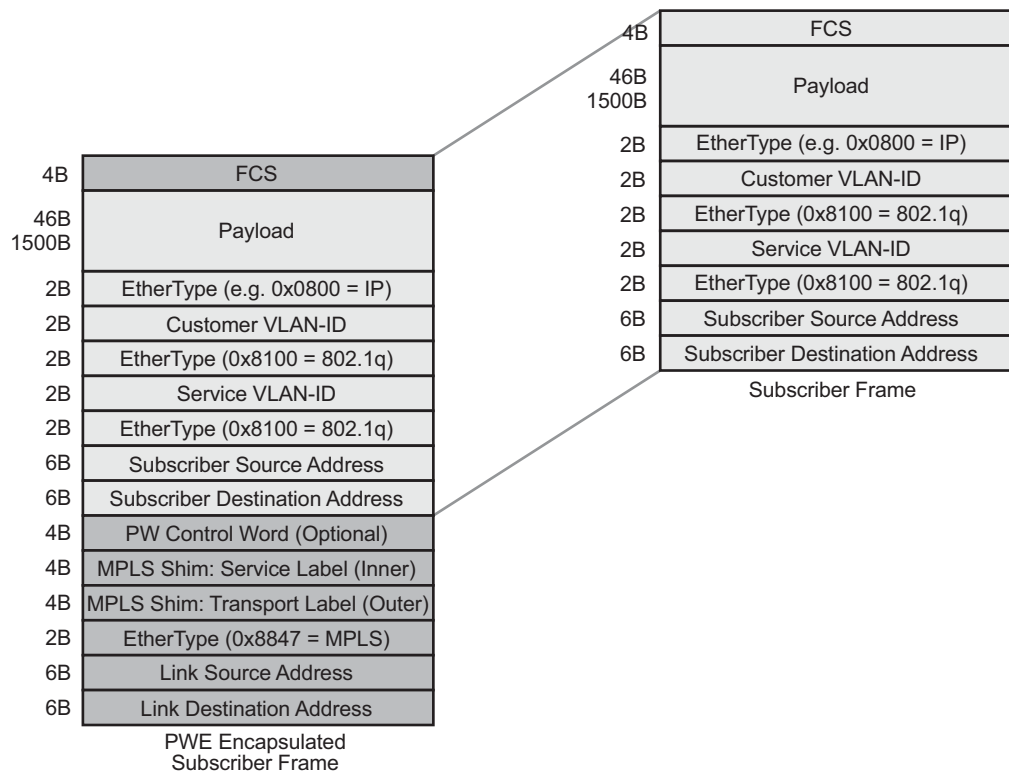


Figure 116: Group Interface Example

Encapsulation

The subscriber frame encapsulated within the pseudowire is shown in [Figure 117](#). Optional control word is not supported. The SDP could be LDP, RSVP or LDP over RSVP. Hash labels are not supported. SDP is bound to a port or a LAG. In case the SDP is re-routed, the corresponding PW-ports are operationally brought down. The PW-ports are associated with the SDP by configuration.



al\_0068

Figure 117: Subscriber Frame with PWE Encapsulation

ESM Configuration with PW-Ports and PW-SAPs

BNG requires configuration for PW-ports. The VC-label for configured PW-port is dynamically signaled using T-LDP with the far-end A-PE. The configuration for the PW-port includes the port-id (unique identifier within the chassis), vc-id (virtual circuit identifier, which is signaled to the peer), and the vc-type (Ether or VLAN, which is signaled to the peer). The vc-id and vc-type must match with the configuration of the PW on the far-end. The vc-id MUST be unique across PW-ports. The encapsulation type (dot1Q or q-in-q) on the PW-port is configurable. The default value

for vc-type is Ether, and the default encapsulation type is dot1Q. With vc-type vlan, the vc-vlan-tag can be configured. vc-type vlan forwarding mode can only be set if encapsulation type is dot1Q. On the BNG, the vc-vlan-tag is only relevant for transport, and not service delineation and ESM. On the BNG, with vc-type vlan configured on the PW-port, the configured vc-type-vlan tag is inserted when forwarding traffic into the PW (i.e. in downstream direction), and top dot1Q tag is stripped when forwarding traffic out of the PW (i.e. in upstream direction). On the BNG, with vc-type ether configured on the PW-port, the received tags (max two, including any provider tag inserted by the far-end) are preserved and passed for PW-SAP lookup or creation. In the downstream direction the PW-SAP tags are inserted and passed back to the far-end.

The following output displays an ESM configuration with PW-ports and PW-SAPs.

```
config>service#
  customer 1 create
    description "Default Customer"
  exit
  sdp 1 mpls create
    description "Default sdp description"
    far-end 10.20.1.2
    ldp
    keep-alive
    shutdown
  exit
  binding
    port 1/1/3
    pw-port 11 vc-id 11 create
      vc-type vlan      #### default encaps-type dot1Q
      no shutdown
    exit
    pw-port 44 vc-id 2 create #### default vc-type Ether, encaps-type dot1Q
      no shutdown
    exit
  exit
  no shutdown
exit
vppls 1 customer 1 vpn 1 create
  sap pw-11:* capture-sap create
    trigger-packet arp dhcp dhcp6 pppoe #
    msap-defaults
      group-interface "grpif-pw-11"
      policy "msap-policy1"
      service 3
    exit
    authentication-policy "base_authpolicy"
  exit
  no shutdown
exit

ies 3 customer 1 vpn 3 create
  description "Default ies description for sevice id 3"
  subscriber-interface "subif" create
    address 11.11.1.1/16
    address 44.44.1.1/16
    group-interface "grpif-pw-11" create
      arp-populate
      dhcp
      server 10.20.1.2
```

```
        gi-address 11.11.1.1
        no shutdown
    exit
    authentication-policy "base_authpolicy"
    sap pw-11:11 create
        sub-sla-mgmt
            def-sub-profile "sub_prof_1"
            def-sla-profile "sla_prof_1"
            no shutdown
        exit
    exit
exit
group-interface "grpif-pw-44" create
    arp-populate
    dhcp
        server 10.20.1.2
        gi-address 11.11.1.1
        no shutdown
    exit
    sap pw-44:44 create
        sub-sla-mgmt
            def-sub-profile "sub_prof_1"
            def-sla-profile "sla_prof_1"
            no shutdown
        exit
    exit
    no shutdown
exit
```

## QoS Support

QoS is supported for ESM over PW-SAPs as with ESM over regular SAPs, and includes currently supported models.

- FC to queue mapping
- H-QOS
  - Per-subscriber HQOS (service scheduler child to port-scheduler parent).
  - PW-SAP queues attached to H-QOS scheduler by parent statement.
  - Scheduler attached to port scheduler by “port-parent” statement.
- Direct service queue to port-scheduler.
  - Aggregate-rate-limit.

## Bandwidth Control at PW-Port Level via Vport

Bandwidth control per PW-port (per AN or per AN/ per service), via vport.

- The vport can be created on the binding port.
- The vport can be associated with the PW-port either via static assignment or dynamic selection via inter-dest-id (returned from RADIUS or DHCP for a host).
- Aggregate-rate-limit can be configured to shape the egress traffic across all hosts associated with the vport via inter-dest-sting match or via static association of underlying PW-port with the vport.

The following output displays a dynamic vport selection based on an inter-dest-id configuration.

```
config>
  Port 1/1/1
    ethernet
      mode hybrid
      encaps-type dot1Q
      mtu 1540
      access
        egress
          vport "v1" create
            agg-rate-limit 1000
            host-match dest "dslam-1"      ##### hosts will be associated with
            exit                          ##### vport based on inter-dest-id
          exit
        exit
      exit
    exit
  exit

config>service>sdp>binding
  pw-port 11 vc-id 11 create
    egress
      shaping int-dest-id "dslam-1"      ##### dynamic vport selection based on
                                          ##### int-dest-id.
```

The following output displays a static assignment of PW-port to vport configuration.

```
config>
  Port 1/1/2
    ethernet
      mode hybrid
      encap-type dot1Q
      access
        egress
          vport "v2" create
            agg-rate-limit 1000
          exit
        exit
      exit
    exit
  exit

config>service>sdp>binding
  pw-port 20 vc-id 20 create
    egress
      shaping vport "v2"      #### static assignment of pw-port to vport.
    exit
  exit
```

---

## Last Mile Shaping

With normal Ethernet aggregation in the next-mile, when last-mile shaping is on, fixed encapsulation-offset is calculate based on the last-mile encapsulation type and the next-mile encapsulation (26 Bytes with q-in-q). This offset is applied to the frame, and the ATM overhead is then dynamically calculated on the adjusted size. The resulting dynamically calculated overhead in the data-path is then applied to the queue-rates and the subscriber aggregate-rate.

With this feature of backhauling subscriber sessions using MPLS PW in the aggregation network, the encapsulation is shown in Fig 3. The last mile does not see any MPLS PW overhead. The next-mile includes overhead due to the PW encapsulation shown in [Figure 117](#). Therefore, when last mile shaping is enabled, the fixed encapsulation-offset is calculated based on the difference between last-mile encapsulation type and next-mile encapsulation, The next-mile encapsulation takes into account the additional PW overhead, which includes:

14B Ethernet header + [4B] (optional network interface Q-tag) + MPLS Labels (variable)

In the data-path the actual PW encapsulation overhead, taking into account the MPLS labels which could be variable (with FRR or PHP) is tracked, and is applied to the computed “encapsulation offset”. This adjusted “encapsulation offset” is applied to the frame. The ATM overhead is then dynamically calculated on the adjusted size, and applied for last mile shaping (to queue-rates and subscriber-aggregate-rate). Note that there is no change from ESM over normal SAPs, in how last-mile shaping is triggered or how the last mile encapsulation type is determined (via configuration in egress context of subscriber profile or dynamically learned from Access-Loop-Encapsulation sub-TLV in vendor specific PPPoE tags).

## BNG Redundancy with ESM over Pseudowire

This feature provides support for stateful BNG redundancy (when the far-end aggregation PE (A-PE) is dual-homed to two BNGs terminating subscriber sessions over MPLS pseudowires (pws) that are initiated from the A-PE and provides ESM). Subscriber state between BNGs is synced using MCS.

### EPIPE Based Aggregation Service

For an EPIPE based aggregation service, the redundancy is based on active/standby PWs from A-PE to dual BNGs. A-PE signals active/standby pseudowire status to peer BNGs. An SRRP instance per PW-Port (group-interface) is required on the BNG with messaging SAP on each PW-Port. BNG terminating active PW assumes the mastership for the SRRP instance on the corresponding PW-Port. SRRP state is tied to the state of the messaging SAP. The messaging SAP goes down when the underlying PW-Port goes down, based on PW status bit signaled by the A-PE.

In this model, there is no SRRP message exchange between the two BNGs, as there is no L2 path between the BNGs. The purpose of SRRP is to get SRRP-aware routing for subscriber routes and managed routes, and/or to be able to use the redundant (shunt) interface. Downstream traffic for a subscriber that ingresses the backup BNG can only be shunted to the active BNG, if the corresponding subscriber-interface on the backup BNG is operationally UP. This can be achieved by creating a second empty group-interface (without SAPs) on the same subscriber-interface with the knob 'oper-up-while-empty' configured. Multiple PWs with endpoint configuration is not supported on the BNG.




Figure 118: BNG Redundancy Based on Active/Standby PW Signaling

### Sample Configuration on Master BNG

```
config>
  pw-port 2 create
  exit

config>redundancy#
  multi-chassis
  peer 10.20.1.3 create
  source-address 10.20.1.2
  sync
  srrp
  sub-mgmt ipoe pppoe
  port pw-2 sync-tag "tag2" create
  exit
  no shutdown
  exit
  no shutdown
  exit
  exit
  exit

config>service>ies#
  redundant-interface "redundant-interface" create
  address 10.10.30.2/24 remote-ip 10.10.30.3
  spoke-sdp 23:1000 create
  no shutdown
  exit
  exit

config>service#
  sdp 1 mpls create
  far-end 10.20.1.2
  ldp
  keep-alive
  shutdown
  exit
  binding
  port 1/1/1
  pw-port 2 vc-id 2 create
  vc-type vlan      ##### default encaps-type dot1Q
  no shutdown
  exit
  exit
  no shutdown
  exit

config>service#
  subscriber-interface "subif" create
  address 11.11.1.2/16 gw-ip-address 11.11.1.1 populate-host-routes
  group-interface "grpif" create
  authentication-policy "base_authpolicy"
  redundant-interface "redundant-interface"
  sap pw-2:1000 create
  description "sap-grp-3"
  exit
  srrp 1 create
  message-path pw-2:1000
```



```

        no shutdown
    exit
    arp-host
        host-limit overall 8000
        min-auth-interval 1
        no shutdown
    exit
exit
exit
exit

```

---

## Sample Configuration on Slave BNG

```

config>
    pw-port 2 create
    exit
config>redundancy#
    multi-chassis
        peer 10.20.1.2 create
            source-address 10.20.1.3
            sync
                srrp
                sub-mgmt ipoe pppoe
                port pw-2 sync-tag "tag2" create
            exit
        exit
        no shutdown
    exit
exit
exit
config>service>ies#
    redundant-interface "redundant-interface" create
        address 10.10.30.3/24 remote-ip 10.10.30.2
        spoke-sdp 32:1000 create
            no shutdown
        exit
    exit
config>service#
    sdp 1 mpls create
        far-end 10.20.1.2
        ldp
        keep-alive
        shutdown
    exit
    binding
        port 1/1/1
        pw-port 2 vc-id 2 create
            vc-type vlan        #### default encaps-type dot1Q
            no shutdown
        exit
    exit
    no shutdown
exit
config>service#
    subscriber-interface "subif" create
        address 11.11.1.3/16 gw-ip-address 11.11.1.1 populate-host-routes

```

```
group-interface "grpif" create
authentication-policy "base_authpolicy"
redundant-interface "redundant-interface"
sap pw-2:1000 create
description "sap-grp-3"
exit
srrp 1 create
keep-alive-interval 1
message-path pw-2:1000
no shutdown
exit
arp-host
host-limit 8000
min-auth-interval 1
no shutdown
exit
exit
group-interface "dummy" create
oper-up-while-empty
exit
exit
exit
```

---


## Sample Configuration on A-PE

```
config>service>epipe#
description "Default epipe description for service id 103"
service-mtu 1492
service-name "XYZ Epipe 103"
endpoint "x" create
standby-signaling-master
exit
sap 1/1/3 create
description "Default sap description for service id 103"
exit
spoke-sdp 1:2 endpoint "x" create
description "Description for Sdp Bind 1 for Svc ID 103"
precedence primary
no shutdown
exit
spoke-sdp 2:2 endpoint "x" create
description "Description for Sdp Bind 2 for Svc ID 103"
no shutdown
exit
no shutdown
```

## VPLS Based Aggregation Service

With VPLS based aggregation service from A-PE, normal SRRP message exchange can take place between the primary and backup BNGs. Master-ship decision and switch-over is based on SRRP. SRRP instance is configured per group-interface corresponding to PW-Port. Fate-sharing groups (FSG) can be configured for a set of SRRP instances (for example, SRRP instances corresponding to PW-Ports sharing the same subnet). Standard **oper-group** *grp-id* would need to be configured with messaging SAPs for all PW-Ports that are in the same FSG, and **monitor-oper-group** *grp-id* would need to be configured under each SRRP instance in same FSG. Existing SRRP support defined in Triple-play services guide for ESM over regular group-interfaces and subscriber SAPs is applicable identically to ESM over PW-Ports and PW-SAPs.

**Note:** With PW over ESM, redundancy in the aggregation network based on MC-LAG between A-PE and dual BNGs is not supported.



**Figure 119: BNG Redundancy with VPLS Based Aggregation Service**

## Sample BNG Redundancy (SRRP/MCS) Configuration with VPLS Service on A-PE

```
config>
  pw-port 1 create
  exit

config>redundancy#
  multi-chassis
    peer 10.20.1.2 create
      source-address 10.20.1.3
      sync
      srrp
      sub-mgmt ipoe pppoe
      port pw-1 sync-tag "tag1" create
      exit
    exit
  no shutdown
  exit
exit
exit
exit

config>service>ies
  redundant-interface "red-1-1" create
    address 1.1.1.2/24 remote-ip 1.1.1.1
    spoke-sdp 1:1 create
      no shutdown
    exit
  exit

  subscriber-interface "sub-1-1" create
    address 20.1.2.2/16 gw-ip-address 20.1.255.254 track-srrp 1
    address 20.2.2.2/16 gw-ip-address 20.2.255.254 track-srrp 2
    dhcp
      gi-address 20.1.2.2
    exit
    group-interface "grp-1-1-1" create
      srrp-enabled-routing
      arp-populate
      dhcp
        server 10.20.1.2
        trusted
        lease-populate 32767
        client-applications dhcp ppp
        gi-address 20.1.2.2
        no shutdown
      exit
      authentication-policy "iesAuthPol"
      redundant-interface "red-1-1"

      sap pw-1:1.1 create
        sub-sla-mgmt
          def-sub-profile "sub_prof_1"
          def-sla-profile "sla_prof_1"
          no shutdown
        exit
      sap pw-1:4000.1 create
        oper-group "1"
```

```

exit
srrp 1 create
  gw-mac 00:00:5e:00:01:01
  keep-alive-interval 50
  message-path pw-1:4000.1
  monitor-oper-group "1" priority-step 10
  no shutdown
exit
exit

```

---

## A-PE configuration with VPLS Aggregation Service (A-PE1)

```

config>service
  customer 1 create
    description "Default customer"
  exit
  sdp 1000 mpls create
    far-end 10.20.1.2
    lsp "lsp_1"
    path-mtu 1600
    keep-alive
    no shutdown
  exit
  sdp 1002 mpls create
    far-end 10.20.1.3
    lsp "lsp_3"
    path-mtu 1600
    keep-alive
    no shutdown
  exit
  vpls 1 customer 1 create
    service-mtu 1600
    stp
    sap 1/1/2 create // to Access-Node
    exit
    sap 1/1/3 create; //to A-PE2
    exit
    spoke-sdp 1000:1 create // to BNG1
    no shutdown
    exit
    no shutdown
  exit
exit

```

---

## A-PE Configuration with VPLS Aggregation Service (A-PE2)

```

config>service
  customer 1 create
    description "Default customer"
  exit
  sdp 1002 mpls create
    far-end 10.20.1.3
    lsp "lsp_2"

```

## BNG Redundancy with ESM over Pseudowire

```
path-mtu 1600
keep-alive
no shutdown
exit

vpls 1 customer 1 create
service-mtu 1600
stp
sap 1/1/2 create // to Access-Node
exit
sap 1/1/3 create; //to A-PE1
exit
spoke-sdp 1002:1 create // to BNG2
no shutdown
exit
no shutdown
exit
exit
```

## Show Commands Related to Active/Standby Pseudowire on Dual BNGs

The following example shows SRRP status, subscriber host, and routing information on master BNG:

```
A:Dut-B>config>redundancy# show srrp 1
```

```
=====
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            : 10.20.1.2
Service ID           : VPRN 3
Group If             : grpif                  MAC Address         : 1c:85:ff:00:00:00
Grp If Description   : N/A
Grp If Admin State   : Up                      Grp If Oper State   : Up
Subscriber If        : subif
Sub If Admin State   : Up                      Sub If Oper State   : Up
Address              : 11.11.1.2/16            Gateway IP           : 11.11.1.1
Redundant If         : redundant-interfa*
Red If Admin State   : Up                      Red If Oper State   : Up
Address              : 10.10.30.2/24
Red Spoke-sdp        : 23:1000
Msg Path SAP         : pw-2:1000
Admin Gateway MAC    :                      Oper Gateway MAC    : 00:00:5e:00:01:01
Config Priority       : 100                    In-use Priority      : 100
Master Priority       : 100
Keep-alive Interval  : 1 deci-seconds          Master Since        : 05/29/2012 07:22:26
Fib Population Mode  : all
VRRP Policy 1        : None                    VRRP Policy 2       : None
=====
* indicates that the corresponding row element may have been truncated.
```

```
A:Dut-B>config>redundancy# show service id 3 arp-host
```

```
=====
ARP host table, service 3
=====
IP Address      Mac Address      Sap Id      Remaining      MC
                MC               Time
                Stdbdy
-----
11.11.1.11      00:80:00:00:00:01 [pw-2:11]    03h35m47s
11.11.1.12      00:80:00:00:00:02 [pw-2:12]    03h35m47s
-----
Number of ARP hosts : 2
=====
```

```
A:Dut-B>config>redundancy# show router 3 route-table 11.11.1.11
```

```
=====
Route Table (Service: 3)
=====
Dest Prefix[Flags]      Type      Proto      Age      Pref
Next Hop[Interface Name]      Metric
-----
```

## BNG Redundancy with ESM over Pseudowire

```
11.11.1.11/32                               Remote Sub Mgmt 00h24m26s  0
      [grpif]                                0
-----
No. of Routes: 1
Flags: L = LFA nexthop available    B = BGP backup route available
      n = Number of times nexthop is repeated
=====
A:Dut-B>config>service>vprn#
```

The following shows SRRP status, subscriber host, and routing info in slave BNG:

```
A:Dut-C>config>redundancy# show srrp 1
```

```
=====
SRRP Instance 1
=====
Description      : (Not Specified)
Admin State      : Up                      Oper State      : initialize
Preempt         : yes                    One GARP per SAP : no
Monitor Oper Group : None
System IP       : 10.20.1.3
Service ID      : VPRN 3
Group If        : grpif                  MAC Address     : 1c:87:ff:00:00:00
Grp If Description : N/A
Grp If Admin State : Up                  Grp If Oper State: Down
Subscriber If    : subif
Sub If Admin State : Up                  Sub If Oper State: Up
Address         : 11.11.1.3/16          Gateway IP      : 11.11.1.1
Redundant If     : redundant-interfa*
Red If Admin State : Up                  Red If Oper State: Up
Address         : 10.10.30.3/24
Red Spoke-sdp   : 32:1000
Msg Path SAP    : pw-2:1000
Admin Gateway MAC :                      Oper Gateway MAC : 00:00:5e:00:01:01
Config Priority  : 1                      In-use Priority  : 1
Master Priority  : 1
Keep-alive Interval : 1 deci-seconds    Master Since     : 05/29/2012 07:22:26
Master Down Interval: 0.000 sec (Expires in 0.000 sec)
Fib Population Mode : all
VRRP Policy 1    : None                  VRRP Policy 2    : None
=====
* indicates that the corresponding row element may have been truncated.
```

```
A:Dut-C>config>redundancy# show service id 3 arp-host
```

```
=====
ARP host table, service 3
=====
IP Address      Mac Address      Sap Id      Remaining      MC
                Stdby
-----
11.11.1.11      00:80:00:00:00:01 [pw-2:11]    03h38m01s     Yes
11.11.1.12      00:80:00:00:00:02 [pw-2:12]    03h38m02s     Yes
-----
Number of ARP hosts : 2
=====
```

```
A:Dut-C>config>redundancy# show router 3 route-table 11.11.1.11
```



```
=====
Route Table (Service: 3)
=====
Dest Prefix[Flags]                Type    Proto    Age          Pref
      Next Hop[Interface Name]                Metric
-----
11.11.1.11/32                    Remote  Sub Mgmt  00h22m03s    0
      [redundant-interface]                      0
-----
No. of Routes: 1
Flags: L = LFA nexthop available    B = BGP backup route available
      n = Number of times nexthop is repeated
=====
```

## On-Demand Subnet Allocation (ODSA)


---

### DHCP pool *subnet-binding-key*

To share a DHCP pool among BNG, the DHCP server will first require a **gi-address**. This feature is used in conjunction with **use-gi-address** from the **local-dhcp-server** and the scope should be *pool* to allow allocation of all subnets within the pool. The DHCP discovery must contain any of the three vendor specific options in Option 82: service ID, service-plus-system ID, or custom string. The intent is to bind a subnet from the shared pool to one of the three required parameters from the vendor specific option (VSO).

When starting to use the shared pool for the first time, a DHCP request should arrive with one of the three required VSOs. The DHCP server looks for a free subnet to bind to the VSO, and then an address is offered to the subscriber. Only subscribers utilizing the same DHCP VSO are allowed to request addresses from the same registered subnet. All new DHCP discovery VSOs are matched against VSOs bound to existing subnets. For the case where a DHCP discovery VSO matches a VSO bound to a subnet, an address from the subnet is offered to subscriber until exhaustion. Once exhausted, the DHCP looks for a free available subnet (if any) for binding and the same process continues until the new subnet is exhausted. Subscribers with non-matching DHCP VSOs are prohibited to request address from the bounded subnet. For the case where a DHCP discovery VSO fails to match any of the current binding, a new unbound subnet is searched for binding (if any).

Each subscriber interface sharing a DHCP pool, must utilize a unique **gi-address**. This is to ensure that the DHCP offer can correctly route back to the subscriber interface as shown in [Figure 120](#). The number of subscriber interfaces sharing the DHCP pool, must have at least the equivalent number of subnets in the DHCP pool. The number of subnets available for sharing should always be equal or greater than the number of subscriber interfaces.



**Figure 120: Subscriber Interfaces Sharing a DHCP Pool**

## ODSA Subnet Advertisement and Routing

To avoid advertising the same IP subnet from multiple subscriber interfaces. A combination of router aggregation and route policies are used to ensure BNGs do not advertisement overlaps subnets. If BNG1, BNG2, and BNG N are sharing a DHCP pool with subnet A,B,C,...,Z. All BNGs must first provision the same subnets on their subscriber interfaces. These subscriber interfaces must not be included in any IGP interfaces, because that will lead to multiple BNG advertising the same subnets which is undesirable. The key is to advertise a subnet only if a subscriber has successfully allocated an address from that particular subnet. Aggregate route can be used to accomplish this. First, the aggregation route will specify a smaller (more specific) subnet of the subscriber interface. For example if Subnet A is 1.1.1.1/24 then the aggregation route would be split into the next smallest subnet 1.1.1.0/25 and 1.1.1.128/25 respectively. A Route policy is then used to advertise the two /25 into IGP only if these aggregate routes become active. A subscriber that has successfully obtained an IP address from the shared pool will activate the aggregate route in the RIB. Once activated in the RIB, the route policy exports these routes by means of IGP or BGP. This in conjunction with ODSA, allows a subnet to be shared between BNGs and provides proper routing for all subscriber traffic.

## ODSA with SRRP

For PPPoE SRRP setups, it is mandatory to use “string” as the subnet-binding key. For IPoE SRRP setups, “string” is preferred but not a mandatory requirement. The benefits of using “string” as a binding key are discussed later in the failover and recovery sections.

For an ODSA PPPoE SRRP setup, there are two mandatory requirements. They are:

- Configure a DHCP vendor specific string under the group interface PPPoE option.
- ODSA DHCP pools must use “string” as the subnet-binding key.

Each group interface can have its own customized string. This will result in each group-interface requesting for its own subnet. Another possibility is to share subnets among all group interfaces under a single subscriber interface. This is accomplished by using the same string on all group-interface under the single subscriber interface. A pair of SRRP group interface between two BNGs should use the same string.

For an IPoE setup, both SRRP and non SRRP, it is possible to insert a custom string via the DHCP relay (configured under the dhcp-relay option). Remember to also use “string” as the subnet-binding.

---

## ODSA SRRP Failover DHCP Behavior

ODSA allows the option of binding a subnet to any of 3 following keys: system-id, system-id + svc-id, or string. For SRRP setups, while PPPoE only allow the use of “string”, IPoE allow the use of any of the three keys for subnet binding. Notice that first two keys both use the element “system-id” for binding. During a SRRP failover, the slave system-id takes over and the new system-id will in turn bind to new subnets. Existing subnet bounded by the old master system-id are non-accessible even if there are still free addresses available. Subscribers created before the failover still requires DHCP renews, requests, and releases of their addresses. When the DHCP server receives these, renews, requests, and releases, the gi-address and system-id might no longer match. The DHCP server will still answer these DHCP messages as long as the IP and MAC address matches the one registered on the server. This allows the subscriber to experience a seamless connection during a SRRP failover.

The other option is subnet-binding with “string” which is mandatory for ODSA PPPoE SRRP setup and optional for ODSA IPoE SRRP setup. Unlike the system-id, the “string” is customizable and can match between two SRRP group interfaces across two different BNGs. During a failover, the subnet-binding key will remain the same. New subscribers can reuse existing subnet already bound by the group interface. No free addresses are wasted. The key requirement is that each pair of SRRP group interface must have a custom unique string.

## ODSA SRRP Recovery DHCP Behavior

When SRRP is repaired, one of the nodes will become the slave. All DHCP relays will come from the SRRP master. For example, if node 1.1.1.3 becomes the slave, then all DHCP relay messages will have the gi-address of 1.1.1.2. Old subscribers that utilize 1.1.1.3 to retrieve DHCP address still require DHCP requests, renews, and releases from the DHCP pool. In the case of IPoE, DHCP server will allow old subscribers with matching IP and MAC to perform DHCP requests, renews, and releases even though the system-id does not match the previous one. This is to ensure subscribers will have uninterrupted services during a recovery.

To recover the addresses/subnets from the slave node, DHCP drain can be used to ensure that IP addresses are released back to the pool when the leases expire. Otherwise, it is best to wait for the DHCP leases to expire for subscribers to ensure that services are uninterrupted. Once expired, subscribers will route through the proper Master SRRP using the correct system-id to retrieve DHCP addresses.

In the case where the subnet-binding key is “string”, the SRRP recovery is more seamless. Although utilizing a different gi-address, DHCP relays will utilize the same “string” for all DHCP transaction. The same subnets will continue to supply, renew, and release DHCP addresses. There is no need to manually drain unused subnets because the same subnets will be used.

## Logical Link Identifier (LLID)

This feature enables service providers to track subscribers on the basis of a virtual-port known as logical line ID (LLID). The LLID (an alphanumeric string) is a logical identification of a subscriber line. Mapping of physical line of a subscriber to LLID is performed via pre-authentication with a separate AAA server than the AAA server used for authenticating the subscriber session during normal access authentication.

LLID serves the purpose of abstracting the physical line of the user from the ISP. If the user moves to a new physical line, the RADIUS server database maintaining the physical line of the subscriber to LLID is updated. Because a subscriber's LLID remains same regardless of subscriber's physical location, using LLID gives service provider a stable and secure identifier for tracking subscriber.

The local user database assigned to the PPPoE node under the group interface can have both a pre-authentication policy and an authentication policy. The purpose of the pre-authentication policy is to retrieve the LLID from the AAA server. The pre-authentication will only extract the calling-station-id attribute (0x31) which is used as the LLID, anything else returned during pre-authentication are simply ignored. If the pre-authentication is missing the LLID, the session will move on to the authentication policy. In the authentication policy that follows, it is possible to use the LLID as the calling-station ID.

It is possible to convey LLID from the LAC to the LNS. The LLID is retrieved through PPPoE pre-authentication where the returned RADIUS attribute calling station ID is used as the LLID. This LLID is selectable attribute in L2TP as a calling-number (AVP 22) to be passed from LAC to LNS. At the LNS, the subscriber calling station number is retrieved from AVP 22 and can be included as an attribute during authentication.

# Open Authentication Model for DHCP and PPPoE Hosts

---

## Terminology

LUDB – Local User Database configured within 7x50

- IP Address Assignment via DHCP Relay — IP address assignment request (DHCP or IPCP) from the host is relayed to an internal or external DHCP server. Gi-address must be present in this relayed request while the pool name is optional. The internal 7x50 DHCP server may select the IP address from its local pool based on the gi-address or based on the pool-name present in the request. The IP address selection method is configuration dependent. Third party DHCP servers may consider additional fields in IP address selection process (mac address, circuit-id, etc).
  - IP Address Assignment via DHCP Proxy — A preconfigured IP address in LUDB or RADIUS server is handed out to the host via a 7x50 DHCP proxy function. This proxy function responds natively using DHCP protocol to the IPoE host. Although PPPoE hosts are not utilizing DHCP protocol, the DHCP proxy functionality within 7x50 is still needed for successful IP address delegation to PPPoE hosts.
- 

## LUDB and RADIUS Access Models

During the subscriber-host instantiation phase in 7x50, various parameters for the hosts are gathered from a single or multiple sources. These parameters represent the level of service within 7x50 to which the host is entitled. Some of the parameters are mandatory for subscriber instantiation while others are optional. The following lists the parameter sources in the order of priority:

- LUDB
- RADIUS
- DHCP Server => DHCP server directly queries LUDB
- DHCP option processed on DHCP ACK that is indicated in subscriber identification policy.
- Extraction from the DHCP Ack via Python (IPv4 only)
- Defaults that are statically configured on the 7x50 node (SAP, msap-policy, capture-sap, and subscriber-identification-policy).

In most cases, the host IP address assignment process is controlled by the parameters returned via LUDB or RADIUS. As such, the IP address delegation is integral part of the host instantiation process and will consequently be described in the following sections.

## No Authentication

IPoE and PPPoE v4/v6 hosts on static SAPs can be instantiated without the need to access LUDB or RADIUS server. In this case, the default subscriber host parameters (sla-profile, sub-profile, subscriber-id) must be provisioned statically under the SAP. The IP address assignment is provided by internal or external DHCP server. The IP address selection on 7x50 based DHCP server is based on the gi-address while third party DHCP servers may provide additional means to select the IP address (*mac-address, circuit-id*, etc.).

A DHCP pool name cannot be provided by 7x50 DHCP relay agent, since the LUDB and/or RADIUS are not utilized.

This model does not support IP address delegation via DHCP Proxy function since there is no LUDB or RADIUS server available that can supply pre-configured IP address.

Host instantiation without LUDB or RADIUS access on dynamic VLANs (capture SAP and consequently mSAP) is not supported.

---

## LUDB Only Access

Subscriber-host authentication, identification and IP address assignment can be performed via LUDB without the need to access the RADIUS server.

The LUDB is normally configured under the group-interface>ppp/dhcp hierarchy and can provide subscriber-identification parameters as well as IP addressing parameters:

Pool names for DHCP relay function (IPv4, IPv6 IA-NA, IPv6 IA-PD)

Fixed IP addresses – IPv4, IPv6 IA-NA, IPv6 IA-PD and IPv6 SLAAC prefix.

In case of capture SAP, the LUDB name configured under the capture SAP must match the LUDB name under the group-interface>dhcp/ppp hierarchy. If the LUDB names do not match, the subscriber-host instantiation will fail.

---

## LUDB Access via DHCPv4 Server

In case that the IPv4 addressing assignment is facilitated by the DHCPv4 relay and an internal DHCPv4 server, the DHCPv4 server itself can query the LUDB for IPv4 address information. LUDB can provide a v4 pool name and IPv4 DHCP options to the DHCPv4 server or it can instruct it to use the gi-address as the IPv4 address selection mechanism.

ESM strings can also be provided via LUDB queried by the DHCPv4 server.



If LUDB access via DHCPv4 server is provided in addition to other authentication means (another LUDB under the group-interface, or RADIUS server), the ESM strings from the LUDB under the grp-interface or from the RADIUS server will have priority over the ESM strings configured under the LUDB accessed by the DHCPv4 server. On the other hand, the IPv4 addressing information will have the highest priority from the LUDB accessed directly by the DHCPv4 server.

Accessing LUDB directly via DHCPv4 server should be used in rare and exceptional cases.

LUDB access under the group-interface, possibly complemented by the RADIUS server will provide necessary means for subscriber-host instantiation in majority of use cases.

---

## RADIUS Only Access

Similar to LUDB-only access, RADIUS server can provide all the necessary information for subscriber-host instantiation, including the IP addressing parameters (pool names or IP addresses/prefixes). Authentication-policy which defines the RADIUS access must be applied to the group-interface.

In case of capture SAP, the authentication policy must be applied under the capture SAP. This authentication policy name must match the authentication policy name that is configured under the group-interface. Otherwise, the host instantiation will fail.

---

## Consecutive Access to LUDB and RADIUS

LUDB and RADIUS access can be combined during subscriber-host instantiation phase.

Configuration wise, LUDB must be referenced under the **group-interface>dhcp/ppp/pppoe** hierarchy (and possibly under the capture SAP), while the authentication-policy is specified within the LUDB. In this fashion, LUDB access is followed by RADIUS access. The subscriber-host parameters retrieved from both sources are combined with LUDB parameters being prioritized over RADIUS parameters in case that both sources return the same parameters.

In case that LUDB and authentication policy are configured simultaneously under the group-interface (and possibly under the capture SAP), the RADIUS authentication policy will be evaluated and LUDB will be ignored.

## **RADIUS Fallback**

In case that RADIUS server is not accessible (non-responsive), the host instantiation phase can be:

Terminated in the case the there is no fallback action within authentication policy specified.

Continued within LUDB if the fallback action within the authentication-policy references LUDB.

Continued without any response from RADIUS. Subscriber-host will be instantiated if defaults parameters are statically configured or the instantiation will fail in case that the defaults are not available.

The fallback action takes effect once the preconfigured RADIUS timeout period expires.

RADIUS fallback is currently not supported for DHCPv6 hosts.

## Subscriber Services

Subscriber services enable an operational model to activate and deactivate subscriber functions from RADIUS through an Access-Accept or CoA message. Using the flexible RADIUS Python script interface, the operator defines the subscriber service functionality by populating a data structure using a parameter list received in a RADIUS Vendor Specific Attribute (VSA). The format and content of the parameter list VSA is defined by the operator. Each subscriber service instance can have a dedicated RADIUS accounting session; an accounting start/stop is sent when the subscriber service is activated/deactivated. Optionally, interim updates are sent with an interim update interval that can be specified per subscriber service instance. Accounting interim update and stop messages contain the subscriber service related statistics (time or volume-and-time).

Subscriber services can be activated on a dual-stack PPPoE session or a single stack IPv4 host. Subscriber service functionality is supported for subscriber QoS overrides: changing queues or policer parameters like rate or burst sizes and adapting root arbiter or subscriber aggregate rates. For example, an operator defines a service to boost the downstream rate using the parameters ("rate-limit":downstream-rate-in-mbps). A subscriber service is activated through a subscriber service activate VSA with value "rate-limit:20" that is received for a PPPoE session. This triggers the operator-defined RADIUS Python script to populate the subscriber-service data-structure variable that changes the subscriber aggregate downstream rate to 20 Mbps. Optionally an accounting start is sent for the subscriber service. Later, when a subscriber service deactivate VSA with the same value "rate-limit:20" is received for the same PPPoE session, the original subscriber aggregate downstream rate is restored, and optionally an accounting stop sent.

## Flexible Subscriber-Interface Addressing (Unnumbered Subscriber-Interfaces)

### Terminology

Subscriber host — A 7x50 representation of an external host requesting a service. Each such host is fully instantiated within the 7x50 for the purpose of providing traffic control and billing services (for example, QoS, filtering, antispooftng, accountitng). The external hosts may represent variety of devices such as regular PCs, STBs, residential gateways, CPEs, VoIP devices. In most cases, the external host will run a DHCPv4/v6 or PPPoEv4/v6 client. DHCP and PPPoE initiation messages from such clients will trigger host instantiation within 7x50. For this the subscriber host term can be interchangeably used with a term DHCP client or PPPoE client.

### Flexible Subscriber-Interface Addressing for IPOE/PPPOE v4/v6 Subscribers

In certain wholesale/retail environments, the wholesale provider that own the 7x50-BNG does not know the IP addresses that the retailers will assign to their clients in advance. For this reason, wholesaler's 7x50-BNG must accept any IP address from retailers and consequently pass it to the client during subscriber-host initiation phase.




Figure 121: Use Case for Flexible IP Addressing Model

Flexible addressing of the subscriber-interface assumes two deployment scenarios:

1. Subscriber-interface is unnumbered — For example, there is no explicit assigned IP address. Instead the subscriber-interface borrows the IP address from an existing interface that is operationally UP and is located in the same routing instance (router | vpn)<sup>2</sup>.

In this case any IP address can be assigned to the subscriber host under the unnumbered subscriber-interface. The subscriber IPv4 address will be installed in the FIB as /32 route while IPv6 address will be installed as an entry of the length anywhere between 64 and 128 bits.

2. Subscriber-interface is numbered — The IP address/prefix is explicitly configured and solely owned by the subscriber-interface.

In this case, all subscriber IP addresses/prefixes that fall under the subnet/prefix dictated by the configured subscriber-interface IP address/prefix will be directly aggregated under the subscriber-interface subnet. As such they will occupy a single entry in the FIB. The rest of the subscriber hosts with IP addresses/prefixes that fall outside of the configured range will be installed in the FIB as individual entries (/32 for IPv4 and an entry of the length anywhere between 64 and 128 bits for IPv6 hosts).

## Default Gateway in IPv4 Flexible Addressing

In scenarios where subscriber host IPv4 address lies within the configured subscriber-interface subnet, the default-gw IPv4 address for the host will be one of the subscriber-interface IPv4 addresses. In this case, the service provider is aware of the IPv4 addressing scheme in the 7x50-BNG and as such it will supply the DHCP client with the appropriate default-gw IPv4 address via LUDB, RADIUS or DHCP Server (in that order of priority).

In scenarios where the retail service provider wants to maintain independence from the IPv4 addressing scheme deployed in the 7x50-BNG (that is controlled by wholesaler), the retailer can always supply its own IPv4 address, the subnet mask and the default-gw IPv4 address. But if the default-gw IPv4 address and/or subnet mask is not supplied by the retailer, then they will be auto-generated by 7x50-BNG. Once the default-gw IPv4 address is auto-generated, it will be sent to the requesting DHCP client via DHCP offer in option 3 (RFC 2132, Router Option, section 3.5). There is no additional configuration needed for this action. 7x50-BNG will automatically detect whether the default-gw IPv4 address is supplied via LUDB, RADIUS or DHCP server and it will act correspondingly.

The default-gw IPv4 address is auto-generated based on the assigned IPv4 address/mask by setting the last bit of the assigned host IPv4 address to binary 01 or binary 10. For example if the subscriber host's assigned IPv4 address is 10.10.10.10 255.255.255.0, then the default-gw IPv4

2. Note that an interface must have an IP address assigned in order to be operationally UP. Therefore, an unnumbered subscriber-interface must reference another existing interface that is operationally UP in the same routing instance. The subscriber-interface will borrow the IP address from the referenced interface.

## Default Gateway in IPv4 Flexible Addressing

address is set to 10.10.10.1. If the assigned IPv4 address is 10.10.10.1 255.255.255.0 , then the auto-generated default gateway IPv4 is set to 10.10.10.2.

The default gateway IPv4 address will always have to be within the subscriber's subnet. If it is not, the behavior might be inconsistent. For example:

1. RADIUS (or DHCP) returns IP@, mask and def-gw:

- IP 10.10.10.1
- Subnet mask 255.255.255.0
- Def-gw 10.10.0.254

The subscriber will be successfully instantiated in 7x50-BNG but the client may not ARP for a default-gw outside of its configured subnet. Whether the client will or will not ARP for a default-gw outside of its configured subnet will depend on the implementation in the RG and CPE.

2. RADIUS returns IP@ and subnet mask.

- In this case the auto-generated default-gw IPv4 address will always be within subscriber's subnet.

Flexible IPv4 addressing with auto-generated default-gw is supported only in Routed Central Office (RCO) model with routed residential gateways (RGs) or CPEs. In RCO model with bridged residential gateways or CPEs, the default-gw IPv4 addresses and the assigned IPv4 addresses may overlap. Once the IPv4 address of the default-gw is auto-generated, it is possible that the second host behind the bridged residential gateway or CPE is assigned the same IPv4 address as the IPv4 address of the default gateway of the first host. Such hosts would not be able to communicate with outside world.

For example:

RADIUS or DHCP server assigns IPv4 address and subnet mask to the first host in a bridged environment:

IP1: 10.10.10.1

Auto-generated default-gw IPv4 address: 10.10.10.2

Since the RADIUS and DHCP Server are not aware of the auto-generated default-gw, they may assign the following IPv4 address to the second host that comes on-line:

IP 2: 10.10.10.2 (same IPv4 address as the default-gw IPv4 address of the first host)

Auto-generated default-gw IPv4 address: 10.10.10.1

Now the first host will forward all traffic outside of the configured subnet to the second hosts which will discard this traffic, effectively rendering this operation model non-deployable. And vice versa.

## IPv4 Subnet Sharing

Subnet sharing between the hosts in flexible IPv4 addressing model is supported. In other words, in flexible IPv4 addressing model the operator can assign all IPv4 addresses (minus one – the default-gw IPv4 address) from a given subnet. In this fashion, all subscribers (routed RGs or CPEs) within a single subnet can share the same default gateway.

For example if the operator owns the IPv4 subnet 10.10.10.0/24, then one IPv4 address can be set aside for the default-gw (for example 10.10.10.254) and the remaining addresses can be assigned to the subscriber (routed RGs or CPEs). An example would be:

RG1: IP=10.10.10.1/24 def-gw 10.10.10.254

RG2: IP=10.10.10.2/24 def-gw 10.10.10.254

RG3: IP=10.10.10.3/24 def-gw 10.10.10.254

:

RG100: IP=10.10.10.100/24 def-gw 10.10.10.254

The subnet sharing is also supported in conjunction with auto-generated default-gw IPv4 address. The implication of this is that the IPv4 address of the default-gw can collide with the same IPv4 address already assigned to an existing subscriber. This is not an issue for routed RGs or CPEs since 7x50-BNG will always answers ARPs for the IPv4 address of the default-gw with its own (7x50) MAC address. However, local-proxy ARP functionality in 7x50-BNG MUST be enabled to support this.

This behavior can be further clarified with the following example.

Let's assume that we have scenario with two routed RGs:

RG-1, IP=10.10.10.0/24, default-gw IP=10.10.10.1

RG-2, IP=10.10.10.1/24, default-gw IP=10.10.10.0

Once RG-1 ARPs for its default gateway of 10.10.10.1, 7x50-BNG will reply with its own MAC address.

Now that host RG-1 has resolved ARP for its default-gw (mac address pointing to 7x50), it can send traffic to the outside world via 7x50-BNG. When such traffic arrives to 7x50, the destination IPv4 address of the received packet will determine the forwarding decision within 7x50. If the destination IPv4 address matches the IPv4 address of any subscriber (RG) instantiated within 7x50, the traffic will be forwarded to that RG. This also includes the case where the destination IPv4 address is the default-gw IPv4 address (10.10.10.1), which represents just another RG within 7x50. The traffic will be consequently passed from RG-1 via 7x50 to RG-2.

## IPv4 Subnet Mask Auto-Generation

The subnet mask corresponding to the IPv4 address assigned to the subscriber is auto-generated in case that the IPv4 addressing authority (LUDB, RADIUS or DHCP Server) does not supply it. The subnet mask is derived from the IPv4 address of the subscriber and possibly the default-gw IPv4 address and it is the smallest subnet that contains both, the IPv4 address of the subscriber and the default-gw.

For example if the RADIUS received IPv4 address is 10.10.10.138 and the received default-gw IPv4 address is 10.10.10.170 , then the subnet mask will be auto-generated and set to 255.255.255.192 ( /26).

138 = 10001010

170 = 10101010

192 = 11000000

In case that neither the subnet mask nor the default-gw are returned, then both would be auto-generated:

1. Subnet mask would be set to /31
2. Default-gateway which must belong to the subscriber's subnet would be set to 10.10.10.139.

In cases where the host IPv4 address and the default-gw are directly supplied by the addressing authority but the subnet mask is missing, the subnet mask auto-generation may cause the host part of the default-gw IPv4 address to become a broadcast IPv4 address. If this is an issue, then it can be avoided by directly providing the subnet mask via the addressing authority.



## Local-proxy-arp and arp-populate

Local-proxy-arp and arp-populate are two commands that are relevant only to IPoEv4 hosts.

Local-proxy-arp command ensures that 7x50 answers ARP Requests with its own MAC address for any 'active' IPv4 address under the subnet on which the ARP request arrived. The 'active' IPv4 address is considered the one that is assigned to an already instantiated hosts or the default-gw (even auto-generated).

In absence of local-proxy-arp command, the only ARP Request that 7x50 will answer is the one for the statically configured IPv4 addresses of the subscriber-interface. In flexible IPv4 addressing, the IPv4 address of the default-gw does not necessarily match any of the configured subscriber-interface IPv4 addresses. The ARP Request for such default-gw IPv4 address would go unanswered. Consequently, the subscriber hosts would not be able to communicate with outside world. Therefore, the flexible IPv4 addressing requires that the local-proxy-arp command is configured.

Arp-populate command disables dynamic learning of ARP entries (IPv4<->MAC mapping) on an interface based on the ARP protocol. In this case, the ARP table is populated based on the DHCPv4 lease state table which contains IPv4<->MAC mappings obtained via DHCP processing during the host instantiation phase. Arp-populate functionality is highly desirable in case of flexible IPv4 addressing.

When arp-populate command is disabled the ARP entries are dynamically learned based on the ARP protocol. This, in conjunction with flexible IPv4 addressing may cause certain issues. Consider the following example:

- The subscriber-host is instantiated in 7x50
- The subscriber-interface is unnumbered
- The ARP table does not contain an ARP entry for the subscriber-host

In this case, downstream traffic towards the subscriber host will trigger 7x50 to send ARP Request for the subscriber host IPv4 address. 7x50 needs to know the MAC address of the subscriber-host in order to forward traffic. Since the subscriber-interface is unnumbered, the source IPv4 address of the ARP request is unknown and consequently the ARP Request will not be sent. As a result, downstream traffic will be dropped.

Note however that the above example is an unlikely scenario. If the subscriber host sends the ARP Request for the default-gw first, 7x50 would create an entry in the ARP table for it and the issue would be resolved. This is the most likely outcome since the subscriber host will always try to initiate communication with the outside and therefor ARP for the IPv4 address of the default-gw (which is 7x50).

## Gi-address Configuration Consideration

With flexible IPv4 address assignment, the gi-address can be configured as any IPv4 address that is already assigned to an interface (loopback interface, regular interface attached to physical port or subscriber interface) within the same routing instance (VRF or GRT).

---

## PPPoE Considerations

PPPoE subscriber hosts do not have the concept of default-gw. Consequently the default-gw auto-generation concept does not apply to PPPoE hosts.

---

## IPoEv6 Considerations

The default-gw for IPoEv6 hosts is link-local IPv6 address. Since this address is always present, there is no need for auto-generation during the subscriber instantiation time.

SLAAC hosts are installed as /64 entries, the length of the installed DHCP-PD prefix is dictated by the prefix-length and the DHCP-NA hosts are installed as /128 entries.

---

## General Configuration Guidelines for Flexible IP Address Assignment

Flexible IP addressing for IPoE/PPPoE v4 and v6 hosts is by default disabled. In other words, the subscriber hosts will be instantiated in 7x50-BNG with ability to forward traffic only if their assigned IP addresses belong to one of the configured subnets/prefixes that are associated with the subscriber-interfaces. IPv4 and IPv6 cases will be examined separately:

### IPv4:

By default, IPoE and PPPo subscriber host creation will fail in the following two cases:

1. The subscriber-interface does not have an IPv4 address configured, and therefore it will be operationally down. This configuration is also known as unnumbered subscriber-interface.
2. The subscriber-interface does have an IPv4 address configured but the IPv4 address assigned to the subscriber host itself is outside of the subscriber-interface configured subnet(s). In such case the host will be instantiated but the forwarding will be disabled.

Subscriber host instantiation and forwarding can be explicitly enabled for both cases above with flexible IP addressing functionality.

For case 1, this can be achieved by borrowing an IP address for the subscriber-interface from any interface that is operationally up within the given routing context. This functionality can be enabled with the following command:

```
configure service ies <id>
configure service vprn <id>
subscriber-interface <intf-name>
    unnumbered <ip-addr | interface-name >
```

To enable forwarding for the subscribers whose IP address falls outside of the configured subnet under the subscriber-interface (case 2), the following command must be entered:

```
configure service ies <id>
configure service vprn <id>
subscriber-interface <intf-name>
    allow-unmatching-subnets
```

The above commands (**unnumbered** and **allow-unmatching-subnets**) are mutually exclusive. In addition, the unnumbered command can be configured only if the subscriber-interface does not have an IP address already configured. Otherwise the execution of this command will fail.

In both of these cases the host will be installed in the routing table as /32.

#### IPv6:

For IPv6 there is a single command that will enable flexible IP addressing for both cases:

1. IPv6 prefixes are not configured under the **subscriber-interface>ipv6** node
2. IPv6 prefixes are configured but the actual address or prefix assigned to the subscriber (via DHCP, LUDB or RADIUS) is outside any prefix that is configured under the **subscriber-interface>ipv6** hierarchy.

This single command is:

```
configure service ies <id>
configure service vprn <id>
    subscriber-interface <name>
        ipv6
            allow-unmatching-prefixes
```

To summarize, the following scenarios are possible:

- PPPoEv4
  - An IPv4 address under the subscriber-interface is configured
    - By default hosts outside of the sub-intf subnet are instantiated but they are in a non-forwarding-state. Traffic is dropped.
    - **allow-unmatching-subnets** is configured. This command is allowed only if subscriber-interface has also configured its own IPv4 address(es). In this case the IP address for IPCP negotiation is one of the sub-intf addresses. Hosts outside of the sub-intf subnets are instantiated and forwarded.

## Caveats

- The **unnumbered** *<ip-address | intf>* command is not allowed in this scenario.
  - An IPv4 address under the subscriber-interface is not configured
    - By default, the subscriber-interface is operationally down. Subscribers cannot be instantiated.
    - The **allow-unmatching-subnets** command has no effect since subscriber-interface does not have an IPv4 address configured and is therefore operationally down. No subscribers can be instantiated.
    - The **unnumbered** *<ip-address | intf>* command is the only viable option in this case. The subscriber-interface borrows an IPv4 address from another interface that is operationally UP and consequently this allows subscribers to be instantiated. This command is mutually exclusive with **allow-unmatching-subnets**. In addition, this command can only be configured if the subscriber-interface itself does not have explicitly configured an IPv4 address.
  - IPoEv4

Similar to the PPPoE case above.
  - IPoEv6 and PPPoEv6 — the **allow-unmatching-prefixes** command is independent of any IPv4 command related to flexible IP address assignment (**unnumbered** or **allow-unmatching-subnets**). This command can always be enabled, regardless of the v6 prefixes configured under the **subscriber-interface>ipv6** hierarchy. Any subscriber, regardless of the subscriber-interface prefix configuration will be instantiated and forwarded.
- 

## Caveats

- Auto-generation of the default-gw IPv4 address is supported only in RCO model with routed RGs/CPEs. Bridged RGs/CPEs are not supported.
- A configured IPv4 address cannot be removed from the subscriber-interface when DHCPv4 hosts under the corresponding subnet are instantiated in the system.
- An IPv4 address cannot be configured under the subscriber-interface while (unnumbered) DHCPv4 hosts under that subnet are already instantiated.
- Executing the **no allow-unmatching-subnets** command is only allowed when there are no unnumbered DHCPv4 hosts instantiated under the subscriber-interface.

## uRPF for Subscriber Management

uRPF is supported for IPv4 and IPv6 dual-stack subscribers with framed routes.

For IPv4, uRPF is supported on group interfaces using anti-spoofing filters. A group interface configured for NATed subscribers will be configured with MAC/IP/PPPoE Session-ID anti-spoofing filters.

IPv6 subscribers, which are non-NAT, are always treated as being on a local subnet. For such subscribers, a 7x50 BNG will install a FIB entry for local routes that match either the wan-host prefix, or the delegated prefix, or both. In strict mode for IPv6 ESM, the uRPF check will check not just that the route matching the SA (which should be a local route, i.e. a subnet) would route the packet back out of the interface it came in on, but in addition that we would route the packet out to the same SAP it was received on.

SROS supports the ability to configure a NH-MAC anti-spoof type for non-NATed subscribers. When configured, the datapath performs ingress anti-spoofing based on source MAC address and egress anti-spoof (also referred to as egress subscriber-host look-up) based on the nh-ip address.

The NH-MAC anti-spoof type is configured under the following context:

```
config>service>vprn>if>sap
```

```
config>service>ies>sub-if>grp-if>sap
```

```
config>service>vprn>sub-if>grp-if>sap
```

```
config>subscr-mgmt>msap-policy
```


A uRPF check is also performed that prefixes delegated to a subscriber on that MAC address exist in the FIB.

## IPoE Sessions

The IP stacks of dual stack IPoE devices are set up and configured independently using different protocols such as DHCPv4, DHCPv6 or SLAAC. As opposed to PPPoE, there is no single protocol that binds the IP stacks from a single end device together.

To facilitate subscriber management of dual stack IPoE devices as a single entity similar as for PPPoE sessions instead of handling individual IPoE subscriber hosts, there is a need for a logical IPoE session construct. An IPoE session enables single authentication, session accounting and policy management (mid-session changes) for dual stack IPoE devices.

An IPoE session is a logical grouping of IPoEv4 and IPoEv6 subscriber hosts that represent the different IP stacks of a single end device and that share authentication data such as subscriber id, subscriber and SLA profile, session-timeout, etc. The grouping of subscriber hosts in an IPoE session is based on a configurable session key per group-interface. The IPoE session key includes by default the SAP identifier and MAC address and can be extended with Circuit-Id/Interface-Id or Remote-Id. For DHCPv6 Remote-Id, the enterprise number is excluded from the session-key. Circuit-id/Interface-Id or Remote-id should only be used in the IPoE session key if all subscriber host associated with the IPoE session have this field in their protocol trigger packets. The IPoE session creation or subscriber host association to an IPoE session fails if the Circuit-Id/Interface-Id or Remote-id is not present in a trigger packet while the field is part of the session-key.



**Figure 122: IPoE Session**

An IPoE session represents a single end device and can have following associated IP stacks:

- IPv4 — A single DHCPv4 host.
- IPv6 WAN — One DHCPv6 IA-NA host and/or one SLAAC host.
- IPv6 PD — One DHCPv6 IA-PD host or PD as managed route.

A violation of the above rules will result in a setup failure of the subscriber host when an attempt is made to associate it to the IPoE session.

## Enabling IPoE Sessions

IPoE sessions are supported in a Routed CO environment with Enhanced Subscriber Management (ESM) enabled. To enable the IPoE session instantiation, the **ipoe-session** CLI context on the capture SAP (managed SAP scenario) and/or group-interface must be configured to no shutdown. See also the configuration steps below.

IPoE sessions cannot be enabled on a group-interface with active subscriber hosts.

Disabling IPoE sessions by executing an **ipoe-session no shutdown** or **no ipoe-session** CLI command on a group-interface will delete all active sessions and associated hosts on that group-interface, resulting in service impact for these subscribers.

---

## IPoE Session Authentication

A single authentication is performed for all subscriber hosts that belong to the same IPoE session. The table below lists the packets that trigger an IPoE session authentication.

**Table 18: IPoE Session Authentication Trigger Packets**

IP stack	Trigger packets
IPv4	DHCPv4 Discover
	DHCPv4 Request
IPv6 WAN	DHCPv6 Solicit
	DHCPv6 Request
	DHCPv6 Relay Forward (Solicit)
	DHCPv6 Relay Request (Solicit)
	Router Solicitation
IPv6 PD	DHCPv6 Solicit
	DHCPv6 Request
	DHCPv6 Relay Forward (Solicit)
	DHCPv6 Relay Request (Solicit)

When a trigger packet is received on a capture SAP or group-interface with IPoE sessions enabled, an IPoE session lookup is performed based on the configured IPoE session key:

- If no IPoE session is found, a new session is created and authenticated following the ESM authentication configuration such as local user database lookup, Radius or Diameter authentication, defaults, and such. After successful authentication, the authentication data is stored in the IPoE session state. The subscriber host is created and associated with the session.
- If an IPoE session already exists, and no re-authentication must be performed then the subscriber host is created using the stored IPoE session data. The subscriber host is associated with the session.
- If an IPoE session already exists, and re-authentication must be performed then the session is re-authenticated. When successful, the authentication data for the IPoE session is updated and applied to all associated hosts. The subscriber host is created and associated with the session. When unsuccessful, existing hosts associated with the session are not impacted and the session data is kept unchanged.

Re-authentication is by default disabled for IPoE sessions. To enable re-authentication, a minimum authentication interval must be configured. The `min-auth-interval` CLI parameter configures the maximum frequency of re-authentications by specifying a minimum interval between two non-forced authentications for the same IPoE session. A re-authentication is triggered by the renewal of any host belonging to the IPoE session. Setting the `min-auth-interval` to zero seconds, will always re-authenticate on each trigger packet. The re-authentication CLI knob in a Radius authentication policy is ignored for IPoE session authentication.

A forced authentication is performed when the `Circuit-Id/Interface-Id` or `Remote-Id` in the trigger packet has changed. An empty or absent `Circuit-Id/Interface-Id` or `Remote-Id` is not considered as a change. The default forced authentication behavior is changed with the `force-auth` CLI command in the group-interface `ipoe-session` context: only force authenticate on `Circuit-Id/Interface-Id` change or only force authenticate on `Remote-Id` change or disable forced authentications.

A new local user database config in the `ipoe-session` CLI context on a capture SAP or group interface ensures that all subscriber hosts associated with an IPoE session are using the same database and therefore common match criteria. The per subscriber host type user-db configs, such as `ipv6 dhcp6 user-db`, `dhcp user-db` and `rtr-solicit-user-db` are ignored when IPoE sessions are enabled.

---

## IPoE Session Accounting

All Radius accounting modes can be enabled for IPoE sessions: queue instance accounting, host accounting or session accounting.

With session accounting, a Radius accounting start is generated when the first host of the session is created and an accounting stop when the last host of the session is deleted. The generation and



interval of periodic interim updates can be configured. Optionally, triggered interim update messages can be generated when a host is deleted from the session or an additional host is associated.

A unique accounting session id is generated for the IPoE session and is used in RADIUS session accounting. The IPoE session accounting session id can be included in the Radius Access Request message via CLI **config>subscr-mgmt>auth-plcy# include-radius-attribute acct-session-id session.**

This accounting session ID can also be used in RADIUS CoA or Disconnect Messages to target the IPoE session.

## IPoE Session Mid-Session Changes

Mid-session changes such as those initiated via RADIUS CoA or Diameter Gx RAR are applied to all hosts associated with the 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 host of the session are targeted and the session state is updated with the new data.

The following tools commands are available to manually enforce a mid-session change:

```
# tools perform subscriber-mgmt edit-ipoe-session sap <sap-id> mac <mac-address> [sub-
scriber <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>]
```

## IPoE Session Termination

When the last subscriber host associated with an IPoE session is deleted from the system, then the IPoE session is also deleted.

An IPoE session and all associated subscriber hosts can be deleted via

- CLI clear command: **clear service id <service-id> ipoe session**
- An **ipoe-session no shutdown** CLI command on a group-interface
- A **no ipoe-session** CLI command on a group-interface. This command resets to the default behavior, which is IPoE sessions disabled.
- Session timeout, configured in the IPoE session policy or obtained from AAA
- Idle timeout

- RADIUS Disconnect Message
- Diameter Gx session termination
- Credit Control: Radius or Diameter Gy

## Limiting the Number of IPoE sessions

The number of IPoE sessions per SAP is limited with the **sap-session-limit** CLI command configured in the **group interface ipoe-session** context.

The number of IPoE sessions per group interface or retail subscriber interface is limited with the **session-limit** CLI command configured in the group interface ipoe-session or retail subscriber interface ipoe-session context.

Hosts associated with IPoE sessions are subject to the SLA Profile Instance limits configured in the **sla profile >config>subscr-mgmt>sla-prof>host-limits>** CLI context.

---

## SAP Session Index

The system will keep track of the number of IPoE sessions active on a given SAP and assign a per SAP session index to each such that always the lowest free index is assigned to the next active IPoE session. When RADIUS authentication is used, the SAP session index can be sent to, and received from, the RADIUS server using the [26-6527-180] Alc-SAP-Session-Index attribute.

It should only be used in a subscriber per VLAN model as the session index is per SAP.

The intended use of the SAP session index is to provide the ability for IPoE sessions in a bridged RG environment to have their own set of queues (for QoS and accounting purposes) when using the same SLA profile name received from a RADIUS server. See section [Subscriber per PPPoE Session Index on page 568](#) for details.

---

## Resiliency

For non-redundant BNG deployments, the IPoE session state is stored in the subscriber-mgmt persistency file for recovery from Compact Flash after a node maintenance operation or failure. This is configured at the system persistence CLI context.

For multi chassis redundancy scenarios, the IPoE session state is synchronized via the “sub-mgmt ipoe” Multi Chassis Synchronization (MCS) application.

Notes:

- Static hosts can be configured on a group-interface with IPoE sessions enabled. A static host will not be associated with an IPoE session.
- Up to sixteen Framed-Routes and sixteen Framed-IPv6-Routes can be associated with an IPoE session.

- A fall back action (accept or local user database lookup) when no Radius servers are available for Radius authentication can be specified for IPoE sessions.
- Lawful Intercept sources initiated from Radius always include all IP stacks from the IPoE session regardless the targeted host in the CoA message.
- ARP hosts are not supported in an IPoE session and cannot be instantiated on a group-interface with IPoE sessions enabled.
- The creation of an IPv4 host using the Alc-Create-Host attribute in a Radius CoA message is not supported on a group-interface with IPoE session enabled.
- A local user database host identification based on option60 is ignored when authenticating an IPoE session.
- Radius authentication of an IPoE session fails when the user-name-format is configured to dhcp-client-vendor-options, mac-giaddr or ppp-user-name
- The alc.dtc.setESM() API in the DHCP Transaction Cache (DTC) Python module cannot be used in combination with IPoE sessions.
- The DHCP Python module (alc.dhcp) used to derive subscriber host attributes from a DHCPv4 ACK message is not supported in combination with IPoE sessions.
- A Radius CoA message containing an Alc-Force-Nak or Alc-Force-Renew attribute is not supported for IPoE sessions
- Subscriber Host Connectivity Verification (SHCV) will continue to work on a per-stack basis. In other words, in a dual stack scenario with SHCV action remove enabled for both stacks, a failure in IPv4 connectivity will not clean up the session unless the IPv4 subscriber host was the last associated host.

## Configuration steps

To create an IPoE session policy:

```
configure
  subscriber-mgmt
    ipoe-session-policy "ipoe-policy-1" create
      description "Default IPoE session policy"
      session-key sap mac          # default
      no session-timeout          # default
  exit
```

Enable IPoE sessions on the capture SAP and/or group interface.

If IPoE sessions is enabled on a capture-sap, then it must also be enabled on the target group-interface. If an IPoE session local user database lookup is configured at the capture-sap, then the same local user database lookup must be configured at the target group-interface.

```
configure
  service

    vpls 10 customer 1 create
    ---snip---
    sap 1/1/4:*. * capture-sap create
    ---snip---
    ipoe-session
      description "IPoE sessions - capture-sap"
      ipoe-session-policy "ipoe-policy-1"
      user-db "ludb-1"
      no shutdown
    exit

    ies 1000 customer 1 create
    subscriber-interface "sub-int-1" create
    ---snip---
    group-interface "group-int-1-1" create
    ---snip---
    ipoe-session
      description "IPoE sessions - IES group-interface"
      force-auth cid-change rid-change          # default
      ipoe-session-policy "ipoe-policy-1"
      min-auth-interval infinite                # default
      sap-session-limit 1                       # default
      session-limit 1000
      user-db "ludb-1"
      no shutdown
    exit
```

To display the IPoE session state, use following command:

```
# show service id <service-id> ipoe session [detail]
```



Configuration steps

---

# Subscriber Management Command Reference

---

## Configuration Commands

- [ANCP Commands on page 1085](#)
- [RADIUS Route Download Commands on page 1092](#)
- [RADIUS Accounting Policy Commands on page 1088](#)
- [RADIUS Route Download Commands on page 1092](#)
- [Authentication Policy Commands on page 1086](#)
- [Category Map and Credit Control Policy Commands on page 1096](#)
- [Diameter Policy Commands on page 1093](#)
- [Filter Commands on page 1099](#)
- [Explicit Subscriber Mapping Commands on page 1101](#)
- [Host Tracking Policy Commands on page 1102](#)
- [PIM Policy Commands on page 1103](#)
- [Managed SAP Policy Commands on page 1117](#)
- [Multi-Chassis Redundancy Commands on page 1113](#)
- [SLA Profile Commands on page 1104](#)
- [Subscriber Identification Policy Commands on page 1107](#)
  - [Auto-Generated Subscriber Identification Key Commands on page 1108](#)
  - [Auto-Generated Subscriber Identification Key Service Commands on page 1108](#)
- [Subscriber MCAC Policy Commands on page 1109](#)
- [Subscriber Profile Commands on page 1110](#)
- [Subscriber Management Service Commands on page 1116](#)
  - [VPLS Subscriber Management Configuration Commands on page 1116](#)
- [RIP Commands on page 1121](#)
- [VPort Commands on page 1122](#)
- [Wireless Portal Protocol \(WPP\) Commands on page 1124](#)
- [Multicast Listener Discovery \(MLD\) Commands on page 1127](#)
- [Show Commands on page 1128](#)
- [Clear Commands on page 1132](#)
- [Debug Commands on page 1133](#)

Note: Enhanced Subscriber Management is supported on the redundant chassis model only.

Subscriber management commands are also described in the [Triple Play Services Command Reference on page 65](#) section.



## ANCP Commands

```

config
— subscriber-mgmt
  — ancp
    — ancp-policy name
      — egress
        — rate-adjustment adjusted-percent
        — no rate-adjustment
        — rate-modify agg-rate-limit
        — rate-modify {scheduler scheduler-name | arbiter arbiter-name}
        — no rate-modify
        — rate-monitor kilobit-per-second [alarm]
        — no rate-monitor
        — rate-reduction kilobit-per-second
        — no rate-reduction
      — ingress
        — rate-adjustment adjusted-percent
        — rate-modify {scheduler scheduler-name | arbiter arbiter-name}
        — rate-monitor kilobit-per-second [alarm] [log]
        — rate-reduction kilobit-per-second
      — [no] port-down
        — [no] disable-shcv [alarm] [hold-time seconds]
    — ancp-static-map
      — entry key ancp-string customer customer-id multi-service-site customer-site-name ancp-policy policy-name
      — entry key ancp-string sap sap-id ancp-policy policy-name
      — no entry key ancp-string customer customer-id multi-service-site customer-site-name
      — no entry key ancp-string sap sap-id

```

## Authentication Policy Commands

```

config
  — subscriber-mgmt
    — aaa
      — radius-coa-port {1647|1700|1812|3799}
      — no radius-coa-port
    — [no] authentication-policy name
      — [no] accept-authorization-change
      — accept-script-policy policy-name
      — no accept-script-policy
      — description description-string
      — no description
      — fallback-action accept
      — fallback-action user-db-local-user-name
      — no fallback-action
      — [no] include-radius-attribute
        — [no] access-loop-options
        — [no] acct-session-id
        — calling-station-id
        — calling-station-id {mac | remote-id | sap-id | sap-string}
        — no calling-station-id
        — [no] circuit-id
        — [no] dhcp-options
        — [no] dhcp6-options
        — [no] dhcp-vendor-class-id
        — [no] mac-address
        — [no] nas-identifier
        — nas-port-id [prefix-string string] [suffix suffix-option]
        — nas-port-type
        — nas-port-type [0..255]
        — no nas-port-type
        — [no] pppoe-service-name
        — [no] remote-id
        — [no] sap-session-index
        — [no] tunnel-server-attrs
      — password password
      — no password
      — ppp-user-name append domain-name
      — ppp-user-name default-domain domain-name
      — ppp-user-name replace domain-name
      — ppp-user-name strip
      — no ppp-user-name
      — pppoe-access-method {none | padi | pap-chap}
      — no pppoe-access-method
      — radius-authentication-server
        — access-algorithm {direct | round-robin}
        — no access-algorithm
        — [no] access-loop-options
        — fallback-action accept
        — fallback-action user-db local-user-db-name
        — no fallback-action
        — hold-down-time seconds
        — no hold-down-time
        — retry count

```

```

— no retry
— router router-instance
— router service-name
— no router
— server server-index address ip-address secret key [hash | hash2] [port
port] [coa-only]
— no server server-index
— source-address ip-address
— no source-address
— timeout seconds
— no timeout
— radius-server-policy radius-server-policy-name
— [no] radius-server-policy
    — server server-index address ip-address secret key [hash | hash2] [port
port-num] [coa-only] [pending-requests-limit limit]
        — hold-down-time
— [no] re-authentication
— request-script-policy policy-name
— no request-script-policy
— send-acct-stop-on-fail {[on-request-failure] [on-reject] [on-accept-failure]}
— no send-acct-stop-on-fail
— 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]
— no user-name-format

```

## RADIUS Accounting Policy Commands

```

config
— subscriber-mgmt
    — [no] radius-accounting-policy name
        — [no] custom-record
            — [no] override-counter override-counter-id
                — e-counters [all]
                — no e-counters
                    — [no] in-profile-octets-discarded-count
                    — [no] in-profile-octets-forwarded-count
                    — [no] in-profile-packets-discarded-count
                    — [no] in-profile-packets-forwarded-count
                    — [no] out-profile-octets-discarded-count
                    — [no] out-profile-octets-forwarded-count
                    — [no] out-profile-packets-discarded-count
                    — [no] out-profile-packets-forwarded-count
                — i-counters [all]
                — no i-counters
                    — [no] all-octets-offered-count
                    — [no] all-packets-offered-count
                    — [no] high-octets-discarded-count
                    — [no] high-packets-discarded-count
                    — [no] in-profile-octets-forwarded-count
                    — [no] in-profile-packets-forwarded-count
                    — [no] low-octets-discarded-count
                    — [no] low-packets-discarded-count
                    — [no] out-profile-octets-forwarded-count
                    — [no] out-profile-packets-forwarded-count
            — [no] queue queue-id
                — e-counters [all]
                — no e-counters
                    — [no] in-profile-octets-discarded-count
                    — [no] in-profile-octets-forwarded-count
                    — [no] in-profile-packets-discarded-count
                    — [no] in-profile-packets-forwarded-count
                    — [no] out-profile-octets-discarded-count
                    — [no] out-profile-octets-forwarded-count
                    — [no] out-profile-packets-discarded-count
                    — [no] out-profile-packets-forwarded-count
                — i-counters [all]
                — no i-counters
                    — [no] all-octets-offered-count
                    — [no] all-packets-offered-count
                    — [no] high-octets-discarded-count
                    — [no] high-octets-offered-count
                    — [no] high-packets-discarded-count
                    — [no] high-packets-offered-count
                    — [no] in-profile-octets-forwarded-count
                    — [no] in-profile-packets-forwarded-count
                    — [no] low-octets-discarded-count
                    — [no] low-octets-offered-count
                    — [no] low-packets-discarded-count
                    — [no] low-packets-offered-count

```

- [no] out-profile-octets-forwarded-count
- [no] out-profile-packets-forwarded-count
- [no] uncoloured-octets-offered-count
- [no] uncoloured-octets-offered-count
- **ref-override-counter** *ref-override-counter-id*
- **ref-override-counter** all
- **no ref-override-counter**
  - **e-counters** [all]
  - **no e-counters**
    - [no] in-profile-octets-discarded-count
    - [no] in-profile-octets-forwarded-count
    - [no] in-profile-packets-discarded-count
    - [no] in-profile-packets-forwarded-count
    - [no] out-profile-octets-discarded-count
    - [no] out-profile-octets-forwarded-count
    - [no] out-profile-packets-discarded-count
    - [no] out-profile-packets-forwarded-count
  - **i-counters** [all]
  - **no i-counters**
    - [no] all-octets-offered-count
    - [no] all-packets-offered-count
    - [no] high-octets-discarded-count
    - [no] high-packets-discarded-count
    - [no] in-profile-octets-forwarded-count
    - [no] in-profile-packets-forwarded-count
    - [no] low-octets-discarded-count
    - [no] low-packets-discarded-count
    - [no] out-profile-octets-forwarded-count
    - [no] out-profile-packets-forwarded-count
- **ref-queue** *queue-id*
- **ref-queue** all
- **no ref-queue**
  - **e-counters** [all]
  - **no e-counters**
    - [no] in-profile-octets-discarded-count
    - [no] in-profile-octets-forwarded-count
    - [no] in-profile-packets-discarded-count
    - [no] in-profile-packets-forwarded-count
    - [no] out-profile-octets-discarded-count
    - [no] out-profile-octets-forwarded-count
    - [no] out-profile-packets-discarded-count
    - [no] out-profile-packets-forwarded-count
  - **i-counters** [all]
  - **no i-counters**
    - [no] all-octets-offered-count
    - [no] all-packets-offered-count
    - [no] high-octets-discarded-count
    - [no] high-octets-offered-count
    - [no] high-packets-discarded-count
    - [no] high-packets-offered-count
    - [no] in-profile-octets-forwarded-count
    - [no] in-profile-packets-forwarded-count
    - [no] low-octets-discarded-count
    - [no] low-packets-discarded-count
    - [no] low-octets-offered-count

```

— [no] low-packets-offered-count
— [no] out-profile-octets-forwarded-count
— [no] out-profile-packets-forwarded-count
— [no] uncoloured-octets-offered-count
— [no] uncoloured-packets-offered-count
— significant-change delta
— no significant-change
— description description-string
— no description
— fallback-action accept
— fallback-action user-db local-user-db-name
— no fallback-action accept
— [no] host-accounting [interim-update]
— include-radius-attribute
— [no] access-loop-options
— [no] calling-station-id
— [no] circuit-id
— [no] detailed-acct-attributes
— [no] dhcp-vendor-class-id
— [no] framed-ip-addr
— [no] framed-ip-netmask
— [no] nas-port
— [no] nas-port-id
— [no] remote-id
— [no] sla-profile
— [no] std-acct-attributes
— [no] sub-profile
— [no] subscriber-id
— [no] user-name
— [no] v6-aggregate-stats
— [no] wifi-rssi
— queue-instance-accounting [interim-update]
— no queue-instance-accounting
— radius-authentication-server
— access-algorithm {direct | round-robin}
— no access-algorithm
— retry count
— no retry
— router router-instance
— no router
— server server-index [address ip-address] [secret key] [port port] [pending-requests-limit limit]
— no server server-index
— source-address ip-address
— no source-address
— timeout seconds
— no timeout
— radius-server-policy radius-server-policy-name
— [no] radius-server-policy
— session-accounting [interim-update] [host-update]
— no session-accounting
— session-id-format {description | number}
— no session-id-format
— update-interval minutes
— no update-interval

```

- **update-interval-jitter** **absolute** *seconds*
- **[no]** **update-interval-jitter**

## RADIUS Route Download Commands

```
configure
— aaa
— route-downloader name [create]
— no route-downloader name
   — base-user-name user-name
   — no base-user-name
   — default-metric metric
   — no default-metric
   — default-tag tag
   — no default-tag
   — description description-string
   — no description
   — download-interval minutes
   — no download-interval
   — max-routes routes
   — no max-routes
   — password password [hash|hash2]
   — no password
   — radius-server-policy policy-name
   — no radius-server-policy
   — retry-interval min minimum max maximum
   — no retry-interval
   — [no] shutdown
```



## Diameter Policy Commands

- [AAA Diameter Peer Policy Commands on page 1093](#)
- [Subscriber Management Diameter Application Policy Commands on page 1094](#)

```

config
— aaa
— diameter-peer-policy peer-policy-name [role {client | proxy}] [create]
— no diameter-peer-policy peer-policy-name
— applications {[gx] [gy] [nasreq]}
— no applications
— connection-timer connection-time
— no connection-timer
— description description-string
— no description
— origin-host origin-host-string
— no origin-host
— origin-realm origin-realm-string
— no origin-realm
— peer name [create]
— no peer name
— address ip-address
— no address
— connection-timer connection-time
— no connection-timer
— destination-host destination-host-string
— no destination-host
— destination-realm destination-realm-string
— no destination-realm
— preference preference
— no preference
— [no] shutdown
— transaction-timer seconds
— no transaction-timer
— transport tcp port port
— no transport
— watchdog-timer seconds
— no watchdog-timer
— python-policy [32 chars max]
— no python-policy
— router service service-name
— router router-instance
— no router
— source-address ip-address
— no source-address
— transaction-timer transaction-time
— no transaction-timer
— vendor-support {three-gpp | vodafone}
— no vendor-support
— watchdog-timer wd-time
— no watchdog-timer

```

## Subscriber Management Diameter Application Policy Commands

```

configure
— subscriber-mgmt
— diameter-application-policy application-policy-name [create]
— no diameter-application-policy application-policy-name
— application {gx | gy | nasreq}
— no application
— description description-string
— no description
— diameter-peer-policy peer-policy-name
— no diameter-peer-policy
— gx
— avp-subscription-id origin [type type]
— no avp-subscription-id
— [no] include-avp
— [no] an-gw-address
— [no] called-station-id
— calling-station-id [type {llid | mac | remote-id | sap-id | sap-
string}]
— no calling-station-id
— [no] ip-can-type
— [no] logical-access-id
— nas-port binary-spec
— no nas-port
— nas-port-id [prefix-type {none | user-string}] [prefix-string
prefix-string] [suffix {circuit-id | none | remote-id | user-
string}] [suffix-string suffix-string]
— no nas-port-id
— nas-port-type [ [0..255] ]
— no nas-port-type
— [no] physical-access-id
— [no] rat-type
— [no] supported-features
— user-equipment-info [type ue-info-type]
— no user-equipment-info
— mac-format mac-format
— no mac-format
— [no] report-ip-address-event
— gy
— avp-subscription-id origin [type type]
— no avp-subscription-id
— out-of-credit-reporting {final | quota-exhausted}
— no out-of-credit-reporting
— vendor-support {three-gpp | vodafone}
— no vendor-support
— [no] include-avp
— 3gpp-imsi {circuit-id | imsi | subscriber-id}
— no 3gpp-imsi
— called-station-id [64 chars max]
— no called-station-id
— [no] radius-user-name
— service-context-id name
— no service-context-id
— nasreq

```

- [no] **include-avp**
  - [no] **called-station-id**
  - **calling-station-id** [type {llid | mac | remote-id | sap-id | sap-string}]
  - no **calling-station-id**
  - [no] **circuit-id**
  - **nas-port** *binary-spec*
  - no **nas-port**
  - **nas-port-id** [prefix-type {none | user-string}] [prefix-string [8 chars max]] [suffix-type {circuit-id | none | remote-id | user-string}] [suffix-string [64 chars max]]
  - no **nas-port-id**
  - **nas-port-type** [0..255]
  - **nas-port-type**
  - no **nas-port-type**
  - [no] **remote-id**
- **mac-format** *mac-format*
- no **mac-format**
- **password** *password*
- no **password**
- **user-name-format** *format*
- no **user-name-format**
- **user-name-operation** *operation* [domain *domain-name*]
- no **user-name-operation**
- **on-failure** [failover {enabled | disabled}] [handling {continue | retry-and-terminate | terminate}]
- no **on-failure**
- **tx-timer** *seconds*
- no **tx-timer**

## Category Map and Credit Control Policy Commands

```

config
— subscriber-mgmt
— category-map category-map-name [create]
— no category-map category-map-name
— description description-string
— no description
— activity-threshold kilobits-per-second
— no activity-threshold
— category category-name [create]
— no category category-name
— credit-type-override {volume | time}
— no credit-type-override
— default-credit volume credits bytes|kilobytes|megabytes|gigabytes
— default-credit time seconds
— no default-credit
— description description-string
— no description
— [no] exhausted-credit-service-level
— [no] egress-ip-filter-entries
— entry entry-id [create]
— action drop
— action forward
— action http-redirect url
— no action
— description description-string
— no description
— match [protocol protocol-id]
— no match
— dscp dscp-name
— no dscp
— dst-port {lt|gt|eq} dst-port-number
— dst-port range start end
— no dst-port
— fragment {true|false}
— icmp-code icmp-code
— no icmp-code
— icmp-type icmp-type
— no icmp-type
— ip-option ip-option-value [ip-option-mask]
— no ip-option
— multiple-option {true | false}
— option-present {true | false}
— src-ip {ip-address/mask | ip-address net-
mask}
— no src-ip
— src-port {lt|gt|eq} src-port-number
— src-port range start end
— no src-port
— tcp-ack {true|false}
— no tcp-ack
— tcp-syn {true|false}
— no tcp-syn

```

```

— [no] ingress-ip-filter-entries
  — entry entry-id [create]
    — action drop
    — action forward
    — action http-redirect url
    — no action
    — description description-string
    — no description
    — match [protocol protocol-id]
    — no match
      — dscp dscp-name
      — no dscp
      — dst-ip {ip-address/mask | ip-address net-mask}
      — no dst-ip
      — dst-port {lt|gt|eq} dst-port-number
      — dst-port range start end
      — no dst-port
      — fragment {true|false}
      — icmp-code icmp-code
      — no icmp-code
      — icmp-type icmp-type
      — no icmp-type
      — ip-option ip-option-value [ip-option-mask]
      — no ip-option
      — multiple-option {true|false}
      — option-present {true|false}
      — src-port {lt|gt|eq} src-port-number
      — src-port range start end
      — no src-port
      — tcp-ack {true|false}
      — no tcp-ack
      — tcp-syn {true|false}
      — no tcp-syn
    — pir pir-rate
    — pir max
    — no pir
  — out-of-credit-action-override {continue | block-category | change-service-level}
  — no out-of-credit-action-override
  — policer policer-id {ingress-only|egress-only|ingress-egress}
  — no policer policer-id
  — queue queue-id {ingress-only | egress-only | ingress-egress}
  — no queue queue-id
  — rating-group rating-group-id
  — no rating-group
— credit-exhaust-threshold threshold-percentage
— no credit-exhaust-threshold
— credit-type {volume | time}
— no credit-type
— credit-control-policy policy-name [create]
— no credit-control-policy policy-name
  — credit-control-server radius
  — no credit-control-server
  — default-category-map category-map-name

```

- **no default-category-map**
- **description** *description-string*
- **no description**
- **error-handling-action** {**continue** | **block**}
- **no error-handling-action**
- **out-of-credit-action** *action*
- **no out-of-credit-action**

## Filter Commands

```

config
— filter
— copy {ip-filter | mac-filter | ipv6-filter} src-filter-id [src-entry src-entry-id] to dst-filter-id
  [dst-entry dst-entry-id] [overwrite]
— ip-filter filter-id [create]
— no ip-filter filter-id
  — default-action drop|forward
  — description description-string
  — entry entry-id [time-range time-range-name] [create]
  — no entry entry-id
    — action drop|forward
    — no action
    — log log-id
    — no log
    — match [next-header next-header]
    — no match
      — dscp
      — no dscp
      — dst-ip
      — no dst-ip
      — dst-port
      — no dst-port
      — icmp-code
      — no icmp-code
      — icmp-type
      — no icmp-type
      — src-ip
      — no src-ip
      — src-port
      — no src-port
      — tcp-ack
      — no tcp-ack
      — tcp-syn
      — no tcp-syn
  — group-inserted-entries application application location location
  — renum old-entry-id new-entry-id
  — scope exclusive|template
  — no scope
  — shared-radius-filter-wmark low low-watermark high high-watermark
  — sub-insert-radius start-entry entry-id count count
  — no sub-insert-radius
  — sub-insert-credit-control start-entry entry-id count count
  — no sub-insert-credit-control
  — sub-insert-shared-radius start-entry entry-id count count
  — sub-insert-wmark [low percentage] [high percentage]
  — no sub-insert-wmark
— ipv6-filter ipv6-filter-id [create]
— no ipv6-filter ipv6-filter-id
  — default-action drop|forward
  — description description-string
  — entry entry-id [time-range time-range-name] [create]
  — no entry entry-id
    — action drop|forward

```

- **no action**
- **log** *log-id*
- **no log**
- **match** [**next-header** *next-header*]
- **no match**
  - **dscp**
  - **no dscp**
  - **dst-ip**
  - **no dst-ip**
  - **dst-port**
  - **no dst-port**
  - **icmp-code**
  - **no icmp-code**
  - **icmp-type**
  - **no icmp-type**
  - **src-ip**
  - **no src-ip**
  - **src-port**
  - **no src-port**
  - **tcp-ack**
  - **no tcp-ack**
  - **tcp-syn**
  - **no tcp-syn**
- **group-inserted-entries** **application** *application* **location** *location*
- **renum** *old-entry-id* *new-entry-id*
- **scope** *exclusive|template*
- **no scope**
- **shared-radius-filter-wmark** **low** *low-watermark* **high** *high-watermark*
- **sub-insert-radius** **start-entry** *entry-id* **count** *count*
- **no sub-insert-radius**
- **sub-insert-credit-control** **start-entry** *entry-id* **count** *count*
- **no sub-insert-credit-control**
- **sub-insert-shared-radius** **start-entry** *entry-id* **count** *count*
- **sub-insert-wmark** [**low** *percentage*] [**high** *percentage*]
- **no sub-insert-wmark**



## Explicit Subscriber Mapping Commands

```
config
— subscriber-mgmt
    — explicit-sub-map
        — entry key sub-ident-string [sub-profile sub-profile-name] [alias sub-alias-string] [sla-profile sla-profile-name]
        — no entry key sub-ident-string
```

## Host Tracking Policy Commands

```
config
— subscriber-mgmt
  — host-tracking-policy policy-name [create]
  — no host-tracking-policy policy-name
    — description description-string
    — no description
    — egress-rate-modify agg-rate-limit
    — egress-rate-modify scheduler scheduler-name
    — no egress-rate-modify
```

## PIM Policy Commands

```
config
— subscriber-mgmt
    — pim-policy policy-name [create]
    — no pim-policy policy-name
        — description description-string
        — no description
```

## SLA Profile Commands

```

config
  — subscriber-mgmt
    — [no] sla-profile sla-profile-name
      — category-map category-map-name [create]
      — no category-map category-map-name
        — category category-name [create]
        — no category category-name
          — idle-timeout timeout
          — no idle-timeout
          — idle-timeout-action {shcv-check|terminate}
          — no idle-timeout-action
      — description description-string
      — no description
      — egress
        — [no] ip-filter filter-id
        — qos sap-egress-policy-id [vport-scheduler|port-scheduler] [force]
        — no qos
          — queue queue-id
          — no queue queue-id
            — avg-frame-overhead percent
            — no avg-frame-overhead
            — burst-limit
            — no burst-limit
            — burst-limit size-in-kbytes
            — no burst-limit
            — high-prio-only percent
            — no high-prio-only
            — mbs size-in-kbytes
            — no mbs
            — rate pir-rate [cir cir-rate]
            — no rate
            — stat-mode stat-mode
            — no stat-mode
          — policer policer-id [create]
          — no policer policer-id
            — cbs {size [bytes | kilobytes] | default}
            — no cbs
            — mbs {size [bytes | kilobytes] | default}
            — no mbs
            — packet-byte-offset {add bytes | subtract bytes}
            — no packet-byte-offset
            — rate {max | kilobits-per-second} [cir {max | kilobits-per-second}]
            — no rate
            — stat-mode stat-mode
            — no stat-mode
        — [no] qos-marking-from-sap
        — report-rate agg-rate-limit
        — report-rate scheduler scheduler-name
        — report-rate pppoe-actual-rate
        — report-rate rfc5515-actual-rate
        — no report-rate
        — scheduler-policy scheduler-policy-name

```

```

— no scheduler-policy
— [no] use-ingress-l2tp-dscp
— [no] host-limits
— ipv4-arp max-nr-of-hosts
— no ipv4-arp
— ipv4-dhcp max-nr-of-hosts
— no ipv4-dhcp
— ipv4-overall max-nr-of-hosts
— no ipv4-overall
— ipv4-ppp max-nr-of-hosts
— no ipv4-ppp
— ipv6-overall max-nr-of-hosts
— no ipv6-overall
— ipv6-pd-ipoe-dhcp max-nr-of-hosts
— no ipv6-pd-ipoe-dhcp
— ipv6-pd-overall max-nr-of-hosts
— no ipv6-pd-overall
— ipv6-pd-ppp-dhcp max-nr-of-hosts
— no ipv6-pd-ppp-dhcp
— ipv6-wan-ipoe-dhcp max-nr-of-hosts
— no ipv6-wan-ipoe-dhcp
— ipv6-wan-ipoe-slaac max-nr-of-hosts
— no ipv6-wan-ipoe-slaac
— ipv6-wan-overall max-nr-of-hosts
— no ipv6-wan-overall
— ipv6-wan-ppp-dhcp max-nr-of-hosts
— no ipv6-wan-ppp-dhcp
— ipv6-wan-ppp-slaac max-nr-of-hosts
— no ipv6-wan-ppp-slaac
— lac-overall max-nr-of-hosts
— no lac-overall
— overall max-nr-of-hosts
— no overall
— [no] remove-oldest
— ingress
— [no] ip-filter filter-id
— qos sap-ingress-policy-id [shared-queuing | multipoint-shared | service-
queuing] [force]
— no qos
— queue queue-id
— no queue queue-id
— burst-limit
— no burst-limit
— burst-limit size-in-kbytes
— no burst-limit
— high-prio-only percent
— no high-prio-only
— mbs size-in-kbytes
— no mbs
— rate pir-rate [cir cir-rate]
— no rate
— stat-mode {v4-v6}
— no stat-mode
— policer policer-id [create]
— no policer policer-id

```

- **cbs** {*size* [bytes | kilobytes] | default}
- **no cbs**
- **mbs** {*size* [bytes | kilobytes] | default}
- **no mbs**
- **packet-byte-offset** {add *bytes* | subtract *bytes*}
- **no packet-byte-offset**
- **rate** {max | kilobits-per-second} [cir {max | kilobits-per-second}]
- **no rate**
- **stat-mode** {v4-v6}
- **no stat-mode**
- **report-rate** agg-rate-limit
- **report-rate** scheduler *scheduler-name*
- **no report-rate**
- **one-time-http-redirection** *filter-id*

## Subscriber Identification Policy Commands

```

config
— subscriber-mgmt
  — [no] sub-ident-policy sub-ident-policy-name
    — app-profile-map
      — entry key app-profile-string app-profile app-profile-name
      — no entry key app-profile-string
      — [no] use-direct-map-as-default
    — description description-string
    — no description
    — primary
      — script-url dhcp-primary-script-url
      — no script-url
      — [no] shutdown
    — secondary
      — script-url dhcp-secondary-script-url
      — no script-url
      — [no] shutdown
    — sla-profile-map
      — entry key sla-profile-string sla-profile sla-profile-name
      — no entry key sla-profile-string
      — [no] use-direct-map-as-default
    — sub-profile-map
      — entry key sub-profile-string sub-profile sub-profile-name
      — no entry key sub-profile-string
      — [no] use-direct-map-as-default
    — tertiary
      — script-url dhcp-tertiary-script-url
      — no script-url
      — [no] shutdown

```

## Auto-Generated Subscriber Identification Key Commands

Note: These commands are supported on the 7450 ESS in mixed mode.

```

config
  — subscriber-mgmt
    — auto-sub-id-key
      — ipoe-sub-id-key sub-id-key [sub-id-key...(up to 4 max)]
      — no ipoe-sub-id-key
      — ppp-sub-id-key sub-id-key [sub-id-key...(up to 5 max)]
      — no ppp-sub-id-key

```

## Auto-Generated Subscriber Identification Key Service Commands

Notes: Refer to the 7750 SR OS Services Guide/7450 ESS OS Services Guide for further services commands.

```

config
  — service
    — vprn
      — subscriber-interface ip-int-name
        — group-interface ip-int-name
          — sap sap-id
            — sub-sla-mgmt
              — def-sub-id use-auto-id
              — def-sub-id use-sap-id
              — def-sub-id string sub-id
              — no def-sub-id
      — vpls
        — sap sap-id
          — sub-sla-mgmt
            — def-sub-id use-auto-id
            — def-sub-id use-sap-id
            — def-sub-id string sub-id
            — no def-sub-id
      — ies
        — subscriber-interface ip-int-name
          — group-interface ip-int-name
            — sap sap-id
              — sub-sla-mgmt
                — def-sub-id use-auto-id
                — def-sub-id use-sap-id
                — def-sub-id string sub-id
                — no def-sub-id

  — subscriber-management
    — msap-policy msap-policy-name
      — sub-sla-mgmt

```



- **def-sub-id** **use-auto-id**
- **def-sub-id** **use-sap-id**
- **def-sub-id** **string** *sub-id*
- **no** **def-sub-id**

## Subscriber MCAC Policy Commands

- config**
  - **subscriber-mgmt**
    - **sub-mcac-policy** *sub-mcac-policy-name* [**create**]
    - **no** **sub-mcac-policy** *sub-mcac-policy-name*
      - **description** *description-string*
      - **no** **description**
      - **[no]** **shutdown**
      - **unconstrained-bw** *mandatory-bw mandatory-bw*
      - **no** **unconstrained-bw**

## Subscriber Profile Commands

For information about configuring accounting policies, refer to the **SR OS System Management Guide**.

```

config
  — subscriber-mgmt
    — [no] sub-profile subscriber-profile-name
      — accounting-policy acct-policy-id
      — no accounting-policy
      — ancp
        — ancp-policy name
        — no ancp-policy
          — egress
            — rate-adjustment
            — rate-adjustment adjusted-percent
            — no rate-adjustment
            — rate-modify scheduler scheduler-name
            — no rate-modify
            — rate-monitor kilobit-per-second [alarm]
            — no rate-monitor
            — rate-reduction kilobit-per-second
            — no rate-reduction
          — ingress
            — rate-adjustment adjusted-percent
            — no rate-adjustment
            — rate-modify scheduler scheduler-name
            — no rate-modify
            — rate-monitor kilobit-per-second [alarm]
            — no rate-monitor
            — rate-reduction kilobit-per-second
            — no rate-reduction
      — [no] collect-stats
      — description description-string
      — no description
      — egress
        — agg-rate-limit agg-rate [queue-frame-based-accounting]
        — no agg-rate-limit
        — avg-frame-size bytes
        — no avg-frame-size
        — encap-offset [type type]
        — no encap-offset
        — lag-per-link-hash class {1|2|3} weight 1..1024
        — no lag-per-link-hash
        — policer-control-policy policy-name
        — no policer-control-policy
          — max-rate {kilobits-per-second | max}
          — no max-rate
          — priority-mbs-thresholds
            — min-thresh-separation size [bytes | kilobytes]
            — no min-thresh-separation
            — priority level
            — mbs-contribution size [bytes | kilobytes] [fixed]
        — scheduler-policy scheduler-policy-name
  
```

- **no scheduler-policy**
  - **scheduler** *scheduler-name* **rate** [*pir-rate*] [**cir** *cir-rate*]
  - **no scheduler** *scheduler-name*
- **hsmda**
  - **egress-qos**
    - **agg-rate-limit** *agg-rate*
    - **no agg-rate-limit**
    - **qos** *policy-id*
    - **no qos**
      - **wrr-policy** *weight*
      - **no wrr-policy**
      - **packet-byte-offset** {**add** *add-bytes* | **subtract** *sub-bytes*}
      - **no packet-byte-offset**
      - **queue** *queue-id* [**create**]
      - **no queue** *queue-id*
        - **rate** *pir-rate* [**cir** *cir-rate*]
        - **no rate**
        - **rate** {**max** | *kilobits-per-second*}
        - **no rate**
        - **slope-policy** *hsmda-slope-policy-name*
        - **no slope-policy**
        - **stat-mode** {**v4-v6**}
        - **no stat-mode**
        - **wrr-weight** *weight*
        - **no wrr-weight**
  - **ingress-qos**
    - **qos** *policy-id*
    - **no qos**
      - **queue** *queue-id*
        - **rate** {**max** | *kilobits-per-second*}
        - **no rate**
        - **cbs** {*size* [**bytes** | **kilobytes**] | **default**}
        - **no cbs**
        - **mbs** *kilobits*
        - **no mbs**
        - **stat-mode** {**v4-v6**}
        - **no stat-mode**
      - **policer** *policer-id* [**create**]
        - **cbs** {*size* [**bytes** | **kilobytes**] | **default**}
        - **no cbs**
        - **mbs** *kilobits*
        - **no mbs**
        - **packet-byte-offset** {**add** *add-bytes* | **subtract** *sub-bytes*}
        - **no packet-byte-offset**
        - **rate** {**max** | *kilobits-per-second*}
        - **no rate**
        - **stat-mode** {**v4-v6**}
        - **no stat-mode**
  - **ingress**
    - **policer-control-policy** *policer-control-policy-name*
    - **no policer-control-policy**
      - **max-rate** {*rate* | **max**}
      - **no max-rate**

- **priority-mbs-thresholds**
  - **min-thresh-separation** *size* [bytes|kilobytes]
  - **min-thresh-separation** **default**
  - **no min-thresh-separation**
  - **priority** *level*
  - **mbs-contribution** *size* [bytes | kilobytes]
- **scheduler-policy** *scheduler-policy-name*
- **no scheduler-policy**
  - **scheduler** *scheduler-name* **rate** [*pir-rate*] [*cir cir-rate*]
  - **no scheduler** *scheduler-name*
- **pim-policy** *policy-name*
- **no pim-policy** *policy-name*
- **radius-accounting-policy**
- **no radius-accounting-policy**
- **sla-profile-map**
  - **entry** **key** *sla-profile-string* **sla-profile** *sla-profile*
  - **no entry** **key** *sla-profile-string*
  - **[no] use-direct-map-as-default**
- **sub-mcac-policy** *policy-name*
- **no sub-mcac-policy**
- **volume-stats-type** {ip|default}
- **no volume-stats-type**

## IPoE Session Policy Commands

- config**
  - **subscriber-mgmt**
    - **ipoe-session-policy** *policy-name* [**create**]
    - **no ipoe-session-policy** *policy-name*
      - **description** *description-string*
      - **no description**
      - **session-key** **sap** **mac** [*cid*] [*rid*]
      - **no session-key**
      - **session-timeout** *timeout*
      - **no session-timeout**

## Multi-Chassis Redundancy Commands

```

config
— redundancy
    — multi-chassis
        — [no] peer ip-address
            — authentication-key [authentication-key | hash-key] [hash | hash2]
            — no authentication-key
            — description description-string
            — no description
            — mc-ipsec
                — bfd-enable service service-id interface interface-name dst-ip
                    ip-address
                — no bfd-enable
                — discovery-interval interval-1 [boot interval-2]
                — no discovery-interval
                — hold-on-neighbor-failure multiplier
                — no hold-on-neighbor-failure
                — keep-alive-interval time-interval
                — no keep-alive-interval
                — tunnel-group group-id [create]
                — no tunnel-group group-id
                    — peer-group group-id
                    — no peer-group
                    — priority priority
                    — no priority
                    — [no] preempt
                    — [no] shutdown
            — [no] mc-lag
                — hold-on-neighbor-failure multiplier
                — no hold-on-neighbor-failure
                — keep-alive-interval interval
                — no keep-alive-interval
                — lag lag-id lacp-key admin-key system-id system-id [remote-lag
                    lag-id] system-priority system-priority
                — no lag lag-id
                — [no] shutdown
        — mc-ring
            — [no] ring sync-tag
                — in-band-control-path
                    — [no] debounce
                    — dst-ip ip-address
                    — no dst-ip
                    — interface ip-int-name
                    — no interface
                    — no max-debounce-time
                    — service-id service-id
                    — no service-id
                — [no] path-b
                    — [no] range vlan-range
                — [no] path-excl
                    — [no] range vlan-range
                — [no] ring-node ring-node-name
                    — connectivity-verify

```

```

— dst-ip ip-address
— no dst-ip
— interval interval
— no interval
— service-id service-id
— no service-id
— [no] shutdown
— src-ip ip-address
— no src-ip
— src-mac ieee-address
— no src-mac
— vlan [0..4094]
— no vlan
— [no] shutdown
— [no] l3-ring sync-tag
— in-band-control-path
— [no] debounce
— dst-ip ip-address
— no dst-ip
— interface ip-int-name
— no interface
— max-debounce-time max-debounce-time
— no max-debounce-time
— service-id service-id
— no service-id
— [no] ring-node ring-node-name
— connectivity-verify
— dst-ip ip-address
— no dst-ip
— interval interval
— no interval
— service-id service-id
— no service-id
— [no] shutdown
— src-ip ip-address
— no src-ip
— src-mac ieee-address
— no src-mac
— vlan [0..4094]
— no vlan
— [no] srrp-instance srrp-id
— [no] shutdown
— [no] shutdown
— source-address ip-address
— no source-address
— [no] sync
— [no] igmp
— [no] igmp-snooping
— [no] ipsec
— [no] local-dhcp-server
— [no] mc-ring
— [no] mld-snooping
— port [port-id|lag-id] [sync-tag sync-tag] [create]
— no port [port-id | lag-id]
— range encap-range [sync-tag sync-tag]

```

- **no range** *encap-range*
- **[no] shutdown**
- **[no] srrp**
- **[no] sub-host-trk**
- **[no] sub-mgmt**
- **tunnel-group** *tunnel-group-id sync-tag tag-name* **[create]**
- **no tunnel-group**

## Subscriber Management Service Commands

### VPLS Subscriber Management Configuration Commands

```

config
  — service
    — [no] vpls service-id [customer customer-id] [vpn vpn-id] [m-vpls]
      — sap sap-id [split-horizon-group group-name] [create] [capture-sap]
      — no sap sap-id
        — arp-host
          — host-limit max-num-hosts
          — no host-limit
          — min-auth-interval min-auth-interval
          — no min-auth-interval
          — [no] shutdown
        — diameter-auth-policy name
        — no diameter-auth-policy
        — host [ip ip-address [mac mac-address]] [subscriber-sap-id | subscriber
          sub-ident-string [sub-profile sub-profile-name [sla-profile sla-profile-
          name [ancp-string ancp-string] [app-profile app-profile-name] [inter-
          dest-id intermediate-destination-id]]
        — no host [ip ip-address [mac mac-address]]
        — no host all
        — [no] sub-sla-mgmt
          — def-app-profile default-app-profile-name
          — no def-app-profile
          — def-sla-profile default-sla-profile-name
          — no def-sla-profile
          — def-sub-id string sub-ident-string
          — def-sub-id use-sap-id
          — no def-sub-id
          — def-sub-profile default-sub-profile-name
          — no def-sub-profile
          — [no] mac-da-hashing
          — multi-sub-sap number-of-sub
          — no multi-sub-sap
          — [no] shutdown
          — single-sub-parameters
            — non-sub-traffic sub-profile sub-profile-name sla-
              profile sla-profile-name [subscriber sub-ident-string]
            — no non-sub-traffic
            — [no] profiled-traffic-only
          — sub-ident-policy sub-ident-policy-name
          — no sub-ident-policy
      — sap sap-id [create] capture-sap
        — ipoe-session
          — description description-string
          — no description
          — ipoe-session-policy policy-name
          — no ipoe-session-policy
          — user-db local-user-db-name
          — no user-db
          — [no] shutdown

```



## Managed SAP Policy Commands

```

config
— subscriber-mgmt
— msap-policy msap-policy-name [create]
— no msap-policy msap-policy-name
— cpu-protection policy-id [mac-monitoring ]
— no cpu-protection
— description description-string
— no description
— no dist-cpu-protection policy-name
— ies-vprn-only-sap-parameters
— anti-spoof {ip-mac | nh-mac}
— no anti-spoof
— igmp-host-tracking
— expiry-time expiry-time
— no expiry-time
— import policy-name
— no import policy-name
— max-num-group max-num-groups
— no max-num-group
— max-num-sources max-num-sources
— no max-num-sources
— max-num-grp-sources [1..32000]
— no max-num-grp-sources [1..32000]
— lag-link-map-profile link-map-profile-id
— no lag-link-map-profile
— sub-sla-mgmt
— def-app-profile app-profile-name
— no def-app-profile
— def-inter-dest-id {string string | use-top-q}
— no def-inter-dest-id
— def-sla-profile sla-profile-name
— no def-sla-profile
— def-sub-id use-sap-id
— def-sub-id string sub-id
— no def-sub-id
— def-sub-profile sub-profile-name
— no def-sub-profile
— multi-sub-sap [limit limit]
— no multi-sub-sap
— single-sub-parameters
— non-sub-traffic sub-profile-name sla-profile sla-profile-name
[subscriber sub-ident-string] [app-profile app-profile-name]
— no non-sub-traffic
— [no] profiled-traffic-only
— sub-ident-policy policy-name
— no sub-ident-policy
— vpls-only-sap-parameters
— arp-host
— host-limit max-num-hosts
— no host-limit
— min-auth-interval min-auth-interval
— no min-auth-interval
— arp-reply-agent [sub-ident]

```

- **no arp-reply-agent**
- **dhcp**
  - **lease-populate** [*nbr-of-leases*]
  - **no lease-populate**
  - **[no] option**
    - **action** *dhcp-action*
    - **no action**
    - **circuit-id** {*ascii-tuple* | *vlan-ascii-tuple*}
    - **remote-id** {*mac* | *string string*}
    - **no remote-id**
    - **[no] vendor-specific-option**
      - **[no] client-mac-address**
      - **[no] sap-id**
      - **[no] service-id**
      - **[no] string**
      - **[no] system-id**
  - **proxy-server**
    - **emulated-server** *ip-address*
    - **no emulated-server**
    - **lease-time** [*days days*] [*hrs hours*] [*min minutes*] [*sec seconds*] [*radius-override*]
    - **no lease-time**
    - **[no] shutdown**
- **egress**
  - **multicast-group** *group-name*
  - **no multicast-group**
- **igmp-snooping**
  - **[no] fast-leave**
  - **import** *policy-name*
  - **no import**
  - **last-member-query-interval** *interval*
  - **no last-member-query-interval**
  - **max-num-groups** *max-num-groups*
  - **no max-num-groups**
- **mcac**
  - **mc-constraints**
    - **level** *level-id* **bw** *bandwidth*
    - **no level** *level-id*
    - **number-down** *number-lag-port-down* **level** *level-id*
    - **no number-down** *number-lag-port-down*
    - **policy** *policy-name*
    - **no policy**
    - **unconstrained-bw** *bandwidth* **mandatory-bw** *mandatory-bw*
    - **no unconstrained-bw**
- **mvr**
  - **from-vpls** *service-id*
  - **no from-vpls**
  - **query-interval** *seconds*
  - **no query-interval**
  - **query-response-interval** *seconds*
  - **no query-response-interval**
  - **robust-count** *robust-count*
  - **no robust-count**

- [no] **send-queries**
- **version** *version*
- **no version**
- [no] **mac-da-hashing**
- **split-horizon-group** *group-name*
- **no split-horizon-group**

## Subscriber Management Service Commands

```
config
  — service
    — vpls
      — sap
        — default-msap-policy policy-name
        — no default-msap-policy
        — trigger-packet [dhcp] [pppoe] [arp] [dhcp6] [ppp]
        — no trigger-packet
        — eval-msap {policy msap-policy-name | msap sap-id}
```

## Service Subscriber Interface, Group Interface IPoE Commands

```
config
  — service
    — vpn service-id [customer customer-id]
    — no vpn service-id
    — ies service-id [customer customer-id]
    — no ies service-id
    — [no] ipoe-session
      — description description-string
      — no description
      — force-auth [cid-change] [rid-change]
      — force-auth disabled
      — no force-auth
      — ipoe-session-policy policy-name
      — no ipoe-session-policy
      — min-auth-interval [days days] [hrs hours] [min minutes] [sec seconds]
      — min-auth-interval infinite
      — no min-auth-interval
      — sap-session-limit sap-session-limit
      — no sap-session-limit
      — session-limit session-limit
      — no session-limit
      — [no] shutdown
      — user-db local-user-db-name
      — no user-db
```

## RIP Commands

```

configure
  — subscriber-mgmt
    — rip-policy policy policy-name [create]
    — no rip-policy policy-name
      — authentication-key authentication-key|hash-key [hash|hash2]
      — no authentication-key
      — authentication-type {none|password|message-digest|message-digest-20}
      — no authentication-type
      — description description-string
      — no description
    — local-user-db local-user-db-name [create]
    — no local-user-db local-user-db-name
      — ppp
        — host host-name [create]
        — no host host-name
          — rip-policy policy-name
          — no rip-policy
      — ipoe
        — host host-name [create]
        — no host host-name
          — rip-policy policy-name
          — no rip-policy

configure
  — service
    — vprn
      — [no] rip
        — [no] group name
        — [no] neighbor ip-int-name

configure
  — router
    — [no] rip
      — [no] group name
      — [no] neighbor ip-int-name

```

## VPort Commands

Refer to the SR OS Interfaces Guide for further information on card, Media Dependent Adapter (MDA), MCM (MDA Carrier Module), CMA (Compact Media Adapter) and port provisioning.

```

config
  — port port-id
    — ethernet
      — access
        — egress
          — vport vport-name [create]
          — no vport vport-name
            — agg-rate-limit limit
            — no agg-rate-limit
            — description description-string
            — no description
            — [no] egress-rate-modify
            — host-match dest string [create]
            — no host-match dest string
            — port-scheduler-policy port-scheduling-policy-name
            — no port-scheduler-policy
            — scheduler-policy scheduler-policy-name
            — no scheduler-policy

```

```

config
— service
— sdp sdp-id [gre | mpls]
— no sdp sdp-id
— binding
— port [port-id | lag-id]
— no port
— pw-port pw-port-id [vc-id vc-id] [create]
— no pw-port pw-port-id
— description description-string
— no description
— egress
— shaper
— int-dest-id int-dest-id
— no int-dest-id int-dest-id
— vport vport-name
— no vport vport-name
— encap-type {dot1q|qinq}
— no encap-type
— [no] shutdown
— vc-type {ether | vlan}
— no vc-type
— vlan-vc-tag vlan-id
— no vlan-vc-tag

```

## Wireless Portal Protocol (WPP) Commands

```

config
  — router
    — wpp
      — portals
        — portal wpp-portal-name address ip-address [create]
        — portal wpp-portal-name
        — no portal wpp-portal-name
          — [no] shutdown
        — [no] shutdown

config
  — service
    — vpn service-id [customer customer-id]
    — no vpn service-id
      — wpp
        — portals
          — portal wpp-portal-name address ip-address [create]
          — portal wpp-portal-name
          — no portal wpp-portal-name
            — [no] shutdown
          — [no] shutdown

config
  — service
    — ies service-id [customer customer-id] [vpn vpn-id]
    — vpn service-id [customer customer-id]
    — no vpn service-id
      — wpp
        — [no] enable-triggered-hosts
        — initial-app-profile app-profile-name
        — no initial-app-profile
        — initial-sla-profile sla-profile-name
        — no initial-sla-profile
        — initial-sub-profile sub-profile-name
        — no initial-sub-profile
        — lease-time [days days] [hrs hours] [min minutes] [sec
          seconds]
        — no lease-time
        — portal router router-instance name wpp-portal-name
        — no portal
        — restore-disconnected {restore|no-restore}
        — no restore-disconnected
        — [no] shutdown
        — user-db local-user-db-name
        — no user-db
        — [no] shutdown

config
  — subscriber-mgmt
    — local-user-db local-user-db-name [create]

```



```

— no local-user-db local-user-db-name
    — dhcp
        — host host-name [create]
        — no host host-name
            — wpp
                — initial-app-profile app-profile-name
                — no initial-app-profile
                — initial-sla-profile sla-profile-name
                — no initial-sla-profile
                — initial-sub-profile sub-profile-name
                — no initial-sub-profile
                — lease-time [days days] [hrs hours] [min minutes] [sec
                    seconds]
                — no lease-time
                — portal router router-instance name wpp-portal-name
                — no portal
                — restore-disconnected {restore|no-restore}
                — no restore-disconnected

```

## Multiple PPOE Session QoS Commands

```

config
  — qos
      — sap-egress policy-id
          — parent-location {default | sla}
          — no parent-location

      — scheduler-policy name [create]
          — tier 1
              — parent-location {none | sub | vport}
              — no parent-location

config
  — port
      — ethernet
          — access
              — egress
                  — vport name
                      — scheduler-policy scheduler-policy-name

config
  — subscr-mgmt
      — sla-prof
          — egress$
              — scheduler-policy scheduler-policy-name
                  — scheduler scheduler-name rate pir-rate [cir cir-rate]

configure
  — subscriber-mgmt
      — authentication-policy name
          — include-radius-attribute
              — [no] sap-session-index
  
```

## Multicast Listener Discovery (MLD) Commands

For more information about MLD commands, refer to the SR OS Routing Protocols Guide.

```

config
— subscrip-mgmt
— mld-policy mld-policy-name [create]
— no mld-policy mld-policy-name
— description description-string
— no description
— egress-rate-modify agg-rate-limit
— egress-rate-modify scheduler scheduler-name
— no egress-rate-modify
— [no] fast-leave
— import policy-name
— no import
— max-num-groups max-num-groups
— no max-num-groups
— max-num-grp-sources [1..32000]
— no max-num-grp-sources
— max-num-sources max-num-sources
— no max-num-sources
— [no] per-host-replication
— redirection-policy policy-name
— no redirection-policy
— static
— [no] group grp-ipv6-address
— [no] source ipv6-address
— [no] starg
— version version
— no version

```

## Show Commands

```

show
  — service
    — active-subscribers detail
    — active-subscribers mirror
    — active-subscribers [summary]
      — credit-control credit-control [subscriber sub-ident-string]
      — credit-control out-of-credit [action action] [summary]
      — filter [subscriber sub-ident-string] [origin origin]
      — hierarchy [subscriber sub-ident-string]
      — host-tracking [subscriber sub-ident-string]
      — host-tracking [subscriber sub-ident-string] detail
      — host-tracking [subscriber sub-ident-string] summary
      — host-tracking [subscriber sub-ident-string] statistics
        — groups [group group-ip-address]
        — groups group group-ip-address detail
        — groups group group-ip-address summary
      — igmp [subscriber sub-ident-string][detail]
      — subscriber sub-ident-string
      — subscriber sub-ident-string detail
      — subscriber sub-ident-string mirror
      — subscriber sub-ident-string sap sap-id sla-profile sla-profile-name
      — subscriber sub-ident-string sap sap-id sla-profile sla-profile-name detail
      — subscriber sub-ident-string sap sap-id sla-profile sla-profile-name mirror
  — id service-id
    — arp-host [wholesaler service-id] [sap sap-id | interface interface-name | ip-address
      ip-address[/mask] | mac ieee-address | {[port port-id] [no-inter-dest-id | inter-
      dest-id inter-dest-id]}] [detail]
    — arp-host statistics [sap sap-id | interface interface-name]
    — arp-host summary [interface interface-name]
    — authentication
      — statistics
    — dhcp
      — lease-state [wholesaler service-id] [sap sap-id | sdp sdp-id:vc-id | inter-
        face 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]}] [detail]
      — statistics [sap sap-id] | [ sdp [sdp-id[:vc-id]] ]
      — summary
    — subscriber-hosts [sap sap-id] [ip ip-address[/mask]] [mac ieee-address] [sub-pro-
      file sub-profile-name] [sla-profile sla-profile-name] [detail]
    — gsm
      — neighbors group [name] [ip-address]
      — sessions [group name] neighbor ip-address [ port port-number] [associ-
        ation] [statistics]
    — host [sap sap-id] [wholesaler service-id] [port port-id] [inter-dest-id intermediate-
      destination-id] [detail]
    — host [sap sap-id] [wholesaler service-id] [port port-id] no-inter-dest-id [detail]
    — host summary
    — host [detail] wholesaler service-id
    — interface [{[ip-address|ip-int-name] [interface-type] [detail] [family]}]summary]
    — ipoe session [sap sap-id] [mac ieee-address] [circuit-id circuit-id] [remote-id
      remote-id] [interface ip-int-name|ip-address] [inter-dest-id intermediate-destina-

```

```

tion-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]
— ipoe session [sap sap-id] [mac ieee-address] [circuit-id circuit-id] [remote-id remote-id] [interface ip-int-name|ip-address] [inter-dest-id intermediate-destination-id] [no-inter-dest-id] [ip-address ip-prefix[/prefix-length]] [port port-id] [subscriber sub-ident-string] [sap-session-id sap-session-index] [wholesaler service-id] detail
— sdp sdp-id pw-port [pw-port-id]
— sdp sdp-id pw-port
— sdp sdp-id pw-port [pw-port-id] [statistics]
— sdp [consistent | inconsistent | na] egressifs
— sdp sdp-id keep-alive-history
— sdp far-end ip-address | ipv6-address keep-alive-history
— sdp [sdp-id] detail
— sdp far-end ip-address | ipv6-address detail
— subscriber-using [service-id service-id] [sap-id sap-id] [interface ip-int-name] [ip ip-address[/mask]] [mac ieee-address] [sub-profile sub-profile-name] [sla-profile sla-profile-name] [app-profile app-profile-name]

show
— pw-port
— pw-port [pw-port-id] [detail]
— pw-port sdp [sdp-id]
— pw-port sdp none

show
— qos
— port-scheduler-policy [port-scheduler-policy-name] [association]
— port-scheduler-policy port-scheduler-policy-name network-policy network-queue-policy-name
— port-scheduler-policy port-scheduler-policy-name sap-egress policy-id
— port-scheduler-policy port-scheduler-policy-name scheduler-policy scheduler-policy-name
— port-scheduler-policy port-scheduler-policy-name scheduler-policy scheduler-policy-name sap-egress policy-id
— sap-egress [policy-id] [association|detail]
— sap-ingress [policy-id] [association | match-criteria | detail]
— scheduler-hierarchy
— customer customer-id site customer-site-name [scheduler scheduler-name] [ingress|egress] [detail]
— port port-id vport name [scheduler scheduler-name] [detail]
— sap sap-id [scheduler scheduler-name] [ingress|egress] [detail]
— subscriber sub-ident-string [scheduler scheduler-name] [ingress|egress] [detail]
— subscriber sub-ident-string sla-profile sla-profile-name sap sap-id [scheduler scheduler-name] [detail]
— scheduler-name scheduler-name
— scheduler-policy [scheduler-policy-name] [association|sap-ingress policy-id|sap-egress policy-id]
— scheduler-stats
— customer customer-id site customer-site-name [scheduler scheduler-name] [ingress|egress] [detail]
— vport port-id vport name [scheduler scheduler-name]

```

```

— sap sap-id [scheduler scheduler-name] [ingress|egress] [detail]
— sla-profile sub-ident-string sla-profile sla-profile-name sap sap-id [scheduler
scheduler-name]
— shared-queue shared-queue-policy-name [detail]

show
— redundancy
— multi-chassis all
— multi-chassis mc-lag
— multi-chassis sync
— mc-ipsec addr [tunnel-group grp-id]
— mc-ring peer ip-address statistics
— mc-ring peer ip-address [ring sync-tag [detail|statistics] ]
— mc-ring peer ip-address ring sync-tag ring-node [ring-node-name [detail|statistics] ]
— mc-ring global-statistics

show
— router
— wpp
— wpp [portal wpp-portal-name] [host ip-address] hosts
— wpp portal wpp-portal-name
— wpp statistics

show
— aaa
— radius-configuration

show
— subscriber-mgmt
— ancp-string [policy-name]
— ancp-string policy-name association
— ancp-string
— ancp-string ancp-string
— ancp-string customer customer-id site customer-site-name
— ancp-string sap sap-id
— authentication policy-name association
— authentication [policy-name]
— authentication [policy-name] statistics
— authentication coa-statistics
— diameter-application-policy [name]
— explicit-subscriber-map
— host-lockout-policy
— host-lockout-policy policy-name association
— host-lockout-policy policy-name
— host-lockout-policy policy-name all
— host-lockout-policy policy-name sap sap-id [circuit-id | mac | remote-id]
— ipoe-session-policy ipoe-session-policy-name association
— ipoe-session-policy
— local-user-db local-user-db-name association [dhcp] [ppp] [l2tp] [radius] [pppoe] [dhcp6]
[capture-sap] [rtr-solicit] [wpp] [ipoe]
— local-user-db local-user-db-name ipoe-host ipoe-host-name
— local-user-db local-user-db-name ipoe-all-hosts
— local-user-db local-user-db-name ipoe-unmatched-hosts

```

- **local-user-db** [*local-user-db-name*]
- **local-user-db** *local-user-db-name* **ppp-all-hosts**
- **local-user-db** *local-user-db-name* **ppp-host** *pppoe-host-name*
- **local-user-db** *local-user-db-name* **ppp-unmatched-hosts**
- **msap-policy** [*msap-policy-name* [**association**]]
- **sla-profile** [*sla-profile-name* [**association**]]
- **statistics iom** (*slot* | **all**) [**host**|**session**|**subscriber**|**summary**] [**non-zero-value-only**]
- **statistics mda** (*mda* | **all**) [**host**|**session**|**subscriber**|**summary**] [**non-zero-value-only**]
- **statistics port** (*port-id* | **all**) [**host**|**session**|**subscriber**|**summary**] [**non-zero-value-only**]
- **statistics pw-port** (*pw-port* | **all**) [**host**|**session**|**subscriber**|**summary**] [**non-zero-value-only**]
- **statistics system** [**host**|**session**|**subscriber**|**summary**] [**non-zero-value-only**]
- **sub-ident-policy** [*sub-ident-policy-name* [**association**]]
- **sub-ident-policy** *sub-ident-policy-name* **script** {**primary** | **secondary** | **tertiary**}
- **sub-profile** [*sub-profile-name* [**association**]]

### Monitor Commands

```
monitor
— service
— subscriber sub-ident-string sap sap-id sla-profile sla-profile-name [base | ingress-queue-id ingress-queue-id | egress-queue-id egress-queue-id] [interval seconds] [repeat repeat] [absolute | rate]
```

### Clear Commands

```
clear
— subscriber-mgmt
— ancp
— ancp-sub-string string
— authentication [policy-name]
— authentication coa-statistics
— msap-policy msap-policy-name
— peakvalue-stats iom (slot | all) [recursive]
— peakvalue-stats mda (mda | all) [recursive]
— peakvalue-stats port (port-id | all)
— peakvalue-stats pw-port (pw-port | all)
— peakvalue-stats system [recursive]
— radius-accounting [policy-name]
— service
— id service-id
— arp-host
— arp-host { 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]
— arp-host statistics [sap sap-id | interface interface-name]
— authentication
— statistics
— ipoe session [sap sap-id] [interface ip-int-name|ip-address] [mac ieee-address] [circuit-id circuit-id] [remote-id remote-id] [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]
— ipoe session all
— msap msap-id
— msap-policy msap-policy-name
— statistics
— subscriber sub-ident-string
— qos
— scheduler-stats
— sla-profile sub-ident-string sla-profile sla-profile-name sap sap-id [scheduler scheduler-name]
clear
— router
— srrp
— interface subscriber-interface [id srrp-id]
— statistics interface subscriber-interface [id srrp-id]
clear
— aaa
— route-downloader name [vprn vprn] [family family]
```



## Debug Commands

```

debug
  — [no] diameter
    — dest-realm realm
    — no dest-realm
    — detail-level level
    — no diameter-peer
    — diameter-peer peer [psm-events]
    — no diameter-peer-policy
    — diameter-peer-policy policy
    — message-type [ccr] [cca] [cer] [cea] [dwr] [dwa] [dpr] [dpa] [rar] [raa] [asr] [asa] [aar]
      [aaa]
    — message-type all
    — no message-type
    — origin-realm realm
    — no origin-realm

debug
  — service
    — id service-id
      — arp
      — [no] arp-host
      — one-time-http-redirection
      — [no] ppp
        — [no] circuit-id circuit-id
        — [no] event
          — dhcp-client [terminate-only]
          — no dhcp-client
          — l2tp [terminate-only]
          — no l2tp
          — local-address-assignment [terminate-only]
          — no local-address-assignment
          — ppp [terminate-only]
          — no ppp
        — [no] mac ieee-address
        — [no] msap msap-id
        — [no] packet
          — detail-level {low|medium|high}
          — no detail-level
          — [no] dhcp-client
          — discovery [padi] [pado] [padr] [pads] [padt]
          — no discovery
          — mode {dropped-only|ingr-and-dropped|egr-ingr-and-
            dropped}
          — no mode
          — ppp [lcp] [pap] [chap] [ipcp] [ipv6cp]
          — no ppp
        — [no] remote-id remote-id
        — [no] sap sap-id
        — [no] username username

    — subscriber-mgmt

```

```

— [no authentication] [policy policy-name] [mac-addr ieee-address] [circuit-id circuit-id]
— [no sub-ident-policy policy-name]
    — script-all-info
    — [no] script-compile-error
    — [no] script-export-variables
    — [no] script-output
    — [no] script-output-on-error
    — [no] script-runtime-error
— router
    — [no] srpp
        — [no] events [interface ip-int-name]
        — [no] packets [interface ip-int-name]
    — [no] radius
        — detail-level {low|medium|high}
        — no detail-level
        — packet-type [authentication] [accounting] [coa]
        — no packet-type
        — radius-attr type attribute-type [transaction]
        — radius-attr type attribute-type [transaction] {address|hex|integer|string} value attribute-value
        — radius-attr vendor vendor-id type attribute-type [transaction encoding-type]
        — radius-attr vendor vendor-id type attribute-type [transaction encoding-type] {address|hex|integer|string} value attribute-value
        — no radius-attr type attribute-type
        — no radius-attr type attribute-type {address|hex|integer|string} attribute-value
        — no radius-attr vendor vendor-id type attribute-type
        — no radius-attr vendor vendor-id type attribute-type {address |string} value attribute-value
        — [no] server-address ip-address
    — [no] wpp
        — [no] packet
            — detail-level detail-level
        — [no] portal wpp-portal-name
            — [no] packet
                — detail-level detail-level

```

## Tools Commands

```

tools
— dump
  — redundancy
    — multi-chassis
      — force-switchover tunnel-group local-group-id
      — mc-ipsec
      — mc-ring
      — srrp-sync-data [instance instance-id] [peer ip-address]
      — srrp-sync-data pppoe
      — srrp-sync-data dhcp pppoe
      — sync-database [peer ip-address] [port port-id | lag-id] [sync-tag sync-tag] [application {dhcps|igmp|igmp-snooping|mc-ring|srrp|sub-mgmt|mld-snooping}] [detail] [type {alarm-deleted|local-deleted}]
— perform
  — aaa
    — route-downloader start [force]
  — persistence
    — downgrade target-version target [reboot]
  — subscriber-mgmt
    — credit-reset sap sap-id subscriber sub-ident-string sla-profile sla-profile-name {category category-name|all-categories}
    — credit-reset sap sap-id ip ip-address {category category-name| all-categories}
    — credit-reset svc service-id ip ip-address {category category-name| all-categories}
    — 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]
    — 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]
    — eval-lease-state sap sap-id ip ip-address [subscriber sub-ident-string] [sub-profile-string sub-profile-string] [sla-profile-string sla-profile-string] [app-profile-string app-profile-string] [inter-dest-id intermediate-destination-id] [ancp-string ancp-string]
    — eval-lease-state svc-id service-id ip ip-address [subscriber sub-ident-string] [sub-profile-string sub-profile-string] [sla-profile-string sla-profile-string] [app-profile-string app-profile-string] [inter-dest-id intermediate-destination-id] [ancp-string ancp-string]
    — re-ident-sub old-sub-ident-string to new-sub-ident-string

```



---

## Triple Play Subscriber Management Configuration Commands

- [Generic Commands on page 1138](#)
- [ANCP and GSMP Commands on page 1140](#)
- [RADIUS Policy Commands on page 1145](#)
- [RADIUS Route Download Commands on page 1179](#)
- [Category Map Commands on page 1183](#)
- [Diameter Commands on page 1196](#)
- [Filter Commands on page 1215](#)
- [PIM Policy Commands on page 1223](#)
- [Managed SAP Policy Commands on page 1224](#)
- [Multi-Chassis Redundancy Commands on page 1251](#)
- [SLA Profile Commands on page 1266](#)
- [Subscriber Identification Policy Commands on page 1292](#)
- [Subscriber Profile Commands on page 1300](#)
- [Subscriber Management Service Commands on page 1324](#)
- [RIP Commands on page 1339](#)
- [Vport Commands on page 1342](#)
- [Clear Commands on page 1415](#)
- [Show Commands on page 1359](#)
- [Tools Commands on page 1420](#)
- [Debug Commands on page 1426](#)
- [Monitor Commands on page 1435](#)

## Generic Commands

### description

<b>Syntax</b>	<b>description</b> <i>description-string</i> <b>no description</b>
<b>Context</b>	config>subscr-mgmt>authentication-policy config>subscr-mgmt>pim-policy config>subscr-mgmt>sla-profile config>subscr-mgmt>sla-profile>egress>ip-filter>entry config>subscr-mgmt>sla-profile>ingress>ip-filter>entry config>subscr-mgmt>sub-ident-policy config>subscr-mgmt>sub-profile config>subscr-mgmt>mld-policy config>service>vpls>gsmp>group config>log>accounting-policy config>service>ies>redundant-interface config>service>ies>subscriber-interface config>service>ies>subscriber-interface>group-interface config>service>ies>subscriber-interface>grp-if>dhcp config>service>ies>sub-if>grp-if>srrp config>service>vprn>sub-if>grp-if>ipoe-session config>redundancy>multi-chassis>peer config>subscr-mgmt>cat-map config>subscr-mgmt>ipoe-session-policy config>service>vpls>sap> ipoe-session config>sub-mgmt>diameter-policy config>sub-mgmt>credit-control-policy config>sub-mgmt>host-lockout>policy config>subscr-mgmt>sub-mcac-policy config>aaa>route-downloader config>aaa>diam-peer-pol config>port>ethernet>access>egress config>subscr-mgmt>cat-map>category config>subscr-mgmt>cat-map>category>exh-lvl>egr-ip
<b>Description</b>	<p>This command creates a text description stored in the configuration file for a configuration context.</p> <p>The <b>description</b> command associates a text string with a configuration context to help identify the context in the configuration file.</p> <p>The <b>no</b> form of this command removes any description string from the context.</p>
<b>Default</b>	No description is associated with the configuration context.
<b>Parameters</b>	<i>description-string</i> — A text string describing the entity. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters excluding double quotes. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

## shutdown

<b>Syntax</b>	<b>[no] shutdown</b>
<b>Context</b>	<pre> config&gt;subscr-mgmt&gt;sub-ident-policy&gt;primary config&gt;subscr-mgmt&gt;sub-ident-policy&gt;secondary config&gt;subscr-mgmt&gt;sub-ident-policy&gt;tertiary config&gt;service&gt;vpls&gt;sap&gt;sub-sla-mgmt config&gt;service&gt;vpls&gt;gsmp config&gt;service&gt;vpls&gt;gsmp&gt;group config&gt;service&gt;vpls&gt;gsmp&gt;group&gt;neighbor config&gt;redundancy&gt;multi-chassis&gt;peer config&gt;redundancy&gt;multi-chassis&gt;peer&gt;mc-lag config&gt;redundancy&gt;multi-chassis&gt;peer&gt;sync config&gt;service&gt;ies&gt;sub-if&gt;dhcp config&gt;subscr-mgmt&gt;sub-mcac-policy config&gt;aaa&gt;route-downloader configure&gt;aaa&gt;diam-peer-pol&gt;peer </pre>
<b>Description</b>	<p>The <b>shutdown</b> command administratively disables the entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics. Many entities must be explicitly enabled using the <b>no shutdown</b> command.</p> <p>The <b>shutdown</b> command administratively disables an entity. The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they can be deleted.</p> <p>Unlike other commands and parameters where the default state is not indicated in the configuration file, <b>shutdown</b> and <b>no shutdown</b> are always indicated in system generated configuration files.</p> <p>The <b>no</b> form of the command puts an entity into the administratively enabled state.</p>
<b>Default</b>	no shutdown

## subscriber-mgmt

<b>Syntax</b>	<b>subscriber-mgmt</b>
<b>Context</b>	config
<b>Description</b>	<p>This command enables the context to configure subscriber management entities. A subscriber is uniquely identified by a subscriber identification string. Each subscriber can have several DHCP sessions active at any time. Each session is referred to as a subscriber host and is identified by its IP address and MAC address.</p> <p>All subscriber hosts belonging to the same subscriber are subject to the same hierarchical QoS (HQoS) processing. The HQoS processing is defined in the <b>sub-profile</b> (the subscriber profile). A sub-profile refers to an existing scheduler policy (configured in <b>the configure&gt;qos&gt;scheduler-policy</b> context) and offers the possibility to overrule the rate of individual schedulers within this policy.</p> <p>Because all subscriber hosts use the same scheduler policy instance, they must all reside on the same complex.</p>

---

## ANCP and GSMP Commands

### ancp

<b>Syntax</b>	<b>ancp</b>
<b>Context</b>	config>subscr-mgmt config>subscr-mgmt>sub-prof
<b>Description</b>	This command enables the context to configure Access Node Control Protocol (ANCP) parameters.

### ancp-policy

<b>Syntax</b>	<b>ancp-policy</b> <i>name</i>
<b>Context</b>	config>subscr-mgmt>ancp
<b>Description</b>	<p>This command creates an Access Node Control Protocol (ANCP) policy. The policy is associated with either the ANCP string (static case) or subscriber-profile (dynamic case) and defines the behavior of the hosts belonging to these profiles.</p> <p>ANCP polices control rates and subscribers based on port-up/port-down messages from the access node. When configured, the 7450 ESS should stop SHCV to a host that is part of a port defined to be down (by port-down message). When the node receives a port-up message for a port that was in port-down, state the node will initiate the SHCV process immediately to verify connectivity.</p> <p>When ANCP is used with Enhanced Subscriber Management, the ANCP string last associated with the subscriber will be used. All hosts of a subscriber will be updated with the new ANCP string.</p>
<b>Default</b>	No policies are defined.
<b>Parameters</b>	<i>name</i> — Configures the ANCP policy name.

### ancp-policy

<b>Syntax</b>	<b>ancp-policy</b> <i>name</i>
<b>Context</b>	config>subscr-mgmt>sub-prof>ancp
<b>Description</b>	This command specifies an existing Access Node Control Protocol (ANCP) policy to associate with the subscriber profile. The policy is associated with either the ANCP string (static case) or subscriber-profile (dynamic case) and defines the behavior of the hosts belonging to these profiles.
<b>Default</b>	No policies are defined.
<b>Parameters</b>	<i>name</i> — Specifies an existing ANCP policy name.



## ingress

<b>Syntax</b>	<b>ingress</b>
<b>Context</b>	config>subscr-mgmt>sla-prof>ingress config>subscr-mgmt>ancp>ancp-policy
<b>Description</b>	This command configures ingress ANCP policy parameters.

## rate-adjustment

<b>Syntax</b>	<b>rate-adjustment</b> <i>adjusted-percent</i> <b>no rate-adjustment</b>				
<b>Context</b>	config>subscr-mgmt>ancp>ancp-policy>ingress config>subscr-mgmt>ancp>ancp-policy>egress				
<b>Description</b>	This command configures a rate adjustment for the scheduler. The <b>rate-adjustment</b> command should be used when the rate returned by the DSLAM is calculated with different encapsulation than the 7450 ESS. The node will adjust the rate by the percent specified as:  $\text{DSLAM\_RATE} * \text{adjust-rate} / 100 \text{ — rate-reduction.}$ The <b>no</b> form of the command returns the default value.				
<b>Default</b>	none				
<b>Parameters</b>	<i>adjusted-percent</i> — Specifies a rate adjustment for the scheduler.  <table> <tr> <td><b>Values</b></td><td>1 — 200</td></tr> <tr> <td><b>Default</b></td><td>100</td></tr> </table>	<b>Values</b>	1 — 200	<b>Default</b>	100
<b>Values</b>	1 — 200				
<b>Default</b>	100				

## rate-reduction

<b>Syntax</b>	<b>rate-reduction</b> <i>kilobit-per-second</i> <b>no rate-reduction</b>
<b>Context</b>	config>subscr-mgmt>ancp>ancp-policy>ingress config>subscr-mgmt>ancp>ancp-policy>egress
<b>Description</b>	This command defines a constant rate reduction to the rate specified by the DSLAM. The <b>rate-reduction</b> command should be used if the node should adjust the rate to a value that is offset (for example by a fixed multicast dedicated bandwidth) compared to the total available on the DSLAM.  When set, the rate will be:  $\text{DSLAM\_RATE} * \text{adjust-rate} / 100 \text{ — rate-reduction}$
<b>Default</b>	none

## rate-monitor

<b>Syntax</b>	<b>rate-monitor</b> <i>kilobit-per-second</i> [ <b>alarm</b> ] <b>no rate-monitor</b>
<b>Context</b>	config>subscr-mgmt>ancp>ancp-policy>ingress config>subscr-mgmt>ancp>ancp-policy>egress
<b>Description</b>	This command configures the rate monitor level.
<b>Default</b>	none
<b>Parameters</b>	<i>kilobit-per-second</i> — Specifies the rate, in kilobits, below which the system will generate an event. <b>alarm</b> — When the monitored rate is below the configured value the system generates an alarm (trap) to the management system. The trap includes the rate as well as the ANCP policy name and the ANCP string.

## rate-modify

<b>Syntax</b>	<b>rate-modify</b> { <b>scheduler</b> <i>scheduler-name</i>   <b>arbiter</b> <i>arbiter-name</i> } <b>no rate-modify</b>
<b>Context</b>	config>subscr-mgmt>ancp>ancp-policy>ingress
<b>Description</b>	This command configures ingress rate modify scheduler parameters.
<b>Default</b>	none
<b>Parameters</b>	<b>scheduler</b> <i>scheduler-name</i> — Specifies a scheduler name. <b>arbiter</b> <i>arbiter-name</i> — Specifies an arbiter name

## egress

<b>Syntax</b>	<b>egress</b>
<b>Context</b>	config>subscr-mgmt>ancp>ancp-policy
<b>Description</b>	This command configures egress ANCP policy parameters.

## rate-modify

<b>Syntax</b>	<b>rate-modify</b> { <b>scheduler</b> <i>scheduler-name</i>   <b>arbiter</b> <i>arbiter-name</i> } <b>rate-modify</b> <b>agg-rate-limit</b> <b>no rate-modify</b>
<b>Context</b>	config>subscr-mgmt>ancp>ancp-policy>egress
<b>Description</b>	This command configures egress rate modify scheduler parameters.

<b>Default</b>	none
<b>Parameters</b>	<p><b>agg-rate-limit</b> — specifies that the maximum total rate for all subscriber egress queues for each subscriber associated with the policy.</p> <p><b>scheduler</b> <i>scheduler-name</i> — Specify a scheduler name.</p> <p><b>arbiter</b> <i>arbiter-name</i> — Specifies an arbiter name</p>

## port-down

<b>Syntax</b>	<b>[no] port-down</b>
<b>Context</b>	config>subscr-mgmt>ancp>ancp-policy
<b>Description</b>	This command specifies the number of GSMP portdown messages received in this ANCP session.

## disable-shcv

<b>Syntax</b>	<b>[no] disable-shcv [alarm] [hold-time seconds]</b>
<b>Context</b>	config>subscr-mgmt>ancp>ancp-policy>port-down
<b>Description</b>	<p>When this command is configured, the node will suspend SHCV for the hosts defined with this ANCP policy until the access node sends a port-up message. When the hold-time parameter is used, the node will suspend SHCV for the period of time defined. If the hold-time parameter is not defined the node will suspend SHCV until a port-up message is received.</p> <p>If the optional alarm flag is used the node should send a SHCV alarm before suspending.</p>
<b>Default</b>	no disable-shcv

## ancp-static-map

<b>Syntax</b>	<b>ancp-static-map</b>
<b>Context</b>	config>subscr-mgmt>ancp
<b>Description</b>	This command enables the context to configure a static ANCP name map.
<b>Default</b>	ancp-static-map

## entry

<b>Syntax</b>	<p><b>entry key</b> <i>ancp-string</i> <b>customer</b> <i>customer-id</i> <b>multi-service-site</b> <i>customer-site-name</i> <b>ancp-policy</b> <i>policy-name</i></p> <p><b>entry key</b> <i>ancp-string</i> <b>sap</b> <i>sap-id</i> <b>ancp-policy</b> <i>policy-name</i></p> <p><b>no entry key</b> <i>ancp-string</i> <b>customer</b> <i>customer-id</i> <b>multi-service-site</b> <i>customer-site-name</i></p> <p><b>no entry key</b> <i>ancp-string</i> <b>sap</b> <i>sap-id</i></p>
---------------	---

<b>Context</b>	config>subscr-mgmt>ancp>static-map
<b>Description</b>	<p>This command configures an ANCP name. When ANCP is configured to provide rate adaptation without the use of enhanced subscriber management, this command will define how to map an ANCP key (usually the circuit-id of the DSLAM port) to either a SAP and a scheduler name (when a Multi-Service Site (MSS) is not used) or a customer, site and scheduler name when MSS is used.</p> <p>Different ANCP names may be used with the same SAPs or customer ID/MSS combinations to allow schedulers within the policy to be mapped to the ANCP names. An ANCP string and SAP combination may reference only one ancp-policy. An ANCP string and customer and site-name combination may reference a single ancp-policy.</p>
<b>Default</b>	none
<b>Parameters</b>	<p><b>key</b> <i>ancp-string</i> — Specify the ASCII representation of the DSLAM circuit-id name.</p> <p><b>customer</b> <i>customer-id</i> — Specify the associated existing customer name.</p> <p><b>multi-service-site</b> <i>customer-site-name</i> — Specify the associated customer's configured MSS name.</p> <p><b>ancp-policy</b> <i>policy-name</i> — Specify an existing ANCP policy name.</p> <p><b>sap</b> <i>sap-id</i> — Specifies the physical port identifier portion of the SAP definition. See <a href="#">Common Service Commands on page 2062</a> for <i>sap-id</i> command syntax.</p>

## RADIUS Policy Commands

### radius-coa-port

<b>Syntax</b>	<b>radius-coa-port</b> {1647 1700 1812 3799} <b>no radius-coa-port</b>
<b>Context</b>	config>aaa
<b>Description</b>	This command configures the system-wide UDP port number that RADIUS is listening on for CoA and Disconnect messages.  The <b>no</b> form of the command resets the default UDP port to 3799.
<b>Default</b>	3799
<b>Parameters</b>	{1647 1700 1812 3799} — Specifies the udp port number for RADIUS CoA and Disconnect Messages.

### authentication-policy

<b>Syntax</b>	<b>authentication-policy</b> <i>name</i> [create] <b>no authentication-policy</b>
<b>Context</b>	config>subscr-mgmt
<b>Description</b>	This command creates the context to configure RADIUS server parameters for session authentication. The policies can be applied to an IES or a VPLS SAP.  The <b>no</b> form of the command removes the RADIUS server configuration for session authentication.  RADIUS servers can be configured for three different applications: <ol style="list-style-type: none"> <li>1. For authentication of dynamic Triple Play subscriber sessions, under config&gt;subscr-mgmt&gt;authentication-plcy</li> <li>2. For 802.1x port authentication, under config&gt;system&gt;security&gt;dot1x&gt;radius-plcy</li> <li>3. For CLI login users, under config&gt;system&gt;radius</li> </ol>
<b>Default</b>	none
<b>Parameters</b>	<i>name</i> — The name of the profile. The string is case sensitive and limited to 32 ASCII 7-bit printable characters with no spaces.

### pim-policy

<b>Syntax</b>	<b>pim-policy</b> <i>policy-name</i> <b>no pim-policy</b> <i>policy-name</i>
---------------	---

## RADIUS Policy Commands

<b>Context</b>	config>subscr-mgmt>sub-prof
<b>Description</b>	This command adds an existing PIM policy to this subscriber profile. The <b>no</b> form of the command removes the specified PIM policy from this subscriber profile.
<b>Default</b>	No PIM policy is added to a subscriber profile by default.
<b>Parameters</b>	<i>policy-name</i> — The name of the PIM policy. The string is case sensitive and limited to 32 ASCII 7-bit printable characters with no spaces.

## radius-accounting-policy

<b>Syntax</b>	<b>radius-accounting-policy</b> <i>name</i> <b>no radius-accounting-policy</b>
<b>Context</b>	config>subscr-mgmt config>subscr-mgmt>sub-prof
<b>Description</b>	This command specifies a subscriber RADIUS based accounting policy.
<b>Parameters</b>	<i>name</i> — The name of the policy. The string is case sensitive and limited to 32 ASCII 7-bit printable characters with no spaces.

## accept-authorization-change

<b>Syntax</b>	<b>[no] accept-authorization-change</b>
<b>Context</b>	config>subscr-mgmt>auth-policy
<b>Description</b>	This command specifies whether or not the system should handle the CoA messages initiated by the RADIUS server, and provide for mid-session interval changes of policies applicable to subscriber hosts.
<b>Default</b>	no accept-authorization-change

## accept-script-policy

<b>Syntax</b>	<b>accept-script-policy</b> <i>policy-name</i> <b>no accept-script-policy</b>
<b>Context</b>	config>subscr-mgmt>auth-policy
<b>Description</b>	This command configures a RADIUS script policy used to change the RADIUS attributes of the incoming Access-Accept messages.
<b>Parameters</b>	<i>policy-name</i> — Configures a Python script policy to modify Access-Accept messages.

## access-loop-options

<b>Syntax</b>	<b>[no] access-loop-options</b>
<b>Context</b>	config>subscr-mgmt>auth-plcy>include-radius-attribute config>subscr-mgmt>acct-plcy>include-radius-attribute
<b>Description</b>	<p>This command enables inclusion of access loop information: Broadband Forum (BBF) access loop characteristics, DSL line state and DSL type. The BBF access loop characteristics are returned as BBF specific RADIUS attributes where DSL line state and DSL type are returned as Alcatel-Lucent specific RADIUS VSA's.</p> <p>Information obtained via the ANCP protocol has precedence over information received in PPPoE Vendor Specific BBF tags or DHCP Vendor Specific BBF Options.</p> <p>If ANCP is utilized and interim accounting update is enabled, any "Port Up" event from GSMP will initiate in an interim update. "Port Up" messages can include information such as an update on the current subscriber actual-upstream-speed. The next interim accounting message will be from "port up" triggering point.</p>
<b>Default</b>	no access-loop-options

## host-accounting

<b>Syntax</b>	<b>[no] host-accounting [interim-update]</b>
<b>Context</b>	config>subscr-mgmt>acct-plcy
<b>Description</b>	<p>This command enables per host accounting mode. In host accounting mode, the acct-session-id is generated per host. This acct-session-id is uniformly included in all accounting messages (START/INTERIM-UPDATE/STOP) and it can be included in RADIUS Access-Request message.</p> <p>Accounting counters are based on the queue counters and as such are aggregated for all host sharing the queues within an sla-profile instance (non HSMDA) or a subscriber (HSMDA). CoA and LI is supported based on the acct-session-id of the host.</p>
<b>Default</b>	no host-accounting
<b>Parameters</b>	<b>interim-update</b> — Without this keyword only START and STOP accounting messages are generated when the host is established/terminated. This is equivalent to a time-based accounting where only the duration of the session is required.

## include-radius-attribute

<b>Syntax</b>	<b>[no] include-radius-attribute</b>
<b>Context</b>	config>subscr-mgmt>auth-plcy config>subscr-mgmt>acct-plcy
<b>Description</b>	This command enables the context to specify the RADIUS parameters that the system should include into RADIUS authentication-request messages.

## all-authorized-session-addresses

<b>Syntax</b>	<b>[no] all-authorized-session-addresses</b>
<b>Context</b>	config>subscr-mgmt>acct-plcy>include-radius-attribute
<b>Description</b>	<p>Applicable for session-accounting mode only.</p> <p>With this flag enabled, all IP address attributes explicitly enabled to be included are the following:</p> <ul style="list-style-type: none"> <li>• delegated-ipv6-prefix</li> <li>• framed-ip-address</li> <li>• framed-ip-netmask</li> <li>• framed-ipv6-prefix</li> <li>• ipv6-address</li> </ul> <p>These are included if the corresponding addresses or prefixes are authorized (via access-accept or ludb) and independent if they are used or not.</p>
<b>Default</b>	no all-authorized-session-addresses

## calling-station-id

<b>Syntax</b>	<b>calling-station-id</b> <b>calling-station-id {mac   remote-id   sap-id   sap-string}</b> <b>no calling-station-id</b>
<b>Context</b>	config>service>ies>if>sap config>service>ies>sub-if>grp-if>sap config>service>vpls>sap config>subscr-mgmt>auth-plcy>include-radius-attribute config>subscr-mgmt>acct-plcy>include>include-radius-attribute
<b>Description</b>	This command enables the inclusion of the calling-station-id attribute in RADIUS authentication requests and RADIUS accounting messages. The value inserted is set at the SAP level. If no <b>calling-station-id</b> value is set at the SAP level, the <b>calling-station-id</b> attribute will not be sent.
<b>Default</b>	no calling-station-id
<b>Parameters</b>	<p><b>mac</b> — Specifies that the mac-address will be sent.</p> <p><b>remote-id</b> — Specifies that the remote-id will be sent.</p> <p><b>sap-id</b> — Specifies that the sap-id will be sent.</p> <p><b>sap-string</b> — Specifies that the value is the inserted value set at the SAP level. If no <b>calling-station-id</b> value is set at the SAP level, the <b>calling-station-id</b> attribute will not be sent.</p>

## access-loop-options

<b>Syntax</b>	<b>[no] access-loop-options</b>
---------------	---------------------------------



**Context** config>subscr-mgmt>auth-plcy>include-radius-attribute  
config>subscr-mgmt>acct-plcy>include-radius-attribute

**Description** This command enables inclusion of access loop information: Broadband Forum (BBF) access loop characteristics, DSL line state and DSL type. The BBF access loop characteristics are returned as BBF specific RADIUS attributes where DSL line state and DSL type are returned as Alcatel-Lucent specific RADIUS VSA's.

Information obtained via the ANCP protocol has precedence over information received in PPPoE Vendor Specific BBF tags or DHCP Vendor Specific BBF Options.

## acct-session-id

**Syntax** [no] acct-session-id

**Context** configure>subscr-mgmt>auth-plcy>include-radius-attribute

**Description** The **acct-session-id** attribute for each subscriber host will be generated at the very beginning of the session initiation. This command will enable or disable sending this attribute to the RADIUS server in the Access-Request messages regardless of whether the accounting is enabled or not. The **acct-session-id** attribute can be used to address the subscriber hosts from the RADIUS server in the CoA Request.

The **acct-session-id** attribute will be unique per subscriber host network wide. It is a 22bytes long field comprised of the system MAC address along with the creation time and a sequence number in a hex format.

**Default** Disabled

## circuit-id

**Syntax** [no] circuit-id

**Context** config>subscr-mgmt>auth-policy>include-radius-attribute  
config>subscr-mgmt>acct-plcy>include-radius-attribute

**Description** This command enables the generation of the agent-circuit-id for RADIUS.

## detailed-acct-attributes

**Syntax** [no] detailed-acct-attributes

**Context** config>subscr-mgmt>auth-plcy>include-radius-attribute

**Description** This command enables detailed reporting of per queue and per policer octet and packet counters using RADIUS VSAs. Enabled by default. It can be enabled simultaneously with aggregate counters (std-acct-attributes).

The **no** form of the command excludes the detailed counter VSAs from the RADIUS accounting messages.

## RADIUS Policy Commands

**Default** detailed-acct-attributes

### dhcp-options

**Syntax** [no] dhcp-options

**Context** config>subscr-mgmt>auth-plcy>include-radius-attribute

**Description** This command enables insertion of RADIUS VSA containing all dhcp-options from dhcp-discover (or dhcp-request) message. The VSA contains all dhcp-options in a form of the string. If required (the total length of all dhcp-options exceeds 255B), multiple VSAs are included.

**Default** no dhcp-options

### dhcp6-options

**Syntax** [no] dhcp6-options

**Context** configure>subscr-mgmt>auth-policy>include

**Description** This command will copy DHCPv6 options from received DHCPv6 messages on ingress access and pass them to the RADIUS server in Accept-Request. The messages will be carried in the ALU VSA Alc-ToServer-Dhcp6-Options.

**Default** no dhcp6-options

### dhcp-vendor-class-id

**Syntax** [no] dhcp-vendor-class-id

**Context** config>subscr-mgmt>auth-plcy>include-radius-attribute

**Description** This command includes the “[26-6527-36] Alc-DHCP-Vendor-Class-Id” attribute in RADIUS accounting messages. The content of the DHCP Vendor-Class-Identifier option (60) is mapped in this attribute.

**Default** no dhcp-vendor-class-id

### framed-ip-addr

**Syntax** [no] framed-ip-addr

**Context** config>subscr-mgmt>acct-plcy>include-radius-attribute

**Description** This command enables the inclusion of the framed-ip-addr attribute.

## framed-ip-netmask

<b>Syntax</b>	<b>[no] framed-ip-netmask</b>
<b>Context</b>	config>subscr-mgmt>acct-plcy>include-radius-attribute
<b>Description</b>	This command enables the inclusion of the framed-ip-netmask attribute.

## mac-address

<b>Syntax</b>	<b>[no] mac-address</b> config>subscr-mgmt>auth-policy>include-radius-attribute config>subscr-mgmt>acct-plcy>include-radius-attribute
<b>Description</b>	This command enables the generation of the client MAC address RADIUS attribute.

## nas-identifier

<b>Syntax</b>	<b>[no] nas-identifier</b>
<b>Context</b>	config>subscr-mgmt>auth-policy>include-radius-attribute config>subscr-mgmt>acct-plcy>include-radius-attribute
<b>Description</b>	This command enables the generation of the nas-identifier RADIUS attribute.

## nas-port

<b>Syntax</b>	<b>[no] nas-port <i>bit-specification binary-spec</i></b>	
<b>Context</b>	config>subscr-mgmt>auth-policy>include-radius-attribute config>subscr-mgmt>acct-plcy>include-radius-attribute	
<b>Description</b>	This command enables the generation of the nas-port RADIUS attribute. You enter decimal representation of a 32-bit string that indicates your port information. This 32-bit string can be compiled based on different information from the port (data types). By using syntax number-of-bits data-type you indicate how many bits from the 32 bits are used for the specific data type. These data types can be combined up to 32 bits in total. In between the different data types 0's and/or 1's as bits can be added.  The <b>no</b> form of this command disables your nas-port configuration.	
<b>Parameters</b>	<i>bit-specification binary-spec</i> — Specifies the NAS-Port attribute	
	<b>Values</b>	binary-spec      <bit-specification> <binary-spec> bit-specification    0   1   <bit-origin> bit-origin            *<number-of-bits><origin> number-of-bits       1 — 32 origin                o   i   s   m   p outer VLAN ID i      inner VLAN ID

s slot number  
m MDA number  
p port number or lag-id

Sample

```
*12o*12i00*2s*2m*2p => 0000 0000 0000 iiii iiii iiii 00ss mmpp
If outer vlan = 0 & inner vlan = 1 & slot = 3 & mda = 1 & port = 1
=> 0000 0000 0000 0000 0000 0001 0011 0101 => nas-port = 309
```

nas-port-id

Syntax	[no] nas-port-id [prefix-string <i>string</i> ] [suffix <i>suffix-option</i> ]
Context	config>subscr-mgmt>auth-policy>include-radius-attribute config>subscr-mgmt>acct-plcy>include-radius-attribute
Description	This command enables the generation of the nas-port-id RADIUS attribute. Optionally, the value of this attribute (the SAP-id) can be prefixed by a fixed string and suffixed by the circuit-id or the remote-id of the client connection. If a suffix is configured, but no corresponding data is available, the suffix used will be 0/0/0/0/0.
Parameters	<b>prefix-string</b> <i>string</i> — Specifies that a user configurable string will be added to the RADIUS NAS port attribute, up to 8 characters in length. <b>suffix</b> <i>suffix-option</i> — Specifies the suffix type to be added to the RADIUS NAS port attribute. <b>Values</b> circuit-id, remote-id

nas-port-type

Syntax	nas-port-type nas-port-type [0..255] no nas-port-type
Context	config>subscr-mgmt>auth-plcy>include-radius-attribute config>subscr-mgmt>acct-plcy>include-radius-attribute
Description	This command enables the generation of the nas-port-type RADIUS attribute. If set to <b>nas-port-type</b> , the following will be sent: values: 32 (null-encap), 33 (dot1q), 34 (qinq), 15 (DHCP hosts). The <b>nas-port-type</b> can also be set as a specified value, with an integer from 0 to 255. The <b>no</b> form of the command reverts to the default.
Default	no nas-port-type
Parameters	<b>0 — 255</b> — Specifies an enumerated integer that specifies the value that will be put in the RADIUS nas-port-type attribute.

## pppoe-service-name

<b>Syntax</b>	<b>[no] pppoe-service-name</b>
<b>Context</b>	config>subscr-mgmt>auth-policy>include-radius-attribute config>subscr-mgmt>acct-plcy>include-radius-attribute
<b>Description</b>	This command enables the generation of the pppoe-service-name RADIUS attribute.

## remote-id

<b>Syntax</b>	<b>[no] remote-id</b>
<b>Context</b>	config>subscr-mgmt>auth-policy>include-radius-attribute config>subscr-mgmt>acct-plcy>include-radius-attribute
<b>Description</b>	This command enables the generation of the agent-remote-id for RADIUS.

## sap-session-index

<b>Syntax</b>	<b>[no] sap-session-index</b>
<b>Context</b>	config>subscr-mgmt>acct-plcy>include-radius-attribute
<b>Description</b>	This command enables the generation of the per-SAP unique session index. The <b>no</b> form of the command excludes <b>sap-session-index</b> attributes.

## tunnel-server-attrs

<b>Syntax</b>	<b>[no] tunnel-server-attrs</b>
<b>Context</b>	config>subscr-mgmt>auth-policy>include-radius-attribute
<b>Description</b>	This command includes tunnel-server attribute.

## sla-profile

<b>Syntax</b>	<b>[no] sla-profile</b>
<b>Context</b>	config>subscr-mgmt>acct-plcy>include-radius-attribute
<b>Description</b>	This command specifies that SLA profile attributes should be included into RADIUS accounting messages.

### std-acct-attributes

<b>Syntax</b>	<b>[no] std-acct-attributes</b>
<b>Context</b>	config>subscr-mgmt>acct-plcy>include-radius-attribute
<b>Description</b>	This command enables reporting of aggregated forwarded octet and packet counters using standard Radius attributes. Disabled by default. It can be enabled simultaneously with detailed per queue/policer counters (detailed-acct-attributes).
<b>Default</b>	no std-acct-attributes

### sub-profile

<b>Syntax</b>	<b>[no] sub-profile</b>
<b>Context</b>	config>subscr-mgmt>acct-plcy>include-radius-attribute
<b>Description</b>	This command specifies that subscriber profile attributes should be included into RADIUS accounting messages.

### subscriber-id

<b>Syntax</b>	<b>[no] subscriber-id</b>
<b>Context</b>	config>subscr-mgmt>acct-plcy>include-radius-attribute
<b>Description</b>	This command specifies that subscriber ID attributes should be included into RADIUS accounting messages.

### user-name

<b>Syntax</b>	<b>[no] user-name</b>
<b>Context</b>	config>subscr-mgmt>acct-plcy>include-radius-attribute
<b>Description</b>	This command enables the inclusion of the user-name attribute. The <b>no</b> form of the command disables the inclusion of the user-name attribute.
<b>Default</b>	no user-name

### v6-aggregate-stats

<b>Syntax</b>	<b>[no] v6-aggregate-stats</b>
<b>Context</b>	config>subscr-mgmt>acct-plcy>include-radius-attribute

<b>Description</b>	This command enables reporting of IPv6 aggregated forwarded octet and packet counters using RADIUS VSAs. Disabled by default. It requires <b>stat-mode v4-v6</b> for policers and queues for which the IPv6 aggregate forwarded packets should be counted.
<b>Default</b>	no v6-aggregate-stats

## wifi-rssi

<b>Syntax</b>	<b>[no] wifi-rssi</b>
<b>Context</b>	config>subscr-mgmt>acct-plcy>include-radius-attribute
<b>Description</b>	This command enables the inclusion of the 802.11 Received Signal Strength Indication attribute.

## password

<b>Syntax</b>	<b>password <i>password</i> [hash   hash2]</b> <b>no password</b>
<b>Context</b>	config>subscr-mgmt>auth-policy
<b>Description</b>	<p>This command sets a password that is sent with <b>user-name</b> in every RADIUS authentication request sent to the RADIUS server upon receipt of DHCP discover or request messages. If no password is configured, no password AVP will be sent.</p> <p>The <b>no</b> form of the command reverts to the default value.</p>
<b>Default</b>	none
<b>Parameters</b>	<p><i>password</i> — A text string containing the password. Allowed values are any string up to 64 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.</p> <p><b>hash</b> — Specifies the key is entered in an encrypted form. If the <b>hash</b> parameter is not used, the key is assumed to be in a non-encrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the <b>hash</b> parameter specified.</p> <p><b>hash2</b> — Specifies the key is entered in a more complex encrypted form. If the <b>hash2</b> parameter is not used, the less encrypted <b>hash</b> form is assumed.</p>

## password

<b>Syntax</b>	<b>password <i>password</i> [hash   hash2]</b> <b>no password</b>
<b>Context</b>	config>subscr-mgmt>diam-appl-plcy>nasreq
<b>Description</b>	This command sets a password that is sent with <b>user-name</b> in every RADIUS authentication request sent to the RADIUS server upon receipt of DHCP discover or request messages. If no password is provided, an empty password will be sent.

The **no** form of the command reverts to the default value.

<b>Default</b>	no password
<b>Parameters</b>	<p><i>password</i> — A text string containing the password. Allowed values are any string up to 64 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.</p> <p><b>hash</b> — Specifies the key is entered in an encrypted form. If the <b>hash</b> or <b>hash2</b> parameter is not used, the key is assumed to be in a non-encrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the <b>hash</b> or <b>hash2</b> parameter specified.</p> <p><b>hash2</b> — Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, this means that a <b>hash2</b> encrypted variable cannot be copied and pasted. If the <b>hash</b> or <b>hash2</b> parameter is not used, the key is assumed to be in a non-encrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the <b>hash</b> or <b>hash2</b> parameter specified.</p>

### ppp-user-name

<b>Syntax</b>	<p><b>ppp-user-name append</b> <i>domain-name</i></p> <p><b>ppp-user-name default-domain</b> <i>domain-name</i></p> <p><b>ppp-user-name replace</b> <i>domain-name</i></p> <p><b>ppp-user-name strip</b></p> <p><b>no ppp-user-name</b></p>
<b>Context</b>	config>subscr-mgmt>auth-plcy
<b>Description</b>	<p>This command configures the password that is sent with the User-Name in Diameter NASREQ AA-Requests for IPoE hosts.</p> <p>When no password is configured, an empty password will be sent.</p>
<b>Default</b>	no ppp-user-name
<b>Parameters</b>	<p><b>append</b> <i>domain-name</i> — A string specified by <code>tmnxSubAuthPlcyPppDomain</code>, preceded with a '@', is appended to the PAP/CHAP user name.</p> <p><b>default-domain</b> <i>domain-name</i> — The same action is performed as with <code>appendDomain</code>, but only if the PAP/CHAP user name does not already contain a domain name.</p> <p><b>replace</b> <i>domain-name</i> — All characters after a '@' delimiter are replaced with the string specified by <code>tmnxSubAuthPlcyPppDomain</code>.</p> <p><b>strip</b> — Any '@' character and all subsequent characters are removed from the PAP/CHAP user name.</p>

### pppoe-access-method

<b>Syntax</b>	<p><b>pppoe-access-method</b> {<b>none</b>   <b>padi</b>   <b>pap-chap</b>}</p> <p><b>no pppoe-access-method</b></p>
<b>Context</b>	config>subscr-mgmt>auth-plcy



<b>Description</b>	This command indicates the authentication method used towards the RADIUS server in case the policy is used for PPPoE.
<b>Parameters</b>	<p><b>none</b> — Indicates that the client will be authenticated by the local user database defined under the group interface and not through RADIUS.</p> <p><b>padi</b> — Indicates that the client will be authenticated by RADIUS as soon as the PADI packet comes in (there is no PPP authentication done in the session in this case).</p> <p><b>pap-chap</b> — Indicates that the RADIUS authentication of the client will be delayed until the authentication protocol phase in the PPP session (PAP or CHAP) and authentication will be performed with the user name and PAP password / CHAP response supplied by the client.</p>

## queue-instance-accounting

<b>Syntax</b>	<b>queue-instance-accounting [interim-update]</b> <b>no queue-instance-accounting</b>
<b>Context</b>	config>subscr-mgmt>acct-plcy
<b>Description</b>	<p>This command enables per queue-instance-accounting. A stream of accounting messages (START/INTERIM-UPDATE/STOP) is generated per queuing instance. A queuing instance is equivalent to an sla-profile instance on non HSMDA based hardware and to subscriber on HSMDA based hardware. Accounting session id is generated per queuing instance and this accounting session id CANNOT be included in RADIUS Access-Request message. Queue instance counters represent volume based aggregation for all hosts sharing the queuing instance.</p> <p>CoA and LI is supported based on the acct-session-id of the queuing instance.</p>
<b>Default</b>	interim-update
<b>Parameters</b>	<b>interim-update</b> — specifies whether accounting messages are sent for the queue-instance. The queue-instance is the subscriber on High Scale MDA (HSMDA), or the SLA profile instance otherwise.

## radius-authentication-server

<b>Syntax</b>	<b>radius-authentication-server</b>
<b>Context</b>	config>subscr-mgmt>acct-plcy
<b>Description</b>	This command creates the context for defining RADIUS authentication server attributes under a given session authentication policy.

## access-algorithm

<b>Syntax</b>	<b>access-algorithm {direct   round-robin}</b> <b>no access-algorithm</b>
<b>Context</b>	config>subscr-mgmt>auth-plcy-srvr

config>subscr-mgmt>acct-plcy>server

- Description** This command configures the algorithm used to access the list of configured RADIUS servers.
- Parameters** **direct** — Specifies that the first server will be used as primary server for all requests, the second as secondary and so on.
- round-robin** — Specifies that the first server will be used as primary server for the first request, the second server as primary for the second request, and so on. If the router gets to the end of the list, it starts again with the first server.

### fallback-action

- Syntax** **fallback-action accept**  
**fallback-action user-db** *local-user-db-name*  
**no fallback-action**
- Context** config>subscr-mgmt>auth-plcy-srvr  
config>subscr-mgmt>auth-plcy
- Description** This command configures the action when no RADIUS server is available.  
The no form of the command removes the action from the configuration.
- Default** no fallback-action

### hold-down-time

- Syntax** **hold-down-time seconds**  
**no hold-down-time**
- Context** config>subscr-mgmt>auth-plcy>radius-auth-server
- Description** This command determines the interval during which no new communication attempts will be made to a RADIUS server that is marked **down** to prevent immediately overloading the server when it is starting up. The only exception is when all servers in the authentication policy are marked **down**; in that case they will all be used again to prevent failures on new client connections.
- Default** 30
- Parameters** *seconds* — Specifies the hold time before re-using a RADIUS server that was down.
- Values** 30 — 900

### router

- Syntax** **router** *router-instance*  
**router** *service-name*  
**no router**
- Context** config>subscr-mgmt>auth-plcy-srvr

```
config>subscr-mgmt>acct-plcy>server
```

**Description** This command specifies the virtual router instance applicable for the set of configured RADIUS servers. This value cannot be changed once a RADIUS server is configured for this policy. When the value is zero, both base and management router instances are matched.

**Parameters** *router-instance* — Specifies the virtual router instance.

<b>Values</b>	router-name:	Base, management
	service-id:	1 — 2147483647
	service-name:	Specifies the service name up to 64 characters in length.

## retry

**Syntax** **retry** *count*  
**no** **retry**

**Context** config>subscr-mgmt>auth-plcy-srvr  
config>subscr-mgmt>acct-plcy>server

**Description** This command configures the number of times the router attempts to contact the RADIUS server for authentication, if not successful the first time.

The **no** form of the command reverts to the default value.

**Default** 3

**Parameters** *count* — The retry count.

**Values** 1 — 10

## radius-server-policy

**Syntax** **radius-server-policy** *radius-server-policy-name*  
**no** **radius-server-policy**

**Context** config>subscr-mgmt>auth-plcy  
config>subscr-mgmt>acct-plcy

**Description** This command references an existing radius-server-policy (available under the config>aaa context) for use in subscriber management authentication and accounting.

When configured in an authentication-policy, following CLI commands are ignored in the policy to avoid conflicts:

- all commands in the radius-authentication-server context
- accept-authorization-change
- coa-script-policy
- accept-script-policy
- request-script-policy

## RADIUS Policy Commands

When configured in a radius-accounting-policy, following CLI commands are ignored in the policy to avoid conflicts:

- all commands in the radius-accounting-server context
- acct-request-script-policy

The **no** form of the command removes the radius-server-policy reference from the configuration

<b>Default</b>	no radius-server-policy
<b>Parameters</b>	<i>radius-server-policy-name</i> — Specifies the RADIUS server policy.

## server

<b>Syntax</b>	<b>server</b> <i>server-index</i> <b>address</b> <i>ip-address</i> <b>secret</b> <i>key</i> [ <b>hash</b>   <b>hash2</b> ] [ <b>port</b> <i>port-num</i> ] [ <b>coa-only</b> ] [ <b>pending-requests-limit</b> <i>limit</i> ] <b>no server</b> <i>index</i>
<b>Context</b>	config> subscr-mgmt>auth-policy>radius-auth-server config>subscr-mgmt>acct-plcy>server
<b>Description</b>	<p>This command adds a RADIUS server and configures the RADIUS server IP address, index, and key values.</p> <p>Up to sixteen RADIUS servers can be configured at any one time in a RADIUS authentication policy. Only five can be used for authentication, all other servers should be configured as coa-only servers. In a RADIUS accounting policy, up to five RADIUS servers can be configured. RADIUS servers are accessed in order from lowest to highest index for authentication or accounting requests until a response from a server is received. A higher indexed server is only queried if no response is received from a lower indexed server (which implies that the server is not available). If a response from a server is received, no other RADIUS servers are queried.</p> <p>The <b>no</b> form of the command removes the server from the configuration.</p>
<b>Default</b>	No RADIUS servers are configured.
<b>Parameters</b>	<p><i>server-index</i> — The index for the RADIUS server. The index determines the sequence in which the servers are queried for authentication requests. Servers are queried in order from lowest to highest index.</p> <p><b>Values</b>      1 — 16 (a maximum of 5 authentication servers)</p> <p><b>address</b> <i>ip-address</i> — The IP address of the RADIUS server. Two RADIUS servers cannot have the same IP address. An error message is generated if the server address is a duplicate.</p> <p><b>secret</b> <i>key</i> — The secret key to access the RADIUS server. This secret key must match the password on the RADIUS server.</p> <p><b>Values</b>      secret-key: Up to 20 characters in length.                  hash-key: Up to 33 characters in length.                  hash2-ke: Up to 55 characters in length.</p> <p><b>hash</b> — Specifies the key is entered in an encrypted form. If the <b>hash</b> parameter is not used, the key is assumed to be in a non-encrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the <b>hash</b> parameter specified.</p>

**hash2** — Specifies the key is entered in a more complex encrypted form. If the **hash2** parameter is not used, the less encrypted **hash** form is assumed.

**port** *port-num* — Specifies the UDP port number on which to contact the RADIUS server for authentication.

**Values** 1 — 65535

**coa-only** — Specifies Change-of-Authorization Messages only. Servers that are marked with the coa-only flag will not be used for authentication, but they will be able to accept RADIUS CoA messages, independent of the accept-authorization-change setting in the authentication policy.

For authentication purposes, the maximum number of servers is 5. All other servers may only be used as coa-only servers.

**pending-requests-limit** *limit* — Specifies the maximum number of outstanding RADIUS authentication requests for this authentication server.

**Default** The default value when not configured is 4096.

**Values** 1 — 4096

## hold-down-time

<b>Syntax</b>	<b>hold-down-time</b> <b>no hold-down-time</b>
<b>Context</b>	config>aaa>radius-server-policy>servers
<b>Description</b>	This command determines the interval during which no new communication attempts will be made to a RADIUS server that is marked down to prevent immediately overloading the server when it is starting up. The only exception is when all servers in the authentication policy are marked down; in that case, they will all be used again to prevent failures on new client connections.
<b>Default</b>	30s
<b>Parameters</b>	<p><i>days</i> — Specifies the hold time in days before re-using a RADIUS server that was down.</p> <p><b>Values</b> 0 — 3650</p> <p><i>hours</i> — Specifies the hold time in hours before re-using a RADIUS server that was down.</p> <p><b>Values</b> 0 — 23</p> <p><i>minutes</i> — Specifies the hold time in minutes before re-using a RADIUS server that was down.</p> <p><b>Values</b> 0 — 59</p> <p><i>seconds</i> — Specifies the hold time in seconds before re-using a RADIUS server that was down.</p>

## source-address

<b>Syntax</b>	<b>source-address</b> <i>ip-address</i> <b>no source-address</b>
<b>Context</b>	config>subscr-mgmt>auth-plcy-srvr

config>subscr-mgmt>acct-plcy>server

<b>Description</b>	<p>This command configures the source address of the RADIUS packet.</p> <p>The system IP address must be configured in order for the RADIUS client to work. See Configuring a System Interface in the 7450 ESS OS Router Configuration Guide. Note that the system IP address must only be configured if the source-address is not specified. When the <b>no source-address</b> command is executed, the source address is determined at the moment the request is sent. This address is also used in the nas-ip-address attribute: over there it is set to the system IP address if <b>no source-address</b> was given.</p> <p>The <b>no</b> form of the command reverts to the default value.</p>
<b>Default</b>	System IP address
<b>Parameters</b>	<p><i>ip-address</i> — The IP prefix for the IP match criterion in dotted decimal notation.</p> <p><b>Values</b> 0.0.0.0 - 255.255.255.255</p>

### timeout

<b>Syntax</b>	<p><b>timeout seconds</b></p> <p><b>no timeout</b></p>
<b>Context</b>	<p>config&gt;subscr-mgmt&gt;auth-plcy-srvr</p> <p>config&gt;subscr-mgmt&gt;acct-plcy&gt;server</p> <p>This command configures the number of seconds the router waits for a response from a RADIUS server.</p> <p>The <b>no</b> form of the command reverts to the default value.</p>
<b>Default</b>	3 seconds
<b>Parameters</b>	<p><i>seconds</i> — The number of seconds the router waits for a response from a RADIUS server, expressed as a decimal integer.</p> <p><b>Values</b> 1 — 90</p>

### session-accounting

<b>Syntax</b>	<p><b>session-accounting [interim-update] [host-update]</b></p> <p><b>no session-accounting</b></p>
<b>Context</b>	config>subscr-mgmt>acct-plcy
<b>Description</b>	<p>This command enables per session accounting mode. In per session accounting mode, the acct-session-id is generated per session. This acct-session-id is uniformly included in all accounting messages (START/INTERIM-UPDTATE/STOP) and it can be included in RADIUS Access-Request message.</p> <p>This accounting mode of operation can be used only in PPPoE environment with dual-stack host in which case both hosts (IPv4 and IPv6) are considered part of the same session. In addition to regular interim-updates, <i>triggered</i> interim-updates are sent by a host joining or leaving the session.</p>

When an IPv4/v6 address is allocated, or released from a dual-stack host, a triggered interim-update message is immediately sent. This triggered interim-update message reflects the change in the IP address. The triggered interim-update has no effect on the interval at which the regular interim updates are scheduled.

Accounting counters are based on the queue counters and as such are aggregated for all host sharing the queues within an sla-profile instance (non HSMDA) or a subscriber (HSMDA).

CoA and LI is supported based on the acct-session-id of the session.

<b>Default</b>	no session-accounting
<b>Parameters</b>	<p><b>interim-update</b> — Without this keyword only START and STOP accounting messages are generated when the session is established/terminated. This is equivalent to a time-based accounting where only the duration of the session is required.</p> <p><b>host-update</b> — This keyword indicates that host updates messages are sent. INTERIM-UPDATE messages can be generated (volume based accounting) by selecting this keyword..</p>

## session-id-format

<b>Syntax</b>	<b>session-id-format {description   number}</b> <b>no session-id-format</b>
<b>Context</b>	config>subscr-mgmt>acct-plcy
<b>Description</b>	This command specifies the format for the acct-session-id attribute used in RADIUS accounting requests.
<b>Parameters</b>	<p><b>description</b> — Specifies to use a string containing following information &lt;subscriber&gt;@&lt;sap-id&gt;@&lt;SLA-profile&gt;_&lt;creation-time&gt;.</p> <p><b>number</b> — Specifies to use a unique number generated by the OS to identify a given session.</p>

## update-interval

<b>Syntax</b>	<b>update-interval <i>minutes</i></b> <b>no update-interval</b>
<b>Context</b>	config>subscr-mgmt>acct-plcy
<b>Description</b>	<p>This command specifies the interval at which accounting data of subscriber hosts will be updated in a RADIUS Accounting Interim-Update message. Requires interim-update to be enabled when specifying the accounting mode in the radius accounting policy.</p> <p>A RADIUS specified interim interval (attribute [85] Acct-Interim-Interval) overrides the CLI configured value.</p>
<b>Parameters</b>	<p><i>minutes</i> — Specifies the interval, in minutes, at which accounting data of subscriber hosts will be updated.</p> <p><b>Values</b>      5 — 259200</p>

## update-interval-jitter

<b>Syntax</b>	<b>update-interval-jitter absolute seconds</b> <b>no update-interval-jitter</b>
<b>Context</b>	config>subscr-mgmt>acct-plcy
<b>Description</b>	<p>This command specifies the absolute maximum random delay introduced on the update interval between two accounting interim update messages. The effective maximum random delay value is the minimum of the configured absolute jitter value and 10% of the configured update-interval.</p> <p>A value of zero will send the accounting interim update message without introducing an additional random delay.</p> <p>The <b>no</b> form of the command sets the default to 10% of the configured update-interval.</p>
<b>Default</b>	no update-interval-jitter This corresponds with 10% of the configured update-interval
<b>Parameters</b>	<b>absolute seconds</b> — specifies the absolute maximum jitter value in seconds.
	<b>Values</b> 0 — 36000

## re-authentication

<b>Syntax</b>	<b>[no] re-authentication</b>
<b>Context</b>	config>subscr-mgmt>auth-policy
<b>Description</b>	<p>This command enables authentication process at every DHCP address lease renewal s only if RADIUS did not reply any special attributes (for example, authentication only, no authorization).</p> <p>The <b>no</b> form of the command reverts to the default value.</p>
<b>Default</b>	disabled

## request-script-policy

<b>Syntax</b>	<b>request-script-policy policy-name</b> <b>no request-script-policy</b>
<b>Context</b>	config>subscr-mgmt>auth-policy
<b>Description</b>	This command specifies the RADIUS script policy used to change the RADIUS attributes of the outgoing Access-Request messages.
<b>Default</b>	none
<b>Parameters</b>	<i>policy-name</i> — Configures a Python script policy to modify Access-Request messages.



## send-acct-stop-on-fail

<b>Syntax</b>	<b>send-acct-stop-on-fail</b> {[on-request-failure] [on-reject] [on-accept-failure]} <b>no send-acct-stop-on-fail</b>
<b>Context</b>	config>subscr-mgmt>auth-policy
<b>Description</b>	<p>This command activates the reporting of RADIUS authentication failures of a PPPoE session to a RADIUS accounting server with an Accounting Stop message.</p> <p>Three failure categories can be enabled separately:</p> <ul style="list-style-type: none"> <li>• <b>on-request-failure</b>: All failure conditions between the sending of an Access-Request and the reception of an Access-Accept or Access-Reject.</li> <li>• <b>on-reject</b>: When an Access-Reject is received</li> <li>• <b>on-accept-failure</b>: All failure conditions that appear after receiving an Access-Accept and before successful instantiation of the host or session.</li> </ul> <p>The RADIUS accounting policy to be used for sending the Accounting Stop messages must be obtained prior to RADIUS authentication via local user database pre-authentication.</p>
<b>Default</b>	no send-acct-stop-on-fail

## user-name-format

<b>Syntax</b>	<b>user-name-format</b> <i>format</i> [ <b>mac-format</b> <i>mac-format</i> ] <b>user-name-format</b> <i>format</i> <b>append</b> [ <i>domain-name</i> ] [ <b>mac-format</b> <i>mac-format</i> ] <b>user-name-format</b> <i>format</i> <b>append</b> <i>domain-name</i> <b>user-name-format</b> <i>format</i> <b>default-domain</b> <i>domain-name</i> [ <b>mac-format</b> <i>mac-format</i> ] <b>user-name-format</b> <i>format</i> <b>replace</b> <i>domain-name</i> [ <b>mac-format</b> <i>mac-format</i> ] <b>user-name-format</b> <i>format</i> <b>strip</b> [ <b>mac-format</b> <i>mac-format</i> ] <b>no user-name-format</b>				
<b>Context</b>	config>subscr-mgmt>auth-policy config>subscr-mgmt>diam-appl-plcy>nasreq				
<b>Description</b>	<p>This command defines the format of the “user-name” field in the session authentication request sent to the RADIUS server.</p> <p>The <b>no</b> form of the command switches to the default format, <b>mac</b>.</p>				
<b>Default</b>	By default, the MAC source address of the DHCP DISCOVER message is used in the user-name field.				
<b>Parameters</b>	<i>format</i> — Specifies the user name format in RADIUS message. <table> <tr> <td><b>Values</b></td><td> ascii-converted-circuit-id, ascii-converted-tuple, circuit-id, dhcp-client-vendor-opts, mac, mac-giaddr, tuple </td></tr> <tr> <td><b>ascii-converted-circuit-id</b></td><td> Identical to circuit-id, but the user name will be sent to the RADIUS server as a string of hex digits, for use if there is binary data in the circuit-id. </td></tr> </table>	<b>Values</b>	ascii-converted-circuit-id, ascii-converted-tuple, circuit-id, dhcp-client-vendor-opts, mac, mac-giaddr, tuple	<b>ascii-converted-circuit-id</b>	Identical to circuit-id, but the user name will be sent to the RADIUS server as a string of hex digits, for use if there is binary data in the circuit-id.
<b>Values</b>	ascii-converted-circuit-id, ascii-converted-tuple, circuit-id, dhcp-client-vendor-opts, mac, mac-giaddr, tuple				
<b>ascii-converted-circuit-id</b>	Identical to circuit-id, but the user name will be sent to the RADIUS server as a string of hex digits, for use if there is binary data in the circuit-id.				

**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, for use if there is binary data in the circuit-id.

**circuit-id** — If the system serves as a DHCP relay server which inserts option 82 info, the user name will be formatted as defined under DHCP information option. If the system is not a DHCP relay server, the circuit-id will be taken from option 82 in the received DHCP message. If no circuit-id can be found, the DHCP-msg is rejected.

**dhcp-client-vendor-opts** — Creates a concatenation of the DHCP client-identifier option (option 60), a “@” delimiter and the DHCP vendor-class identifier options. The two option strings are parsed for any characters which are non-printing are considered invalid and must be converted to underscore “\_” characters. In addition, any space character (hex 20) and @ character (hex 40) are also converted to underscore. The character set considered valid is ASCII hex 21 through hex 3F, and hex 41 through hex 7E. Any character outside this set will be converted into an underscore (hex 5F) character.

**mac** — The MAC source address of the DHCP DISCOVER message is used in the user-name field. The format of the MAC address string used as the user name in the RADIUS authentication requests uses lowercase hex digits, and “:” as the inter-digit separator, for example, 00:11:22:aa:bb:cc is valid but 00-11-22-AA-BB-CC will return an error. The RADIUS server must be configured accordingly, otherwise the authentication request will fail.

**mac-giaddr** — Specifies that MAC giaddr indicates the format used to identify the user towards the RADIUS server.

**tuple** — The concatenation of MAC source address and circuit-ID are used in the user-name field.

**mac-format** — Specifies how a MAC address is represented when contacting a RADIUS server. This is only used while the value of is equal to the DHCP client vendor options and if the MAC address is used by default of the DHCP client vendor options.

Examples:	ab:	00:0c:f1:99:85:b8	Alcatel-Lucent 7xxx style
	XY-	00-0C-F1-99-85-B8	IEEE canonical style
	mmmm.	0002.03aa.abff	Cisco style

**append** — Specifies the data type which is is an enumerated integer that indicates what needs to be appended to the user-name sent to the RADIUS server.

**Values**

- 1 — nothing
- 2 — domain name

**domain** — In some instances it is desired to add a domain only to usernames which have omitted the domain (@domain). In these instances a default-domain can be appended to usernames which lack a @domain

**append** — Adds a “@” delimiter and the specified string after the PAP/CHAP username. No allowance is made for the presence of an existing domain or @ delimited.

**replace** — Replaces the character-string after the “@” delimiter with the string specified.

**strip** — Removes all characters after and including the “@” delimiter.

Example:

```
Command: append
String: domainA-1.com
```

```

PAP/CHAP User:someuser
Resulting User:someuser@domainA-1.com

Command: append
String: domainA-1.com
PAP/CHAP User:someuser@existing-domain.net
Resulting User:someuser@existing-domain.net@domainA-1.com

Command: strip
String:
PAP/CHAP User:someuser@existing-domain.net
Resulting User:someuser

Command: replace
String: domainA-1.com
PAP/CHAP User:someuser@existing-domain.net
Resulting User:someuser@domainA-1.com

Command: default-domain
String:domainA-1.com
PAP/CHAP User:someuser@existing-domain.net
Resulting User:someuser@existing-domain.net

Command: default-domain
String: domainA-1.com
PAP/CHAP User:someuser
Resulting User:someuser@domainA-1.com

```

## user-name-format

<b>Syntax</b>	<b>user-name-format</b> <i>format</i> <b>no user-name-format</b>
<b>Context</b>	config>subscr-mgmt>diam-appl-plcy>nasreq
<b>Description</b>	This command defines the format of the User-Name AVP value in Diameter NASREQ AA-Requests for IPoE hosts.
<b>Parameters</b>	<i>format</i> — Specifies the format of the User-Name AVP value.
<b>Values</b>	<p><b>mac</b> — The MAC source address of the DHCP DISCOVER message is used in the user-name field. The format of the MAC address string is defined with the mac-format CLI command.</p> <p><b>circuit-id</b> — If the system serves as a DHCP relay server which inserts option 82 info, the user name will be formatted as defined under DHCP information option. If the system is not a DHCP relay server, the circuit-id will be taken from option 82 in the received DHCP message. If no circuit-id can be found, the DHCP-msg is rejected.</p> <p><b>tuple</b> — A concatenation of MAC source address and circuit-ID.</p> <p><b>ascii-converted-circuit-id</b> — Identical to circuit-id, but the user name is a string of hex digits, for use if there is binary data in the circuit-id.</p> <p><b>ascii-converted-tuple</b> — Identical to tuple, but the circuit-id part of the user name is a string of hex digits, for use if there is binary data in the circuit-id.</p> <p><b>dhcp-client-vendor-opts</b> — A concatenation of the DHCP client-identifier option (option 60), “@” as delimiter and the DHCP vendor-class identifier options. Spaces</p>

(hex 20), @ character (hex 40) and non printable characters (all character outside range hex 21 through hex 7E) are converted to underscore “\_” (hex 5F).  
**mac-giaddr** — A concatenation of MAC source address and DHCP gi address.  
**nas-port-id** — the value of the nas-port-id with format defined in the include-avp section.

user-name-operation

<b>Syntax</b>	<b>user-name-operation operation [domain <i>domain-name</i>]</b> <b>no user-name-operation</b>
<b>Context</b>	config>subscr-mgmt>diam-appl-plcy>nasreq
<b>Description</b>	<p>This command enables domain name manipulation of the user name, such as append, strip, replace or add as default.</p> <p>For IPoE, this command only applies when <b>user-name-format</b> is configured to <b>dhcp-client-vendor-opts</b>.</p>
<b>Default</b>	no user-name-operation
<b>Parameters</b>	<p><b>operation</b> — Specifies the user name manipulations with respect to domain name values.</p> <p><b>Values</b></p> <ul style="list-style-type: none"><li><b>append-domain</b> – appends an “@” delimiter with the specified domain-name at the end of the user-name, independent if a domain name was already present.</li><li><b>strip-domain</b> – removes all characters after and including the “@” delimiter.</li><li><b>default-domain</b> – adds an “@” delimiter and the specified domain name to user-names that have no domain name present.</li><li><b>replace-domain</b> – replaces the characters after the “@” delimiter with the specified domain-name.</li></ul> <p><b>domain <i>domain-name</i></b> — Specifies the domain name string to be used in the specified operation. Maximum 128 characters.</p>

---

## RADIUS Accounting Policy Custom Record Commands

### custom-record

<b>Syntax</b>	<b>[no] custom-record</b>
<b>Context</b>	config>subscr-mgmt>acct-plcy
<b>Description</b>	This command enables the context to configure the layout and setting for a custom accounting record associated with this accounting policy. The <b>no</b> form of the command reverts the configured values to the defaults.

### override-counter

<b>Syntax</b>	<b>[no] override-counter</b> <i>override-counter-id</i>
<b>Context</b>	config>log>acct-policy>cr
<b>Description</b>	This command enables the context to configure Application Assurance override counter parameters. The <b>no</b> form of the command removes the ID from the configuration.
<b>Parameters</b>	<i>override-counter-id</i> — Specifies the override counter ID. <b>Values</b> 1 — 8

### e-counters

<b>Syntax</b>	<b>e-counters [all]</b> <b>no e-counters</b>
<b>Context</b>	config>log>acct-policy>cr>override-cntr config>log>acct-policy>cr>queue config>log>acct-policy>cr>ref-override-cntr config>log>acct-policy>cr>ref-queue
<b>Description</b>	This command configures egress counter parameters for this custom record. The <b>no</b> form of the command
<b>Parameters</b>	<b>all</b> — Includes all counters.

### i-counters

<b>Syntax</b>	<b>i-counters [all]</b> <b>no i-counters</b>
<b>Context</b>	config>log>acct-policy>cr>override-cntr config>log>acct-policy>cr>ref-override-cntr config>log>acct-policy>cr>ref-queue
<b>Description</b>	This command configures ingress counter parameters for this custom record. The <b>no</b> form of the command
<b>Parameters</b>	<b>all</b> — Includes all counters.

### queue

<b>Syntax</b>	<b>[no] queue <i>queue-id</i></b>
<b>Context</b>	config>log>acct-policy>cr
<b>Description</b>	This command specifies the queue-id for which counters will be collected in this custom record. The counters that will be collected are defined in egress and ingress counters. The <b>no</b> form of the command reverts to the default value
<b>Parameters</b>	<i>queue-id</i> — Specifies the queue-id for which counters will be collected in this custom record.

### in-profile-octets-discarded-count

<b>Syntax</b>	<b>[no] in-profile-octets-discarded-count</b>
<b>Context</b>	config>log>acct-policy>cr>oc>e-count config>log>acct-policy>cr>roc>e-count config>log>acct-policy>cr>queue>e-count config>log>acct-policy>cr>ref-queue>e-count
<b>Description</b>	This command includes the in-profile octets discarded count. For queues with <b>stat-mode v4-v6</b> , this command includes the IPv4 octets discarded count instead. The <b>no</b> form of the command excludes the in-profile octets discarded count.

### in-profile-octets-forwarded-count

<b>Syntax</b>	<b>[no] in-profile-octets-forwarded-count</b>
<b>Context</b>	config>log>acct-policy>cr>oc>e-count config>log>acct-policy>cr>roc>e-count config>log>acct-policy>cr>queue>e-count config>log>acct-policy>cr>ref-queue>e-count

**Description** This command includes the in-profile octets forwarded count. For queues with **stat-mode v4-v6**, this command includes the IPv4 octets forwarded count instead.

The **no** form of the command excludes the in-profile octets forwarded count.

## in-profile-packets-discarded-count

**Syntax** **[no] in-profile-packets-discarded-count**

**Context** config>log>acct-policy>cr>oc>e-count  
 config>log>acct-policy>cr>roc>e-count  
 config>log>acct-policy>cr>queue>e-count  
 config>log>acct-policy>cr>ref-queue>e-count

**Description** This command includes the in-profile packets discarded count.

For queues with **stat-mode v4-v6**, this command includes the IPv4 packets discarded count instead.

The **no** form of the command excludes the in-profile packets discarded count.

## in-profile-packets-forwarded-count

**Syntax** **[no] in-profile-packets-forwarded-count**

**Context** config>log>acct-policy>cr>oc>e-count  
 config>log>acct-policy>cr>roc>e-count  
 config>log>acct-policy>cr>queue>e-count  
 config>log>acct-policy>cr>ref-queue>e-count

**Description** This command includes the in-profile packets forwarded count.

For queues with **stat-mode v4-v6**, this command includes the IPv4 packets forwarded count instead.

The **no** form of the command excludes the in-profile packets forwarded count.

## out-profile-octets-discarded-count

**Syntax** **[no] out-profile-octets-discarded-count**

**Context** config>log>acct-policy>cr>oc>e-count  
 config>log>acct-policy>cr>roc>e-count  
 config>log>acct-policy>cr>queue>e-count  
 config>log>acct-policy>cr>ref-queue>e-count

**Description** This command includes the out of profile packets discarded count.

For queues with **stat-mode v4-v6**, this command includes the IPv6 octets discarded count instead.

The **no** form of the command excludes the out of profile packets discarded count.

## out-profile-octets-forwarded-count

<b>Syntax</b>	<b>[no] out-profile-octets-forwarded-count</b>
<b>Context</b>	config>log>acct-policy>cr>oc>e-count config>log>acct-policy>cr>roc>e-count config>log>acct-policy>cr>queue>e-count config>log>acct-policy>cr>ref-queue>e-count
<b>Description</b>	This command includes the out of profile octets forwarded count. For queues with <b>stat-mode v4-v6</b> , this command includes the IPv6 octets forwarded count instead. The <b>no</b> form of the command excludes the out of profile octets forwarded count.

## out-profile-packets-discarded-count

<b>Syntax</b>	<b>[no] out-profile-packets-discarded-count</b>
<b>Context</b>	config>log>acct-policy>cr>oc>e-count config>log>acct-policy>cr>roc>e-count config>log>acct-policy>cr>queue>e-count config>log>acct-policy>cr>ref-queue>e-count
<b>Description</b>	This command includes the out of profile packets discarded count. For queues with <b>stat-mode v4-v6</b> , this command includes the IPv6 packets discarded count instead. The <b>no</b> form of the command excludes the out of profile packets discarded count.

## out-profile-packets-forwarded-count

<b>Syntax</b>	<b>[no] out-profile-packets-forwarded-count</b>
<b>Context</b>	config>log>acct-policy>cr>oc>e-count config>log>acct-policy>cr>roc>e-count config>log>acct-policy>cr>queue>e-count config>log>acct-policy>cr>ref-queue>e-count
<b>Description</b>	This command includes the out of profile packets forwarded count. For queues with <b>stat-mode v4-v6</b> , this command includes the IPv6 packets forwarded count instead. The <b>no</b> form of the command excludes the out of profile packets forwarded count.

## all-octets-offered-count

<b>Syntax</b>	<b>[no] all-octets-offered-count</b>
<b>Context</b>	config>log>acct-policy>cr>oc>i-count config>log>acct-policy>cr>roc>i-count



```
config>log>acct-policy>cr>queue>i-count
config>log>acct-policy>cr>ref-queue>i-count
```

- Description** This command includes all octets offered in the count.  
The **no** form of the command excludes the octets offered in the count.
- Default** no all-octets-offered-count

## all-packets-offered-count

- Syntax** [no] all-packets-offered-count
- Context** config>log>acct-policy>cr>oc>i-count  
config>log>acct-policy>cr>roc>i-count  
config>log>acct-policy>cr>queue>i-count  
config>log>acct-policy>cr>ref-queue>i-count
- Description** This command includes all packets offered in the count.  
The **no** form of the command excludes the packets offered in the count.
- Default** no all-packets-offered-count

## high-octets-discarded-count

- Syntax** [no] high-octets-discarded-count
- Context** config>log>acct-policy>cr>oc>i-count  
config>log>acct-policy>cr>roc>i-count  
config>log>acct-policy>cr>queue>i-count  
config>log>acct-policy>cr>ref-queue>i-count
- Description** This command includes the high octets discarded count.  
For queues with **stat-mode v4-v6**, this command includes the IPv4 octets discarded count instead.  
The **no** form of the command excludes the high octets discarded count.
- Default** no high-octets-discarded-count

## high-octets-offered-count

- Syntax** [no] high-octets-offered-count
- Context** config>log>acct-policy>cr>oc>i-count  
config>log>acct-policy>cr>roc>i-count  
config>log>acct-policy>cr>queue>i-count  
config>log>acct-policy>cr>ref-queue>i-count
- Description** This command includes the high octets offered count.

The **no** form of the command excludes the high octets offered count.

### high-packets-discarded-count

<b>Syntax</b>	<b>[no] high-packets-discarded-count</b>
<b>Context</b>	config>log>acct-policy>cr>oc>i-count config>log>acct-policy>cr>roc>i-count config>log>acct-policy>cr>queue>i-count config>log>acct-policy>cr>ref-queue>i-count
<b>Description</b>	This command includes the high packets discarded count.  For queues with <b>stat-mode v4-v6</b> , this command includes the IPv4 packets discarded count instead.  The <b>no</b> form of the command excludes the high packets discarded count.
<b>Default</b>	no high-packets-discarded-count

### high-packets-offered-count

<b>Syntax</b>	<b>[no] high-packets-offered-count</b>
<b>Context</b>	config>log>acct-policy>cr>oc>i-count config>log>acct-policy>cr>roc>i-count config>log>acct-policy>cr>queue>i-count config>log>acct-policy>cr>ref-queue>i-count
<b>Description</b>	This command includes the high packets offered count.  The <b>no</b> form of the command excludes the high packets offered count.
<b>Default</b>	no high-packets-offered -count

### in-profile-octets-forwarded-count

<b>Syntax</b>	<b>[no] in-profile-octets-forwarded-count</b>
<b>Context</b>	config>log>acct-policy>cr>oc>i-count config>log>acct-policy>cr>roc>i-count config>log>acct-policy>cr>queue>i-count config>log>acct-policy>cr>ref-queue>i-count
<b>Description</b>	This command includes the in profile octets forwarded count.  For queues with <b>stat-mode v4-v6</b> , this command includes the IPv4 octets forwarded count instead.  The <b>no</b> form of the command excludes the in profile octets forwarded count.
<b>Default</b>	no in-profile-octets-forwarded-count

## in-profile-packets-forwarded-count

<b>Syntax</b>	<b>[no] in-profile-packets-forwarded-count</b>
<b>Context</b>	config>log>acct-policy>cr>oc>i-count config>log>acct-policy>cr>roc>i-count config>log>acct-policy>cr>queue>i-count config>log>acct-policy>cr>ref-queue>i-count
<b>Description</b>	This command includes the in profile packets forwarded count.  For queues with <b>stat-mode v4-v6</b> , this command includes IPv4 packets forwarded count instead.  The <b>no</b> form of the command excludes the in profile packets forwarded count.
<b>Default</b>	no in-profile-packets-forwarded-count

## low-octets-discarded-count

<b>Syntax</b>	<b>[no] low-octets-discarded-count</b>
<b>Context</b>	config>log>acct-policy>cr>oc>i-count config>log>acct-policy>cr>roc>i-count config>log>acct-policy>cr>queue>i-count config>log>acct-policy>cr>ref-queue>i-count
<b>Description</b>	This command includes the low octets discarded count.  For queues with <b>stat-mode v4-v6</b> , this command includes the IPv6 octets discarded count instead.  The <b>no</b> form of the command excludes the low octets discarded count.
<b>Default</b>	no low-octets-discarded-count

## low-packets-discarded-count

<b>Syntax</b>	<b>[no] low-packets-discarded-count</b>
<b>Context</b>	config>log>acct-policy>cr>oc>i-count config>log>acct-policy>cr>roc>i-count config>log>acct-policy>cr>queue>i-count config>log>acct-policy>cr>ref-queue>i-count
<b>Description</b>	This command includes the low packets discarded count.  For queues with <b>stat-mode v4-v6</b> , this command includes the IPv6 packets discarded count instead.  The <b>no</b> form of the command excludes the low packets discarded count.
<b>Default</b>	no low-packets-discarded-count

### low-octets-offered-count

<b>Syntax</b>	<b>[no] low-octets-offered-count</b>
<b>Context</b>	config>log>acct-policy>cr>oc>i-count config>log>acct-policy>cr>roc>i-count config>log>acct-policy>cr>queue>i-count config>log>acct-policy>cr>ref-queue>i-count
<b>Description</b>	This command includes the low octets discarded count. The <b>no</b> form of the command excludes the low octets discarded count.

### low-packets-offered-count

<b>Syntax</b>	<b>[no] low-packets-offered-count</b>
<b>Context</b>	config>log>acct-policy>cr>oc>i-count config>log>acct-policy>cr>roc>i-count config>log>acct-policy>cr>queue>i-count config>log>acct-policy>cr>ref-queue>i-count
<b>Description</b>	This command includes the low packets discarded count. The <b>no</b> form of the command excludes the low packets discarded count.

### out-profile-octets-forwarded-count

<b>Syntax</b>	<b>[no] out-profile-octets-forwarded-count</b>
<b>Context</b>	config>log>acct-policy>cr>oc>i-count config>log>acct-policy>cr>roc>i-count config>log>acct-policy>cr>queue>i-count config>log>acct-policy>cr>ref-queue>i-count
<b>Description</b>	This command includes the out of profile octets forwarded count. For queues with <b>stat-mode v4-v6</b> , this command includes the IPv6 octets forwarded count instead. The <b>no</b> form of the command excludes the out of profile octets forwarded count.
<b>Default</b>	no out-profile-octets-forwarded-count

### out-profile-packets-forwarded-count

<b>Syntax</b>	<b>[no] out-profile-packets-forwarded-count</b>
<b>Context</b>	config>log>acct-policy>cr>oc>i-count config>log>acct-policy>cr>roc>i-count config>log>acct-policy>cr>queue>i-count

```
config>log>acct-policy>cr>ref-queue>i-count
```

- Description** This command includes the out of profile packets forwarded count.  
For queues with **stat-mode v4-v6**, this command includes the IPv6 packets forwarded count instead.  
The **no** form of the command excludes the out of profile packets forwarded count.
- Default** no out-profile-packets-forwarded-count

## uncoloured-octets-offered-count

- Syntax** **[no] uncoloured-packets-offered-count**
- Context** config>log>acct-policy>cr>queue>i-count  
config>log>acct-policy>cr>ref-queue>i-count
- Description** This command includes the uncoloured octets offered in the count.  
The **no** form of the command excludes the uncoloured octets offered in the count.

## uncoloured-packets-offered-count

- Syntax** **[no] uncoloured-packets-offered-count**
- Context** config>log>acct-policy>cr>queue>i-count  
config>log>acct-policy>cr>ref-queue>i-count
- Description** This command includes the uncoloured packets offered count.  
The **no** form of the command excludes the uncoloured packets offered count.

## ref-aa-specific-counter

- Syntax** **[no] ref-aa-specific-counter any**
- Context** config>log>acct-policy>cr
- Description** This command  
The **no** form of the command

## ref-override-counter

- Syntax** **ref-override-counter** *ref-override-counter-id*  
**ref-override-counter all**  
**no ref-override-counter**
- Context** config>log>acct-policy>cr

## RADIUS Policy Commands

<b>Description</b>	This command configures a reference override counter. The <b>no</b> form of the command reverts to the default value.
<b>Default</b>	no ref-override-counter

### ref-queue

<b>Syntax</b>	<b>ref-queue</b> <i>queue-id</i> <b>ref-queue all</b> <b>no ref-queue</b>
<b>Context</b>	config>log>acct-policy>cr
<b>Description</b>	This command configures a reference queue. The <b>no</b> form of the command reverts to the default value.
<b>Default</b>	no ref-queue

### significant-change

<b>Syntax</b>	<b>significant-change</b> <i>delta</i> <b>no significant-change</b>
<b>Context</b>	config>log>acct-policy>cr
<b>Description</b>	This command configures the significant change required to generate the record.
<b>Parameters</b>	<i>delta</i> — Specifies the delta change (significant change) that is required for the custom record to be written to the xml file. <b>Values</b> 0 — 4294967295

## RADIUS Route Download Commands

### route-downloader

<b>Syntax</b>	<b>route-downloader</b> <i>name</i> [ <b>create</b> ] <b>no route-downloader</b> <i>name</i>
<b>Context</b>	config>aaa
<b>Description</b>	<p>This command creates or enters the configuration of a route-downloader instance. The route-downloader is a process that uses radius access-request messages to a particular server. The server returns either an access-accept or access-deny message. Access-accept messages also contain the prefixes (in the form of static blackhole routes in various formats)</p> <p>The <b>no</b> form of the command removes the name from the configuration. The object must be shutdown prior to deletion. No prefix is needed to delete an existing route-download object.</p>
<b>Default</b>	None. Only a single route-downloader object can be created.
<b>Parameters</b>	<p><i>name</i> — Specifies the name of this RADIUS route downloader.</p> <p><b>create</b> — This keyword is mandatory while creating an instance of the route-download object.</p>

### base-user-name

<b>Syntax</b>	<b>base-user-name</b> <i>user-name</i> <b>no base-user-name</b>
<b>Context</b>	config>aaa>route-downloader
<b>Description</b>	<p>This command sets the prefix for the user name that shall be used as access requests. The actual name used will be a concatenation of this string, the “-” (dash) character and a monotonically increasing integer.</p> <p>The <b>no</b> form of the command removes the user-name from the configuration.</p>
<b>Default</b>	The system’s configured name (system-name).
<b>Parameters</b>	<i>user-name</i> — Specifies the prefix of the username that is used in the RADIUS access requests. The username used in the RADIUS access requests is a concatenation of this string, the dash character and an increasing integer.

### default-metric

<b>Syntax</b>	<b>default-metric</b> <i>metric</i> <b>no default-metric</b>
<b>Context</b>	config>aaa>route-downloader

## RADIUS Route Download Commands

<b>Description</b>	This command sets the default metric that routes imported by the RTM will acquire. The no form of the command removes the metric
<b>Default</b>	2
<b>Parameters</b>	<i>metric</i> — Specifies the default metric of the routes imported. <b>Values</b> 0 — 254

### default-tag

<b>Syntax</b>	<b>default-tag</b> <i>tag</i> <b>no default-tag</b>
<b>Context</b>	config>aaa>route-downloader
<b>Description</b>	This command sets the default tag that routes processed by the AAA route downloader will take. Note that any route received with a specific tag retains the specific tag. The tag value is passed to the Route Table Manager and is available as match condition on the export statement of other routing protocols. The <b>no</b> form of the command reverts to the default.
<b>Default</b>	0
<b>Parameters</b>	<i>tag</i> — Specifies the default tag of the routes imported. <b>Values</b> 0 — 4294967295

### download-interval

<b>Syntax</b>	<b>download-interval</b> <i>minutes</i> <b>no download-interval</b>
<b>Context</b>	config>aaa>route-downloader
<b>Description</b>	This command sets the time interval, in minutes, that the system waits for between two consecutive runs of the route-download process. The time is counted from the start-time of the run, thus, if an route-download process is still ongoing by the time the timer expires, the process will restart from count=1. The no form of the command reverts to the default value.
<b>Default</b>	720
<b>Parameters</b>	<i>minutes</i> — Specifies the time interval, in minutes, between the start of the last route downloader run and the start of the next route downloader run. <b>Values</b> 1 — 1440



## max-routes

**max-routes** *routes*  
**no max-routes**

<b>Context</b>	config>aaa>route-downloader
<b>Description</b>	<p>This command determines the upper limits for total number of routes to be received and accepted by the system. The total number is inclusive of both IPv4 and IPv6 addresses and no differentiation is needed across protocols. It includes the sum of both. Once this limit is reached, the download process stops sending new access-requests until the next download-interval expires.</p> <p>The no form of the command reverts to the default value.</p>
<b>Default</b>	200000
<b>Parameters</b>	<p><i>routes</i> — Specifies the maximum number of the routes imported.</p> <p><b>Values</b>      1 — 200000</p>

## password

**password** *password* [*hash*|*hash2*]  
**no password**

<b>Context</b>	config>aaa>route-downloader
<b>Description</b>	<p>This command specifies the password that is used in the RADIUS access requests. It shall be specified as a string of up to 32 characters in length.</p> <p>The <b>no</b> form of the command resets the password to its default of <b>ALU</b> and will be stored using <i>hash</i>/<i>hash2</i> encryption.</p>
<b>Default</b>	ALU
<b>Parameters</b>	<p><i>password</i> — Specifies a password string up to 32 characters in length.</p> <p><b>hash</b> — Specifies the key is entered in an encrypted form. If the <b>hash</b> parameter is not used, the key is assumed to be in a non-encrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the <b>hash</b> parameter specified.</p> <p><b>hash2</b> — Specifies the key is entered in a more complex encrypted form. If the <b>hash2</b> parameter is not used, the less encrypted <b>hash</b> form is assumed.</p>

## radius-server-policy

<b>Syntax</b>	<b>radius-server-policy</b> <i>policy-name</i> <b>no radius-server-policy</b>
<b>Context</b>	config>aaa>route-downloader

## RADIUS Route Download Commands

<b>Description</b>	This command references an existing radius-server-policy (available under the <b>config&gt;aaa</b> context). The server (or servers) referenced by the policy will be used as the targets for the access-request message.  The <b>no</b> form of the command removes the policy name from the route-downloader configuration.
<b>Default</b>	none
<b>Parameters</b>	<i>policy-name</i> — Specifies the RADIUS server policy.

## retry-interval

<b>Syntax</b>	<b>retry-interval min <i>minimum</i> max <i>maximum</i></b> <b>no retry-interval</b>
<b>Context</b>	config>aaa>route-downloader
<b>Description</b>	This command sets the duration, in minutes, of the retry interval. The retry interval is the interval meant for the system to retry sending an Access Request message after the previous one was unanswered (not with an access reject but rather just a RADIUS failure or ICMP port unreachable). This timer is actually an exponential backoff timer that starts at <b>min</b> and is capped at <b>max</b> minutes.  The <b>no</b> form of the command reverts to the default values.
<b>Default</b>	retry-interval min 10 max 20
<b>Parameters</b>	<b>min <i>minimum</i></b> — Specifies the duration, in minutes, of the retry interval. This duration grows exponentially after each sequential failure.  <b>Values</b> 1 — 1440  <b>Default</b> 10  <b>max <i>maximum</i></b> — Specifies the maximum duration, in minutes, of the retry interval.  <b>Values</b> 1 — 1440  <b>Default</b> 20

## Category Map Commands

### category-map

<b>Syntax</b>	<b>category-map</b> <i>category-map-name</i> [ <b>create</b> ] <b>no category-map</b> <i>category-map-name</i>
<b>Context</b>	config>subscr-mgmt config>subscr-mgmt>sla-prof
<b>Description</b>	This command specifies the category map name.
<b>Default</b>	none
<b>Parameters</b>	<i>category-map-name</i> — Specifies the category map name up to 32 characters in length. <b>create</b> — Mandatory keyword when creating a new category map.

### credit-control-policy

<b>Syntax</b>	<b>credit-control-policy</b> <i>policy-name</i> [ <b>create</b> ] <b>no credit-control-policy</b> <i>policy-name</i>
<b>Context</b>	config>subscr-mgmt
<b>Description</b>	This command creates, configures or deletes a credit control policy.
<b>Parameters</b>	<i>policy-name</i> — Specifies the policy name, 32 characters max.

### credit-control-server

<b>Syntax</b>	<b>credit-control-server</b> radius <b>credit-control-server</b> <b>diameter</b> <i>policy-name</i> <b>no credit-control-server</b>
<b>Context</b>	config>subscr-mgmt>credit-control-policy
<b>Description</b>	This command configures the credit control server to use. In case of RADIUS, the servers defined in the authentication policy are used. For Diameter, the peers defined in the specified Diameter policy are used.
<b>Default</b>	no credit-control-server
<b>Parameters</b>	<b>radius</b> — Use the RADIUS authentication servers defined in the RADIUS authentication policy in the group-interface to report credit usage and obtain new credit. <b>diameter</b> <i>policy-name</i> — Use the diameter peers specified in the diameter <b>policy</b> <i>policy-name</i> to report credit usage and obtain new credit.

## default-category-map

<b>Syntax</b>	<b>default-category-map</b> <i>category-map-name</i> <b>no default-category-map</b>
<b>Context</b>	config>subscr-mgmt>credit-control-policy
<b>Description</b>	This command configures the default category map.
<b>Parameters</b>	<i>category-map-name</i> — Specifies the category map name, 32 chars max.

## error-handling-action

<b>Syntax</b>	<b>error-handling-action</b> { <b>continue</b>   <b>block</b> } <b>no error-handling-action</b>
<b>Context</b>	config>subscr-mgmt>credit-control-policy
<b>Description</b>	This command configures the error handling action for the policy.

## out-of-credit-action

<b>Syntax</b>	<b>out-of-credit-action</b> { <b>continue</b>   <b>disconnect-host</b>   <b>block-category</b>   <b>change-service-level</b> } <b>no out-of-credit-action</b>
<b>Context</b>	config>subscr-mgmt>credit-control-policy
<b>Description</b>	This command configures the action to be performed when out of credit is reached.
<b>Parameters</b>	{ <b>continue</b>   <b>disconnect-host</b>   <b>block-category</b>   <b>change-service-level</b> } — Specifies the action to be taken when out of credit is reached.

## activity-threshold

<b>Syntax</b>	<b>activity-threshold</b> <i>kilobits-per-second</i> <b>no activity-threshold</b>
<b>Context</b>	config>subscr-mgmt>cat-map
<b>Description</b>	This command configures the threshold that is applied to determine whether or not there is activity. This is only valid for credit-type = time (not volume).
<b>Default</b>	0
<b>Parameters</b>	<i>kilobits-per-second</i> — Specifies the activity threshold value in kilobits per second.
<b>Values</b>	1 — 100000000

## category

<b>Syntax</b>	<b>category</b> <i>category-name</i> [ <b>create</b> ] <b>no category</b> <i>category-name</i>
<b>Context</b>	config>subscr-mgmt>cat-map
<b>Description</b>	This command specifies the category name.
<b>Default</b>	none
<b>Parameters</b>	<i>category-name</i> — Specifies the category name up to 32 characters in length. <b>create</b> — Mandatory keyword when creating a new category.

## category-map

<b>Syntax</b>	<b>category-map</b> <i>category-map-name</i> <b>no category-map</b>
<b>Context</b>	config>subscr-mgmt>sla-prof
<b>Description</b>	This command references the category-map to be used for the idle-timeout monitoring of subscriber hosts associated with this sla-profile. The <b>category-map</b> must already exist in the <b>config&gt;subscr-mgmt</b> context.
<b>Parameters</b>	<i>category-map-name</i> — Specifies the name of the category map (up to 32 characters in length) where the activity-threshold and the category is defined for idle-timeout monitoring of subscriber hosts.

## category

<b>Syntax</b>	<b>category</b> <i>category-name</i> [ <b>create</b> ] <b>no category</b> <i>category-name</i>
<b>Context</b>	config>subscr-mgmt>sla-prof>cat-map
<b>Description</b>	This command defines the category in the category-map to be used for the idle-timeout monitoring of subscriber hosts.
<b>Parameters</b>	<i>category-name</i> — Specifies the name (up to 32 characters in length) of the category where the queues and policers are defined for idle-timeout monitoring of subscriber hosts. <b>create</b> — Mandatory keyword when creating a new category

## idle-timeout

<b>Syntax</b>	<b>idle-timeout</b> <i>timeout</i> <b>no idle-timeout</b>
<b>Context</b>	config>subscr-mgmt>sla-prof>cat-map>category

## RADIUS Route Download Commands

<b>Description</b>	This command defines the idle-timeout value.
<b>Default</b>	no idle-timeout – corresponds with an infinite idle-timeout
<b>Parameters</b>	<i>timeout</i> — Specifies the idle-timeout in seconds.
<b>Values</b>	60 — 15552000

### idle-timeout-action

<b>Syntax</b>	<b>idle-timeout-action {shcv-check   terminate}</b> <b>no idle-timeout-action</b>
<b>Context</b>	config>subscr-mgmt>sla-prof>cat-map>category
<b>Description</b>	This command defines the action to be executed when the idle-timeout is reached. The action is performed for all hosts associated with the sla-profile instance.
<b>Default</b>	terminate
<b>Parameters</b>	<b>shcv-check</b> — performs a subscriber host connectivity verification check (IPoE hosts only). Note that host connectivity verification must be enabled on the group-interface where the host is connected.  If the check is successful, the hosts are not disconnected and the idle-timeout timer is reset.  If the check fails, the hosts are deleted, similar as for “idle-timeout-action=terminate”.  <b>terminate</b> — Deletes the subscriber host from the system: for PPP hosts, a terminate request is send; for IPoE hosts a DHCP release is send to the DHCP server.

### credit-type-override

<b>Syntax</b>	<b>credit-type-override {volume   time}</b> <b>no credit-type-override</b>
<b>Context</b>	config>subscr-mgmt>cat-map>category
<b>Description</b>	This command overrides the <b>credit-type</b> configured in the <b>config&gt;subscr-mgmt&gt;cat-map</b> context for the given category.
<b>Default</b>	no credit-type-override
<b>Parameters</b>	<b>volume</b> — If different than the value specified in the <b>credit-type</b> command, the value overrides the credit-type.  <b>time</b> — If different than the value specified in the <b>credit-type</b> command, the value overrides the credit-type.

### default-credit

<b>Syntax</b>	<b>default-credit volume</b> <i>credits</i> bytes   kilobytes   megabytes   gigabytes
---------------	---

	<b>default-credit</b> <i>time seconds</i> <b>no default-credit</b>
<b>Context</b>	config>subscr-mgmt>cat-map>category
<b>Description</b>	This command configures the default credit used by this category.
<b>Default</b>	no default-credit
<b>Parameters</b>	<b>volume</b> <i>credits bytes kilobytes megabytes gigabytes</i> — Specifies the default value for the volume credit and the unit in which the default value is expressed. <b>Values</b> 1 — 4294967295 (minimum 100 megabytes) <b>time</b> <i>seconds</i> — Specifies the default value for the time credit, in seconds. <b>Values</b> 900 — 4294967295 (minimum 900 seconds)

## exhausted-credit-service-level

<b>Syntax</b>	<b>[no] exhausted-credit-service-level</b>
<b>Context</b>	config>subscr-mgmt>cat-map>category
<b>Description</b>	This command enables the context to configure the exhausted credit service level
<b>Default</b>	exhausted-credit-service-level

## egress-ip-filter-entries

<b>Syntax</b>	<b>[no] egress-ip-filter-entries</b>
<b>Context</b>	config>subscr-mgmt>cat-map>category>exh-lvl
<b>Description</b>	This command configures the egress IP filter entries.

## egress-ipv6-filter-entries

<b>Syntax</b>	<b>[no] egress-ipv6-filter-entries</b>
<b>Context</b>	config>subscr-mgmt>cat-map>category>exh-lvl
<b>Description</b>	This command configures the egress IPv6 filter entries.

## ingress-ip-filter-entries

<b>Syntax</b>	<b>[no] ingress-ip-filter-entries</b>
<b>Context</b>	config>subscr-mgmt>cat-map>category>exh-lvl

**Description** This command configures the ingress IP filter entries.

### ingress-ipv6-filter-entries

**Syntax** **[no] ingress-ipv6-filter-entries**

**Context** config>subscr-mgmt>cat-map>category>exh-lvl

**Description** This command configures the ingress IPv6 filter entries.

### pir

**Syntax** **[no] pir**

**Context** config>subscr-mgmt>cat-map>category>exh-lvl

**Description** This command configures the PIR.

### entry

**Syntax** **entry entry-id [create]**

**Context** config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ip  
config>subscr-mgmt>cat-map>category>exh-lvl>egr-ip

**Description** This command configures the IP filter entry.

**Parameters** *entry-id* — Specifies the entry ID.

**Values** 1..65535

### action

**Syntax** **action drop**  
**action forward**  
**action http-redirect url**  
**no action**

**Context** config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ip>entry  
config>subscr-mgmt>cat-map>category>exh-lvl>egr-ip>entry

**Description** This command configures the action for the filter entry.

**Parameters** **drop** — Specifies to drop the IP filter entry.

**forward** — Specifies to forward the IP filter entry.

**http-redirect url** — Specifies the HTTP web address that will be sent to the user's browser. Note that http-redirect is not supported on 7750 SR-1 or 7450 ESS-1 models.



The following displays information that can optionally be added as variables in the portal URL (http-redirect url):

- \$IP – Customer’s IP address
- \$MAC – Customer’s MAC address
- \$URL – Original requested URL
- \$SAP – Customer’s SAP
- \$SUB – Customer’s subscriber identification string
- \$CID – string that represents the circuit-id or interface-id of the subscriber host (hexadecimal format)
- \$RID – string that represents the remote-id of the subscriber host (hexadecimal format)

**Values** 255 characters maximum

## match

<b>Syntax</b>	<b>match</b> [ <i>next-header next-header</i> ] <b>no match</b>
<b>Context</b>	config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ip>entry config>subscr-mgmt>cat-map>category>exh-lvl>egr-ip>entry
<b>Description</b>	This command configures the match criteria for this IP filter entry.
<b>Parameters</b>	<i>protocol-id</i> — Specifies the protocol number accepted in DHB.
	<b>Values</b> 0..255

## dscp

<b>Syntax</b>	<b>dscp</b> <i>dscp-name</i> <b>no dscp</b>
<b>Context</b>	config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ip>entry>match config>subscr-mgmt>cat-map>category>exh-lvl>egr-ip>entry>match
<b>Description</b>	This command configures DSCP match conditions.
<b>Parameters</b>	<i>dscp-name</i> — Specifies the DSCP name.
	<b>Values</b> 32 chars max

## dst-ip

<b>Syntax</b>	<b>dst-ip</b> { <i>ip-address/mask</i>   <i>ip-address netmask</i> } <b>no dst-ip</b>
<b>Context</b>	config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ip>entry>match

## RADIUS Route Download Commands

**Description** This command configures the destination IP match condition.

**Parameters** *ip-address/mask* — Specifies the IPv4 address and mask.

**Values**      *ip-address*      *a.b.c.d*  
                 *mask*    0..32

*netmask* — Specifies the mask, expressed as a dotted quad.

**Values**      *a.b.c.d*

### dst-port

**Syntax**      **dst-port** {*lt* | *gt* | *eq*} *dst-port-number*  
                 **dst-port range** *start end*  
                 **no dst-port**

**Context**      config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ip>entry>match  
                 config>subscr-mgmt>cat-map>category>exh-lvl>egr-ip>entry>match

**Description** This command configures the destination port match condition.

**Parameters** *lt|gt|eq* — Specifies the operator.

*dst-port-number* — Specifies the destination port number as a decimal hex or binary.

**Values**      0..65535

### fragment

**Syntax**      **fragment** {*true* | *false*}

**Context**      config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ip>entry>match  
                 config>subscr-mgmt>cat-map>category>exh-lvl>egr-ip>entry>match

**Description** This command configures the fragmentation match condition.

**Parameters** *true|false* — Sets/resets fragmentation check.

### icmp-code

**Syntax**      **icmp-code** *icmp-code*  
                 **no icmp-code**

**Context**      config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ip>entry>match  
                 config>subscr-mgmt>cat-map>category>exh-lvl>egr-ip>entry>match

**Description** This command configures the ICMP code match condition.

**Parameters** *icmp-code* — Specifies the ICMP code numbers accepted in DHB.

**Values**      0..255

## icmp-type

<b>Syntax</b>	<b>icmp-type</b> <i>icmp-type</i> <b>no icmp-type</b>
<b>Context</b>	config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ip>entry>match config>subscr-mgmt>cat-map>category>exh-lvl>egr-ip>entry>match
<b>Description</b>	This command configures the ICMP type match condition.
<b>Parameters</b>	<i>icmp-type</i> — Specifies the ICMP type numbers accepted in DHB. <b>Values</b> 0..255

## ip-option

<b>Syntax</b>	<b>ip-option</b> <i>ip-option-value</i> [ <i>ip-option-mask</i> ] <b>no ip-option</b>
<b>Context</b>	config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ip>entry>match config>subscr-mgmt>cat-map>category>exh-lvl>egr-ip>entry>match
<b>Description</b>	This command configures the IP option match condition.
<b>Parameters</b>	<i>ip-option-value</i> — Specifies the IP option value as a decimal hex or binary. <b>Values</b> 0..255 <i>ip-option-mask</i> — Specifies the IP option mask as a decimal hex or binary. <b>Values</b> 0..255

## multiple-option

<b>Syntax</b>	<b>multiple-option</b> {true   false}
<b>Context</b>	config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ip>entry>match config>subscr-mgmt>cat-map>category>exh-lvl>egr-ip>entry>match
<b>Description</b>	This command configures the multiple-option match condition.
<b>Parameters</b>	<b>true false</b> — Sets or resets the multiple option check.

## option-present

<b>Syntax</b>	<b>option-present</b> {true   false}
<b>Context</b>	config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ip>entry>match config>subscr-mgmt>cat-map>category>exh-lvl>egr-ip>entry>match
<b>Description</b>	This command configures the option-present match condition.

**Parameters** **true | false** — Sets or resets the option present check.

### src-ip

**Syntax** **src-ip** *{ip-address/mask | ip-address netmask}*  
**no src-ip**

**Context** config>subscr-mgmt>cat-map>category>exh-lvl>egr-ip>entry>match

**Description** This command configures the source IP match condition.

**Parameters** *ip-address/mask* — Specifies the IPv4 address and mask.

<b>Values</b>	ip-address	a.b.c.d
	mask	0 — 32

*netmask* — Specifies the mask, expressed as a dotted quad.

<b>Values</b>	a.b.c.d
---------------	---------

### src-port

**Syntax** **src-port** *{lt | gt | eq}* *src-port-number*  
**src-port range** *start end*  
**no src-port**

**Context** config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ip>entry>match  
config>subscr-mgmt>cat-map>category>exh-lvl>egr-ip>entry>match

**Description** This command configures the source port match condition.

**Parameters** *lt|gt|eq* — Specifies the operators.

*src-port-number* — Specifies the source port number as a decimal hex or binary.

<b>Values</b>	0..65535
---------------	----------

*dst-port-number* — Specifies the destination port number as a decimal hex or binary.

<b>Values</b>	0..65535
---------------	----------

### tcp-ack

**Syntax** **tcp-ack** *{true | false}*  
**no tcp-ack**

**Context** config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ip>entry>match  
config>subscr-mgmt>cat-map>category>exh-lvl>egr-ip>entry>match

**Description** This command configures the TCP ACK match condition. The **no tcp-ack** command disables the checking on the presence or absence of the tcp-ack flag.

**Parameters** **true|false** — True|false indicates that the entry will match on the presence resp. absence of the tcp-ack flag in the received packet. .

## tcp-syn

**Syntax** **tcp-syn {true | false}**  
**no tcp-syn**

**Context** config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ip>entry>match  
 config>subscr-mgmt>cat-map>category>exh-lvl>egr-ip>entry>match

**Description** This command configures the TCP SYN match condition. The **no** tcp-syn command disables the checking on the presence or absence of the tcp-syn flag.

**Parameters** **true|false** — True|false indicates that the entry will match on the presence resp. absence of the tcp-syn flag in the received packet.

## pir

**Syntax** **pir *pir-rate***  
**pir max**  
**no pir**

**Context** config>subscr-mgmt>cat-map>category>svc-lvl

**Description** This command configures the PIR which will be enforced for all queues pertaining to this category.

**Default** no pir

**Parameters** *pir-rate* — Specifies the amount of bandwidth in kilobits per second (thousand bits per second).

**Values** 1 — 40000000

**max** — Specifies to use the maximum amount of bandwidth.

## out-of-credit-action-override

**Syntax** **out-of-credit-action-override {continue | block-category | change-service-level}**  
**no out-of-credit-action-override**

**Context** config>subscr-mgmt>cat-map>category

**Description** This command specifies the action to be taken if the credit is exhausted.

**Default** no out-of-credit-action-override

**Parameters** **continue** — Specifies to continue when running out of credit.  
**block-category** — Specifies to block the category when running out of credit.  
**change-service-level** — Specifies to change the service level when running out of credit.

## policer

<b>Syntax</b>	<b>policer</b> <i>policer-id</i> { <b>ingress-only</b>   <b>egress-only</b>   <b>ingress-egress</b> } <b>no policer</b> <i>policer-id</i>
<b>Context</b>	config>subscr-mgmt>cat-map>category
<b>Description</b>	This command configures a policer in this category.
<b>Parameters</b>	<i>policer-id</i> — Specifies a policer identifier. The parameter <i>policer-id</i> references a <i>policer-id</i> that must be previously created within the SAP QoS policy.  <b>Values</b> 1 — 63  <b>ingress-only</b> — Specifies that ingress policers are defined in this category. <b>egress-only</b> — Specifies that egress policers are defined in this category. <b>ingress-egress</b> — Specifies that ingress and egress policers are defined in this category.

## queue

<b>Syntax</b>	<b>queue</b> <i>queue-id</i> { <b>ingress-only</b>   <b>egress-only</b>   <b>ingress-egress</b> } <b>no queue</b> <i>queue-id</i>
<b>Context</b>	config>subscr-mgmt>cat-map>category
<b>Description</b>	This command configures a queue in this category.
<b>Default</b>	none
<b>Parameters</b>	<i>queue-id</i> — Specifies the queue ID for this instances. Each queue nominated in the category map is monitored for activity (over a period of approximately 60 seconds), should the activity fall below the threshold value then a time is started. Whenever this timer exceeds the configured timeout under the idle-timeout the action (currently disconnect) is executed for that subscriber and all hosts under that given SLA-profile-instance.  <b>Values</b> 1 — 32  <b>ingress-only</b> — Specifies that ingress queues are defined in this category. <b>egress-only</b> — Specifies that egress queues are defined in this category. <b>ingress-egress</b> — Specifies that ingress and egress queues are defined in this category.

## rating-group

<b>Syntax</b>	<b>rating-group</b> <i>rating-group-id</i> <b>no rating-group</b>
<b>Context</b>	config>subscr-mgmt>cat-map>category
<b>Description</b>	This command configures the rating group applicable for this category.

<b>Default</b>	no rating group
<b>Parameters</b>	<i>rating-group-id</i> — Specifies the rating group applicable for this category.

## credit-exhaust-threshold

	<b>credit-exhaust-threshold</b> <i>threshold-percentage</i> <b>no credit-exhaust-threshold</b>
<b>Context</b>	config>subscr-mgmt>cat-map
<b>Description</b>	This command specifies the credit exhaust threshold taken into account to take action as depicted in <a href="#">Figure 113</a> . The <b>no</b> form of the command reverts the configured value to the default.
<b>Default</b>	100
<b>Parameters</b>	<i>threshold-percentage</i> — Specifies the percent to use for the credit exhaust threshold. <b>Values</b> 50 — 100

## credit-type

<b>Syntax</b>	<b>credit-type</b> { <b>volume</b>   <b>time</b> } <b>no credit-type</b>
<b>Context</b>	config>subscr-mgmt>cat-map
<b>Description</b>	This command specifies whether volume or time based accounting is performed.
<b>Default</b>	volume
<b>Parameters</b>	<b>volume</b> — specifies volume-based accounting. <b>time</b> — Specifies time-based accounting.

---

## Diameter Commands

### diameter-peer-policy

<b>Syntax</b>	<b>diameter-peer-policy</b> <i>policy-definition-name</i> <b>no diameter-peer-policy</b>
<b>Context</b>	configure>aaa
<b>Description</b>	This command creates a base diameter policy with up to 5 peers. There is a (TCP) connection created to each peer while only two peers can be active (used by applications) simultaneously. Various diameter applications can reference this policy.
<b>Default</b>	none
<b>Parameters</b>	<i>policy-definition-name</i> — Specifies the name of the policy that is created.

### diameter-application-policy

<b>Syntax</b>	<b>diameter-application-policy</b> <i>application-policy-name</i> [ <b>create</b> ] <b>no diameter-application-policy</b> <i>application-policy-name</i>
<b>Context</b>	configure>subscr-mgmt
<b>Description</b>	This command creates diameter application policy.
<b>Default</b>	none
<b>Parameters</b>	<i>application-policy-name</i> — Specifies the name of the diameter policy up to 32 characters in length.

### diameter-peer-policy

<b>Syntax</b>	<b>diameter-peer-policy</b> <i>referenced-policy-name</i> <b>no diameter-peer-policy</b>
<b>Context</b>	configure>subscr-mgmt>diam-app-pol
<b>Description</b>	This command is used by an application (DCCA, Gx, policy-management application, etc.) to reference a base diameter peer policy that the application will use.
<b>Default</b>	none
<b>Parameters</b>	<i>referenced-policy-name</i> — Specifies the name of the referenced policy.

### applications

<b>Syntax</b>	<b>applications</b> {[gx] [gy] [nasreq]}
---------------	--



**no application**

<b>Context</b>	configure>aaa>diam-peer-plcy
<b>Description</b>	<p>This command specifies which applications are advertised in the Capability Exchange Request (CER) messages sent on the peers.</p> <p>Applications that can be configured on a Diameter peer policy:</p> <ul style="list-style-type: none"> <li>• client and proxy role: <ul style="list-style-type: none"> <li>→ gx</li> <li>→ nasreq</li> <li>→ gx nasreq</li> </ul> </li> <li>• client role only: <ul style="list-style-type: none"> <li>→ gy</li> </ul> </li> </ul> <p><b>Note:</b> gx and nasreq applications can be enabled simultaneously on a single diameter peer.</p>
<b>Default</b>	none
<b>Parameters</b>	<p><b>gx</b> — Gx application support will be advertised in CER.</p> <p><b>gy</b> — Gy (DCCA) application support will be advertised in CER.</p> <p><b>nasreq</b> — NASREQ application support will be advertised in CER.</p>

## application

<b>Syntax</b>	<b>application {gx   gy   nasreq}</b> <b>no application</b>
<b>Context</b>	configure>aaa>diam-appl-pol
<b>Description</b>	<p>This command specifies the Diameter application for which this policy contains the configuration details, such as AVPs to include and their format.</p> <p>Applications are mutually exclusive.</p>
<b>Default</b>	none
<b>Parameters</b>	<p><b>gx</b> — This policy contains Gx application configuration options.</p> <p><b>gy</b> — This policy contains Gy application configuration options.</p> <p><b>nasreq</b> — This policy contains NASREQ application configuration options.</p>

## connection-timer

<b>Syntax</b>	<b>[no] connection-timer</b> <i>connection-time</i>
<b>Context</b>	configure>aaa>diam-peer-pol configure>aaa>diam-peer-pol>peer

## RADIUS Route Download Commands

<b>Description</b>	This command sets maximum amount of time the node attempts to reconnect to a diameter peer after a connection to the peer has been brought down due to a transport failure. There are certain exceptions to this rule, such as peer which terminated the transport connection indicating that it does not wish to communicate. A value of 0 means that the connection will not be retried. The configuration at peer level overrules the value configured at diameter-base level for the given peer.
<b>Default</b>	30 seconds at diameter-base level The default value at peer is taken from diameter-base.
<b>Parameters</b>	<i>connection-time</i> — Specifies the amount of time, in seconds. <b>Values</b> 1 — 1000

### origin-host

<b>Syntax</b>	<b>origin-host</b> <i>origin-host-string</i> <b>no origin-host</b>
<b>Context</b>	configure>aaa>diam-peer-pol
<b>Description</b>	This command configures the origin-realm AVP that will be sent in CER messages and all application based messages. Together with the Origin-Host AVP, these two AVPs form a Diameter Identity.
<b>Parameters</b>	<i>origin-host-string</i> — Specifies the Origin-Host AVP (Attribute Value Pair) used by this policy up to 80 characters in length.

### origin-realm

<b>Syntax</b>	<b>origin-realm</b> <i>origin-realm-string</i> <b>no origin-realm</b>
<b>Context</b>	configure>aaa>diam-peer-pol configure>aaa>diam-peer-pol>peer config>sub-mgmt>diameter-policy>diameter-base>peer
<b>Description</b>	This command configures the <i>origin-realm</i> AVP that will be sent in CER messages and all application based messages. Together with the Origin-Host AVP, these two AVPs form a Diameter Identity.
<b>Parameters</b>	<i>origin-realm-string</i> — Specifies the <i>origin-realm</i> AVP (Attribute Value Pair) used by this policy. up to 80 characters in length.

### peer

<b>Syntax</b>	<b>peer</b> <i>name</i> [create] <b>no peer</b> <i>name</i>
<b>Context</b>	configure>aaa>diam-peer-pol

<b>Description</b>	This command enables the context to configure diameter peer parameters. Up to five diameter peers can be defined inside of a diameter peer policy.
<b>Default</b>	none
<b>Parameters</b>	<i>name</i> — Specifies the peer name, up to a maximum of 32 characters.

## address

<b>Syntax</b>	<b>address</b> <i>ip-address</i> <b>no address</b>
<b>Context</b>	configure>aaa>diam-peer-pol>peer
<b>Description</b>	This command configures the address.
<b>Parameters</b>	<i>ip-address</i> — Specifies the IPv4 or IPv6 address of the diameter peer.

## destination-host

<b>Syntax</b>	<b>destination-host</b> <i>destination-host-string</i> <b>no destination-host</b>
<b>Context</b>	configure>aaa>diam-peer-pol>peer
<b>Description</b>	This command configures the destination-host AVP that will be sent in CCR-i/u and RAA messages. If the destination-host is not explicitly set via configuration, it will be learned from CCA or RAR messages. In other words, the origin-host received in the CCA or RAR message will be used to populate or replace the destination-host for the DCAA or GX session in 7x50.
<b>Parameters</b>	<i>destination-host-string</i> — Specifies the destination host name up 80 characters in length.

## preference

<b>Syntax</b>	<b>preference</b> <i>preference</i> <b>no preference</b>
<b>Context</b>	config>sub-mgmt>diameter-policy>diameter-base>peer config>sub-mgmt>diameter-policy>diameter-base configure>aaa>diam-peer-pol>peer
<b>Description</b>	This command configured preference per peer. Only the two peers with the highest preference in the peer table are considered for use (primary and secondary). Other peers can be the Open state and they just run keepalives (watchdog-request/answer messages). Once the primary peer fails, the secondary peer will be used as long as the last transaction on it has succeeded (stickiness). Another peer in the Open state will become secondary. Load balancing between peers is not supported.  The <b>no</b> form of the command reverts to the default value.
<b>Default</b>	none

**Parameters**     *preference* — Specifies the preference of this DIAMETER policy peer.  
**Values**            1 — 100

### transaction-timer

**Syntax**            **transaction-timer** *seconds*  
                  **no transaction-timer**

**Context**            configure>aaa>diam-peer-pol  
                  configure>aaa>diam-peer-pol>peer

**Description**        This command defines the time-out value for the Base Diameter messages (DWR, CER, DPR). Once the transaction-timer expires, an appropriate action will be taken for each message type.

                  This timer is used in the following cases:

- Opening the TCP connection (and completing the 3-way handshake) - if the TCP ACK is not received within the time specified by the transaction-timer, the TCP connection is closed and the connection-timer is started waiting for the new connection to be initiated.
- Capability Exchange – if the response to the CER message (CEA) is not received within the time specified by the transaction-timer, the peer connection is closed and the connection-timer is started waiting for the new connection to be initiated.
- Peer disconnect Request- if the response to the DPR message is not received (DPA) within the time specified by the transaction-timer, the peer connection is closed.
- DWR Timeout - if the response to the DWR message is not received (DWA) within the time specified by the transaction-timer, the peer connection is NOT closed. Instead the peer will transition into a peer suspended mode and at the same time the watchdog timer is restarted.

**Default**            none

**Parameters**        *seconds* — Specifies the policy peer transaction timer value in seconds.  
**Values**            1 — 1000

### transport

**Syntax**            **transport tcp** *port port*  
                  **no transport**

**Context**            configure>aaa>diam-peer-pol>peer

**Description**        This command defines source tcp port of the connection channel. Only TCP transport is currently supported

**Default**            3868

**Parameters**        **port** *port* — Specifies the transport protocol port number used towards this policy peer.  
**Values**            1 — 65535

## destination-realm

<b>Syntax</b>	<b>destination-realm</b> <i>destination-realm-string</i> <b>no destination-realm</b>
<b>Context</b>	configure>aaa>diam-peer-pol>peer
<b>Description</b>	This command configures the destination-realm AVP that will be sent in CCR-i/u and RAA messages. The Destination-Realm cannot be learned dynamically from the CCA or RAR messages and therefore it should be explicitly configured in 7x50. Once configured, it cannot be changed while peers are open.
<b>Parameters</b>	<i>destination-realm-string</i> — Specifies the destination realm name, maximum 80 displayable characters.

## watchdog-timer

<b>Syntax</b>	<b>watchdog-timer</b> <i>seconds</i> <b>no watchdog-timer</b>
<b>Context</b>	configure>aaa>diam-peer-pol configure>aaa>diam-peer-pol>peer
<b>Description</b>	<p>This command configures the interval between consecutive watchdog messages.</p> <p>On the first timeout of the DWR, 7x50 will resend the DWR message. The peer is still operation during this time.</p> <p>On the second timeout, the peer will transition into a suspended mode and the peer-failover procedure will be initiated (if the peer-failover is enabled via configuration). In this state the peer is not used for new transactions. At the same time, the cooldown procedure is started which means that it would take 3 successful DWR/DWA message exchanges to re-instate the peer in a fully operation state.</p> <p>On the third timeout, the peer is removed and its connection is closed.</p> <p>This behavior is described in RFC 3539, §3.4.1)</p>
<b>Default</b>	30
<b>Parameters</b>	<i>seconds</i> — specifies the device watchdog timer in seconds used by this policy peer.
<b>Values</b>	1 — 1000

## python-policy

<b>Syntax</b>	<b>python-policy</b> [32 chars max] <b>no python-policy</b>
<b>Context</b>	configure>aaa>diam-peer-pol
<b>Description</b>	This command specified the python-policy for Diameter messages received or transmitted on the Diameter peers defined in the diameter-peer-policy.

## RADIUS Route Download Commands

<b>Default</b>	none
<b>Parameters</b>	<i>name</i> — Specifies the name of the Python policy, up to 32 characters long.

### router

<b>Syntax</b>	<b>router</b> <i>router-instance</i> <b>router service</b> <i>service-name</i> <b>no router</b>
<b>Context</b>	configure>aaa>diam-peer-pol
<b>Description</b>	This command references the routing instance from which diameter peering is instantiated. <i>router-instance</i> — Specify one of the following parameters for the router instance: <i>router-name</i> — Specifies a router name up to 32 characters to be used in the match criteria. <b>Values</b> Base, management <b>Default</b> Base <i>service-id</i> — Specifies an existing service ID to be used in the match criteria. <b>Values</b> 1 — 2147483647 <b>service-name</b> <i>service-name</i> — Specifies an existing service name up to 64 characters in length.

### source-address

<b>Syntax</b>	<b>source-address</b> <i>ip-address</i> <b>no source-address</b>
<b>Context</b>	configure>aaa>diam-peer-pol
<b>Description</b>	This command configures the IPv4 source-address of all diameter messages sent to peers.
<b>Parameters</b>	<i>ip-address</i> — The IP prefix for the IP match criterion in dotted decimal notation. <b>Values</b> 0.0.0.0 — 255.255.255.255

### vendor-support

<b>Syntax</b>	<b>vendor-support</b> [three-gpp   vodafone] <b>no vendor-support</b>
<b>Context</b>	config>subscr-mgmt>diam-appl-plcy>gy config>aaa>diam-peer-plcyconfig
<b>Description</b>	In a diameter peer policy, this command specifies the vendor support announced in the capability exchange. In a Gy diameter application policy, this command specifies the vendor specific attributes for the user sessions.

The **no** form of the command reverts to the default value.

<b>Default</b>	three-gpp
<b>Parameters</b>	<b>three-gpp</b> — Specifies the 3GPP diameter policy vendor type. <b>vodafone</b> — Specifies the vodafone diameter policy vendor type.

## include-avp

<b>Syntax</b>	<b>[no] include-avp</b>
<b>Context</b>	config>subscr-mgmt>diam-appl-plcy>gy config>subscr-mgmt>diam-appl-plcy>gx config>subscr-mgmt>diam-appl-plcy>nasreq
<b>Description</b>	This command enables the context to configure AVPs and their format to be included in Diameter Gx, Gy or NASREQ application messages.

## an-gw-address

<b>Syntax</b>	<b>[no] an-gw-address</b>
<b>Context</b>	config>subscr-mgmt>diam-appl-plcy>gx>include-avp
<b>Description</b>	This command configures the IPv4 address of the 7x50.

## called-station-id

<b>Syntax</b>	<b>[no] called-station-id</b>
<b>Context</b>	config>subscr-mgmt>diam-appl-plcy>gx>include-avp config>subscr-mgmt>diam-appl-plcy>nasreq>avp
<b>Default</b>	no called-station-id
<b>Description</b>	This command configures the MAC address of AP in WiFi.

## calling-station-id

<b>Syntax</b>	<b>calling-station-id [type {llid   mac   remote-id   sap-id   sap-string}]</b> <b>no calling-station-id</b>
<b>Context</b>	config>subscr-mgmt>diam-appl-plcy>gx>include-avp config>subscr-mgmt>diam-appl-plcy>nasreq>avp
<b>Description</b>	This command includes the calling-station-id AVP in the specified format.
<b>Default</b>	no calling-station-id

## RADIUS Route Download Commands

<b>Parameters</b>	<b>type</b> — Specifies the format of the Calling-Station-ID AVP.
<b>Values</b>	<b>llid</b> — The LLID (logical link identifier) is the mapping from a physical to logical identification of a subscriber line and supplied by a RADIUS llid-server. <b>mac</b> — Specifies that the mac-address will be sent. <b>remote-id</b> — Specifies that the remote-id will be sent. <b>sap-id</b> — Specifies that the sap-id will be sent. <b>sap-string</b> — Specifies that the value is the inserted value set at the SAP level. If no calling-station-id value is set at the SAP level, the calling-station-id attribute will not be sent.

### circuit-id

<b>Syntax</b>	<b>[no] circuit-id</b>
<b>Context</b>	config>subscr-mgmt>diam-appl-plcy>nasreq>avp
<b>Description</b>	This command includes the Agent-Circuit-Id AVP.

### ip-can-type

<b>Syntax</b>	<b>[no] ip-can-type</b>
<b>Context</b>	config>subscr-mgmt>diam-appl-plcy>gx>include-avp
<b>Description</b>	This command includes the ip-can-type.

### logical-access-id

<b>Syntax</b>	<b>[no] logical-access-id</b>
<b>Context</b>	config>subscr-mgmt>diam-appl-plcy>gx>include-avp
<b>Description</b>	This command includes the logical-access-id.

### nas-port

<b>Syntax</b>	<b>nas-port</b> <i>binary-spec</i> <b>no nas-port</b>
<b>Context</b>	config>subscr-mgmt>diam-appl-plcy>gx>avp config>subscr-mgmt>diam-appl-plcy>nasreq>avp
<b>Description</b>	This command specifies the format of the 32 bit string used as value for the Nas-Port AVP.
<b>Default</b>	no nas-port



<b>Parameters</b>	<i>binary-spec</i> — Specifies the NAS-Port AVP format.	
	<b>Values</b>	
	binary-spec	<bit-specification> <binary-spec>
	bit-specification	0   1   <bit-origin>
	bit-origin	*<number-of-bits><origin>
	number-of-bits	1 — 32
	origin	s   m   p   o   i   v   c
		s - slot number
		m - MDA number
		p - port number or lag-id
		o - outer VLAN ID
		i - inner VLAN ID
		v - ATM VPI
		c - ATM VCI

## nas-port-id

<b>Syntax</b>	<b>nas-port-id</b> [ <b>prefix-type</b> { <b>none</b>   <b>user-string</b> }] [ <b>prefix-string</b> <i>prefix-string</i> ] [ <b>suffix-type</b> { <b>circuit-id</b>   <b>none</b>   <b>remote-id</b>   <b>user-string</b> }] [ <b>suffix-string</b> <i>suffix-string</i> ] <b>no nas-port-id</b>
<b>Context</b>	config>subscr-mgmt>diam-appl-plcy>gx>avp config>subscr-mgmt>diam-appl-plcy>nasreq>avp
<b>Description</b>	This command includes the Nas-Port-Id AVP.
<b>Default</b>	no nas-port-id
<b>Parameters</b>	<p><b>prefix-type</b> — Specifies what type of prefix will be added to the NAS-Port-Id attribute if included in Nas-Port-Id AVP messages.</p> <p><b>Values</b></p> <p><b>none</b> — No prefix is added.</p> <p><b>user-string</b> — Specifies the user configurable string to be added as prefix to the NAS-Port-Id attribute if included in DIAMETER Gx messages.</p> <p><i>prefix-string</i> — Specifies the user configurable string to be added as a prefix.</p> <p><b>suffix-type</b> — specifies the suffix to be added to the NAS-Port attribute NAS-Port AVP.</p> <p><b>Values</b></p> <p><b>one</b> — No suffix is added.</p> <p><b>circuit-id</b> — the circuit-id is added as suffix-string.</p> <p><b>remote-id</b> — the remote-id is added as suffix-string.</p> <p><b>user-string</b> — a user configurable suffix-string is added.</p> <p><i>suffix-string</i> — Specifies the string to be added as suffix. Max. 64 characters.</p>

## nas-port-type

<b>Syntax</b>	<b>nas-port-type</b> <b>nas-port-type</b> [ [0..255] ] <b>no nas-port-type</b>
---------------	--

## RADIUS Route Download Commands

<b>Context</b>	config>subscr-mgmt>diam-appl-plcy>gx>avp config>subscr-mgmt>diam-appl-plcy>nasreq>include-avp
<b>Description</b>	This command includes the Nas-Port-Type AVP.
<b>Default</b>	no nas-port-type
<b>Parameters</b>	<b>none</b> — Values as defined in RFC 2865, <i>Remote Authentication Dial In User Service (RADIUS)</i> , and RFC 4603, <i>Additional Values for the NAS-Port-Type Attribute</i> . <b>0..255</b> — Specifies the integer value between 0..255 for the Nas-Port-Type AVP.

### remote-id

<b>Syntax</b>	<b>[no] remote-id</b>
<b>Context</b>	config>subscr-mgmt>diam-appl-plcy>gx>include-avp config>subscr-mgmt>diam-appl-plcy>nasreq>include-avp
<b>Description</b>	This command enables the generation of the agent-remote-id for RADIUS.

### physical-access-id

<b>Syntax</b>	<b>[no] physical-access-id</b>
<b>Context</b>	config>subscr-mgmt>diam-appl-plcy>gx>include-avp
<b>Description</b>	This command includes the physical access ID.

### rat-type

<b>Syntax</b>	<b>[no] rat-type</b>
<b>Context</b>	config>subscr-mgmt>diam-appl-plcy>gx>include-avp
<b>Description</b>	This command includes the RAT type.

### supported-features

<b>Syntax</b>	<b>[no] supported-features</b>
<b>Context</b>	config>subscr-mgmt>diam-appl-plcy>gx>include-avp
<b>Description</b>	This command includes the supported-features.

## user-equipment-info

<b>Syntax</b>	<b>user-equipment-info</b> [type <i>ue-info-type</i> ] <b>no user-equipment-info</b>
<b>Context</b>	config>subscr-mgmt>diam-appl-plcy>gx>include-avp
<b>Description</b>	This command includes the user-equipment-info.

## mac-format

<b>Syntax</b>	<b>mac-format</b> <i>mac-format</i> <b>no mac-format</b>
<b>Context</b>	config>subscr-mgmt>diam-appl-plcy>gx config>subscr-mgmt>diam-appl-plcy>nasreq
<b>Description</b>	This command configures the format of the MAC address when reported in Gx or NASREQ application message AVPs such as Calling-Station-Id or User-Name.
<b>Default</b>	mac-format "aa:"
<b>Parameters</b>	<i>mac-format</i> — Specifies the MAC address format.
<b>Values</b>	like aa: 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

## report-ip-address-event

<b>Syntax</b>	<b>[no] report-ip-address-event</b>
<b>Context</b>	config>subscr-mgmt>diam-appl-plcy>gx
<b>Description</b>	<p>This command enables triggered CCR-u messages based on IP address allocation/de-allocation for the subscriber-host.</p> <p>In case that the requests for both IP address families (IPv4 and IPv6) arrive at approximately the same time, a single CCR-i will be sent containing the IP addresses from both address families - IPv4 and IPv6 (NA, PD or SLAAC). Otherwise, in case that the requests for IP addresses are not nearly simultaneous, the CCR-i will contain only the IP address that was allocated first (the one that triggered the session creation). The request for the second IP address family will, depending on configuration, trigger an additional CCR-u that will carry the IP address allocation update to the PCRF along with the UE_IP_ADDRESS_ALLOCATE (18) event. Apart from that, the CCR-u content should mirror the content of the CCR-i with exception of already allocated IP address(es).</p> <p>In case that this command is disabled, IP address triggered CCR-u messages will not be sent.</p>
<b>Default</b>	report-ip-addr-event (enabled)

### 3gpp-imsi

<b>Syntax</b>	<b>3gpp-imsi {circuit-id imsi subscriber-id}</b> <b>no 3gpp-imsi</b>
<b>Context</b>	config>subscr-mgmt>diam-appl-plcy>gy>include-avp
<b>Description</b>	This command specifies the origin of the information to send in the DCCA IMSI AVP. The no form of the command reverts to the default value.
<b>Default</b>	subscriber-id
<b>Parameters</b>	<b>circuit-id</b> — Specifies the circuit-id as DCCA IMSI AVP value. <b>subscriber-id</b> — Specifies the subscriber-id as DCCA IMSI AVP value. <b>imsi</b> — Specifies the imsi as DCCA IMSI AVP value.

### called-station-id

<b>Syntax</b>	<b>called-station-id [64 chars max]</b> <b>no called-station-id</b>
<b>Context</b>	config>subscr-mgmt>diam-appl-plcy>gy>include-avp
<b>Description</b>	This command configures the value of the called station ID AVP.
<b>Default</b>	no called-station-id
<b>Parameters</b>	<i>64 chars max</i> — Specifies the called station ID up to 64 characters.

### radius-user-name

<b>Syntax</b>	<b>[no] radius-user-name</b>
<b>Context</b>	config>subscr-mgmt>diam-appl-plcy>gy>include-avp
<b>Description</b>	This command includes the RADIUS user name AVP in the Diameter gy messages.
<b>Default</b>	no radius-user-name

### service-context-id

<b>Syntax</b>	<b>service-context-id name</b> <b>no service-context-id</b>
<b>Context</b>	config>subscr-mgmt>diam-appl-plcy>gy>include-avp
<b>Description</b>	This command configure the value of the service context ID AVP.
<b>Default</b>	no service-context-id

**Parameters** *name* — Specifies the service context ID AVP value up to 32 displayable characters.

## preference

**Syntax** **preference** *preference*  
**no preference**

**Context** configure>aaa>diam-peer-pol>peer

**Description** This command configures the preference given to this policy peer with respect to the other peers associated with this policy.

If multiple peers are available for this policy, only the available peer with the highest preference will be used.

If multiple peers with the same preference are available, one of them will be used.

The **no** form of the command reverts to the default value.

**Default** 50

**Parameters** *preference* — Specifies the preference of this policy peer.

**Values** 1 — 100

## transaction-timer

**Syntax** [**no**] **transaction-timer** *transatcion-time*

**Context** configure>aaa>diam-peer-pol

**Description** This command sets maximum amount of time the node waits for a diameter peer to respond before trying another peer. The configuration at peer level overrules the value configured at diameter-base level for the given peer.

**Default** 30 seconds at diameter-base level

Default value at peer is taken from diameter-base.

**Parameters** *transaction* — Specifies the DIAMETER peer policy transaction timer in seconds.

**Values** 1-1000

## router

**Syntax** **router service** *service-name*  
**router** *router-instance*  
**no router**

**Context** config>sub-mgmt>diameter-policy>diameter-base

## RADIUS Route Download Commands

<b>Description</b>	This command specifies the virtual router in which the diameter connection(s) will be established by this diameter policy.
<b>Parameters</b>	<i>router-instance</i> — Specifies the router name.  <b>Values</b> router-instance: <i>router-name service-id</i> router-name: Base, management service-id: 1 — 2147483647  <b>Default</b> Base <i>service-name</i> — Specifies the VPRN service ID.

## source-address

<b>Syntax</b>	<b>source-address</b> <i>ip-address</i> <b>no source-address</b>
<b>Context</b>	config>sub-mgmt>diameter-policy>diameter-base
<b>Description</b>	This command configures the source address.
<b>Default</b>	no source-address; system-ip address is used instead
<b>Parameters</b>	<i>ip-address</i> — Specifies the UC IPv4 or IPv6 IP address.

## gx

<b>Syntax</b>	<b>gx</b>
<b>Context</b>	config>sub-mgmt>diameter-policy>diameter-base
<b>Description</b>	This command enables the context to configure Gx parameters.

## gy

<b>Syntax</b>	<b>gy</b>
<b>Context</b>	config>sub-mgmt>diameter-policy
<b>Description</b>	This command enables the context to configure Diameter Credit Control Application or Gy-specific options.

## nasreq

<b>Syntax</b>	<b>nasreq</b>
<b>Context</b>	config>sub-mgmt>diameter-policy

**Description** This command enables the context to configure NASREQ application-specific attributes.

## avp-subscription-id

<b>Syntax</b>	<b>avp-subscription-id origin [type type]</b> <b>no avp-subscription-id</b>
<b>Context</b>	config>subscr-mgmt>diam-appl-plcy>gx config>subscr-mgmt>diam-appl-plcy>gy
<b>Description</b>	This command is used to provide identification information to the PCRF for the end user. Subscription-id is a grouped AVP. In case that parameter designated to be the subscription-id is not available, the subscription-avp will not be sent.  The <b>no</b> form of the command reverts to the default value.
<b>Default</b>	none
<b>Default</b>	<b>avp-subscription-id subscriber-id type private</b>
<b>Parameters</b>	<b>origin</b> — Specifies the origin of the information to send in the Subscription-Id-Data AVP.  <div> <b>Values</b> <ul style="list-style-type: none"> <li><b>circuit-id</b> — The circuit ID.</li> <li><b>dual-stack-remote-id</b> — The remote-id for IPv4 and IPv6. The enterprise-id field is stripped off from IPv6 remote-id before it is passed to the PCRF in Gx message.</li> <li><b>imei</b> — The physical ID of the end device.</li> <li><b>imsi</b> — The SIM ID.</li> <li><b>mac</b> — The MAC address of the end device.</li> <li><b>msisdn</b> — The phone number of the end device.</li> <li><b>nas-port-id</b> — nas-port-id can be a prefix or suffix with a custom string to make it unique network wide.</li> <li><b>subscriber-id</b> — The subscriber ID.</li> <li><b>username</b> — The username identifier can be of type <b>private</b> or <b>nai</b>. The username is a ppp-username (PAP/CHAP). In case that ppp-username is not available, the string in the Username attribute returned via RADIUS or NASREQ will be used.</li> </ul> </div> <b>type</b> — Specifies the type of the identifier stored in the Subscription-Id-Data AVP.  <div> <b>Values</b> <ul style="list-style-type: none"> <li><b>e164</b> — The identifier is in international E.164 format (e.g., MSISDN).</li> <li><b>imsi</b> — The identifier is in international IMSI format according to the ITU-T E.212 numbering plan.</li> <li><b>nai</b> — The identifier is in the form of a Network Access Identifier as defined in RFC 2486.</li> <li><b>private</b> — The identifier is a private type identifier.</li> </ul> </div>

## out-of-credit-reporting

<b>Syntax</b>	<b>out-of-credit-reporting {final quota-exhausted}</b> <b>no out-of-credit-reporting</b>
<b>Context</b>	config>subscr-mgmt>diam-peer-plcy>gy

<b>Description</b>	<p>This command changes the reporting reason in an intermediate interrogation when the final granted units have been consumed and a corresponding out-of-credit-action different from "disconnect-host" is started.</p> <p>The no form of the command reverts to the default value</p>
<b>Default</b>	out-of-credit-reporting final
<b>Parameters</b>	<p><b>final</b> — Specifies the reporting reason in an intermediate interrogation when the final granted units have been consumed and a corresponding out-of-credit-action different from <b>disconnect-host</b> is started.</p> <p><b>quota-exhausted</b> — Specifies the reporting reason in an intermediate interrogation when the final granted units have been consumed and a corresponding out-of-credit-action different from <b>disconnect-host</b> is started.</p>

### on-failure

<b>Syntax</b>	<b>on-failure [failover {enabled disabled}] [handling {continue   retry-and-terminate   terminate}]</b> <b>no on-failure</b>
<b>Context</b>	config>subscr-mgmt>diam-peer-plcy
<b>Description</b>	<p>Behavior of the application's session in case of a peer failure can be controlled by the Diameter server through two AVPs carried in CCA messages that are defined in RFC4006:</p> <ul style="list-style-type: none"><li>• CC-Session-Failover AVP<ul style="list-style-type: none"><li>→ FAILOVER_NOT_SUPPORTED</li><li>→ FAILOVER_SUPPORTED</li></ul></li><li>• Credit-Control-Failure-Handling AVP<ul style="list-style-type: none"><li>→ TERMINATE</li><li>→ CONTINUE</li><li>→ RETRY_AND_TERMINATE</li></ul></li></ul>

In case that those AVPs are not provided by the Diameter server, the local configuration provided by this command will take effect. This command defines the following:

- Peer-failover behavior to a secondary peer in case that the primary peer is unresponsive. The primary peer is considered unresponsive in case that the application message sent to it, times out. The failover mechanism defined by this command is only applicable to CCR messages (and not to RAA messages since there is no response expected). The time out of the message is determined by the tx-timer command.

The peer-failover action based on the message timeout is defined per session. In other words, a message timeout for one session cannot cause the failover for some other session.

The maximum number of transmissions per session is hardcoded to 2 and the same message is never re-transmitted to the same TCP socket (a TCP socket is defined as a current peering connection defined by the TCP source/destination IP addresses/ports; closing and then reopening a connection to the same peer will result in creation of a new TCP socket). Once the original message for the session times out on the primary peer, the message will be re-transmitted to the



secondary peer, provided that the secondary peer is available and the failover is enabled with the corresponding handling mechanism. In case that the secondary peer is unavailable, the original message will not be re-transmitted to the same primary peer again.

Once the reply from a peer is received, the session will be tied to that peer until the next timeout. In other words, the session always sticks to the peer from which it received the last response.

- Handling behavior in case that the response from the peer is not received or the peers are not available at all (all peering connections are closed). This behavior is applicable to CCR-i messages in Gx and CCR-i/u messages in Gy. In case of Gx, if the response to a session initiation message (CCR-i) is not received, the fate of the session will depend on the configuration (the session can be terminated or continue to exist with default parameters).

<b>Default</b>	on-failure failover enabled handling terminate
<b>Parameters</b>	<p><b>failover enabled</b> — The session is allowed to switch to the secondary peer.</p> <p><b>failover disabled</b> — The session is NOT allowed to switch to the secondary peer.</p> <p><b>handling continue</b> — The sessions will continue to exist if the response to a transmitted CCR message is not received. Whether the transmitted message will be re-transmitted depends on the failover configuration. In case of session initiation procedure in the Gx case (CCR-i timeout), the subscriber host will be instantiated with the default parameters, assuming that they are provided. In the default parameter are not provided, the subscriber host initiation will fail.</p> <p><b>handling retry-and-terminate</b> — The message will be re-transmitted in case that the peer-failover is enabled and the secondary peer is available. Once the retransmitted message (CCR-i in Gx; CCR-i/u in Gy) is timed-out, the application session will be terminated.</p> <p><b>handling terminate</b> — The session will be terminated if the response to the original message (CCR-I in Gx; CCR-i/u in Gy) is not received. No re-transmissions will be attempted, regardless of whether the failover is enabled or not.</p>

## tx-timer

<b>Syntax</b>	<b>tx-timer seconds</b> <b>no tx-timer</b>
<b>Context</b>	configure>subscr-mgmt>diam-app-pol
<b>Description</b>	<p>This command defines the time-out period for the application's CCR-i/u messages that are waiting for a reply from a peer (message is in a pending state). Peer-failover behavior determines the action that will be taken once the message times out. Peer-failover behavior can be dictated by the PCRF or can be locally configured in 7x50.</p> <p>Per RFC 4006, sec 13, <i>Diameter Credit-Control Application, Credit-Control Application Related Parameters</i>, When real-time credit-control is required, the credit-control client contacts the credit-control server before and while the service is provided to an end user. Due to the real-time nature of the application, the communication delays SHOULD be minimized; e.g., to avoid an overly long service setup time experienced by the end user. The Tx timer is introduced to control the waiting time in the client in the Pending state. When the Tx timer elapses, the credit-control client takes an action to the end user according to the value of the Credit-Control-Failure-Handling AVP or Direct-Debiting-Failure-Handling AVP. The recommended value is 10 seconds.</p>

## RADIUS Route Download Commands

<b>Default</b>	10
<b>Parameters</b>	<i>seconds</i> — specifies the Tx Timer value (in seconds) for this policy.
<b>Values</b>	10 — 1000

### diameter-application-policy

<b>Syntax</b>	<b>diameter-application-policy</b> <i>application-policy-name</i> <b>no diameter-application-policy</b>
<b>Context</b>	configure>service>vpls>sap configure>service>vpn>sub-if>grp-if configure>service>ies>sub-if>grp-if configure>subscr-mgmt>loc-user-db>ipoe>host configure>subscr-mgmt>loc-user-db>pppoe>host
<b>Description</b>	This command associates the specified diameter-application-policy with the processing of the host attachment requests.
<b>Default</b>	none
<b>Parameters</b>	<i>application-policy-name</i> — Specifies the name of the diameter policy up to 32 characters in length.

---

## Filter Commands

### filter

<b>Syntax</b>	<b>filter</b>
<b>Context</b>	configure
<b>Description</b>	This command manages the configuration of filters.

### copy

<b>Syntax</b>	<b>copy {ip-filter   mac-filter   ipv6-filter} <i>src-filter-id</i> [src-entry <i>src-entry-id</i>] to <i>dst-filter-id</i> [dst-entry <i>dst-entry-id</i>] [overwrite]</b>
<b>Context</b>	configure>filter
<b>Description</b>	This command copies filters and its entries.
<b>Parameters</b>	<p><i>src-filter-id</i> — Specifies the source filter ID.</p> <p><b>Values</b> 1..65535</p> <p><i>src-entry-id</i> — Specifies the source entry ID.</p> <p><b>Values</b> 1..65535</p> <p><i>dst-filter-id</i> — Specifies the destination filter ID.</p> <p><b>Values</b> 1..65535</p> <p><i>dst-entry-id</i> — Specifies the destination entry ID.</p> <p><b>Values</b> 1..65535</p> <p><b>overwrite</b> — Specifies an overwrite.</p>

### ip-filter

<b>Syntax</b>	<b>ip-filter <i>filter-id</i> [create] no ip-filter <i>filter-id</i></b>
<b>Context</b>	configure>filter
<b>Description</b>	This command configures an IP filter.
<b>Parameters</b>	<p><i>filter-id</i> — Specifies the filter ID.</p> <p><b>Values</b> 1..65535</p>

## ipv6-filter

<b>Syntax</b>	<b>ipv6-filter</b> <i>ipv6-filter-id</i> [ <b>create</b> ] <b>no ipv6-filter</b> <i>ipv6-filter-id</i>
<b>Context</b>	configure>filter
<b>Description</b>	This command configures an IPv6 filter.
<b>Parameters</b>	<i>filter-id</i> — Specifies the filter ID.
<b>Values</b>	1..65535

## default-action

<b>Syntax</b>	<b>default-action</b> <i>drop</i>   <i>forward</i>
<b>Context</b>	configure>filter>ip-filter configure>filter>ipv6-filter
<b>Description</b>	This command configures default-action for the IP or IPv6 filter.
<b>Parameters</b>	<i>drop</i>   <i>forward</i> — This keyword specifies the filter action.

## entry

<b>Syntax</b>	<b>entry</b> <i>entry-id</i> [ <b>time-range</b> <i>time-range-name</i> ] [ <b>create</b> ] <b>no entry</b> <i>entry-id</i>
<b>Context</b>	configure>filter>ip-filter configure>filter>ipv6-filter
<b>Description</b>	This command configures an IP or IPv6 filter entry.
<b>Parameters</b>	<i>entry-id</i> — Specifies the entry ID.
<b>Values</b>	1..65535
	<i>time-range-name</i> — Specifies the time range name.
<b>Values</b>	32 charas max

## action

<b>Syntax</b>	<b>action</b> <i>drop</i>   <i>forward</i> <b>no action</b>
<b>Context</b>	config>filter>ip-filter>entry config>filter>ipv6-filter>entry
<b>Description</b>	This command configures actions for the IP or IPv6 filter entry.

**Parameters** *drop|forward* — Specifies the filter action.

## log

**Syntax** **log** *log-id*  
**no log**

**Context** config>filter>ip-filter>entry  
 config>filter>ipv6-filter>entry

**Description** This command configures the log for the IP or IPv6 filter entry.

**Parameters** *log-id* — Specifies the log ID.

**Values** 101..199

## match

**Syntax** **match** [**next-header** *next-header*]  
**no match**

**Context** config>filter>ip-filter>entry  
 config>filter>ipv6-filter>entry

**Description** This command configures the match criteria for the IP or IPv6 filter entry.

**Parameters** *next-header* — Specifies the protocol numbers accepted in DHB.

**Values** [1..42|45..49|52..29|61..255]

**Values** none | crtp | crudp | egp | eigrp | encap | ether-i p | gre | icmp | idrp | igmp | igp | ip |  
 ipv6 | ipv6-icmp | ipv6-no-nxt | isis | iso-ip | l2tp | ospf-igp | pim | pnni | ptp | rdp |  
 rsvp | stp | tcp | udp | vrrp \* udp/tcp wildcard

## dscp

**Syntax** [**no**] **dscp**

**Context** config>filter>ip-filter>entry>match  
 config>filter>ipv6-filter>entry>match

**Description** This command configures DSCP match condition.

## dst-ip

**Syntax** [**no**] **dst-ip**

**Context** config>filter>ip-filter>entry>match  
 config>filter>ipv6-filter>entry>match

## RADIUS Route Download Commands

**Description** This command configures the destination IP or IPv6 address match condition.

### dst-port

**Syntax** [no] dst-port

**Context** config>filter>ip-filter>entry>match  
config>filter>ipv6-filter>entry>match

**Description** This command configures the destination port match condition.

### icmp-code

**Syntax** [no] icmp-code

**Context** config>filter>ip-filter>entry>match  
config>filter>ipv6-filter>entry>match

**Description** This command configures the ICMP code match condition.

### icmp-type

**Syntax** [no] icmp-type

**Context** config>filter>ip-filter>entry>match  
config>filter>ipv6-filter>entry>match

**Description** This command configures the ICMP type match condition.

### src-ip

**Syntax** [no] src-ip

**Context** config>filter>ip-filter>entry>match  
config>filter>ipv6-filter>entry>match

**Description** This command configures the source IP or IPv6 address match condition.

### src-port

**Syntax** [no] src-port

**Context** config>filter>ip-filter>entry>match  
config>filter>ipv6-filter>entry>match

**Description** This command configures the source port match condition.

## tcp-ack

<b>Syntax</b>	<b>[no] tcp-ack</b>
<b>Context</b>	config>filter>ip-filter>entry>match config>filter>ipv6-filter>entry>match
<b>Description</b>	This command configures the TCP ACK match condition.

## tcp-syn

<b>Syntax</b>	<b>[no] tcp-syn</b>
<b>Context</b>	config>filter>ip-filter>entry>match config>filter>ipv6-filter>entry>match
<b>Description</b>	This command configures the TCP SYN match condition.

## group-inserted-entries

<b>Syntax</b>	<b>group-inserted-entries application <i>application</i> location <i>location</i></b>
<b>Context</b>	config>filter>ip-filter config>filter>ipv6-filter
<b>Description</b>	This command groups auto-inserted entries.
<b>Parameters</b>	<i>application</i> — Specifies the application. <div style="margin-left: 40px;"><b>Values</b>      radius   credit-control</div> <i>location</i> — Specifies the location. <div style="margin-left: 40px;"><b>Values</b>      top   bottom</div>

## renum

<b>Syntax</b>	<b>renum <i>old-entry-id</i> <i>new-entry-id</i></b>
<b>Context</b>	config>filter>ip-filter config>filter>ipv6-filter
<b>Description</b>	This command renumbers an IP or IPv6 filter entry.
<b>Parameters</b>	<i>old-entry-id</i> — Specifies the old entry ID to be renumbered. <div style="margin-left: 40px;"><b>Values</b>      1..65535</div> <i>new-entry-id</i> — Specifies the new entry ID. <div style="margin-left: 40px;"><b>Values</b>      1..65535</div>

### scope

<b>Syntax</b>	<b>scope exclusive   template</b> <b>no scope</b>
<b>Context</b>	config>filter>ip-filter config>filter>ipv6-filter
<b>Description</b>	This command configures the scope for the IP or IPv6 filter.
<b>Parameters</b>	<b>exclusive   template</b> — Specifies the type of policy.

### shared-radius-filter-wmark

<b>Syntax</b>	<b>shared-radius-filter-wmark low <i>low-watermark</i> high <i>high-watermark</i></b> <b>no shared-radius-filter-wmark</b>
<b>Context</b>	config>filter>ip-filter config>filter>ipv6-filter
<b>Description</b>	This command defines the thresholds that will be used to raise a respective alarm when the number of shared filter copies increases.
<b>Default</b>	no shared-radius-filter-wmark
<b>Parameters</b>	<i>low-watermark</i> — specifies low threshold for the number of shared filter copies <div><b>Values</b> 0-8000</div> <i>high-watermark</i> — specifies high threshold for the number of shared filter copies <div><b>Values</b> 0-8000</div>

### sub-insert-radius

<b>Syntax</b>	<b>sub-insert-radius start-entry <i>entry-id</i> count <i>count</i></b> <b>no sub-insert-radius</b>
<b>Context</b>	config>filter>ip-filter config>filter>ipv6-filter
<b>Description</b>	This command defines the range of filter entries which will be reserved for entries created based on information (match criteria and action) from RADIUS auth-response messages.  The <b>no</b> version of the command disables the insertion, which means that information from auth-response messages cannot be stored in the filter, and the corresponding host will not be created in the system.
<b>Default</b>	per default insertion is disabled
<b>Parameters</b>	<i>entry-id</i> — An integer defining the lowest entry of the range. <i>count</i> — An integer defining the number of entries in the range.



## sub-insert-credit-control

<b>Syntax</b>	<b>sub-insert-credit-control start-entry <i>entry-id</i> count <i>count</i></b> <b>no sub-insert-credit-control</b>
<b>Context</b>	config>filter>ip-filter config>filter>ipv6-filter
<b>Description</b>	This command defines the range of filter entries that will be reserved for entries created based on information (match criteria and action) configured under the category-map configuration tree to enforce reduced-service level in case of credit exhaustion.  The <b>no</b> version of the command disables the insertion, which means that entries will not be installed even though the credit for the given category and subscriber-host has been exhausted.
<b>Default</b>	per default insertion is disabled
<b>Parameters</b>	<i>entry-id</i> — An integer defining the lowest entry of the range. <i>count</i> — An integer defining the number of entries in the range.

## sub-insert-shared-radius

<b>Syntax</b>	<b>sub-insert-shared-radius start-entry <i>entry-id</i> count <i>count</i></b> <b>no sub-insert-shared-radius</b>
<b>Context</b>	config>filter>ip-filter config>filter>ipv6-filter config>filter>ip-filter config>filter>ipv6-filter
<b>Description</b>	This command defines the range of filter entries that will be reserved for shared filter entries received in RADIUS messages.  The no version of the command disables the insertion resulting in a host setup failure when shared filter attributes are received in a RADIUS authentication response.
<b>Default</b>	no sub-insert-shared-radius
<b>Parameters</b>	<i>entry-id</i> — specifies the lowest entry of the range. <b>Values</b> 1-65535 <i>count</i> — specifies the number of entries in the range. <b>Values</b> 1-65535

## sub-insert-wmark

<b>Syntax</b>	<b>sub-insert-wmark [<i>low percentage</i>] [<i>high percentage</i>]</b> <b>no sub-insert-wmark</b>
<b>Context</b>	config>filter>ip-filter

## RADIUS Route Download Commands

<b>Description</b>	This command defines the thresholds that will be used to raise a respective alarm to provide monitoring of the resources for subscriber-specific filter insertion.  The <b>no</b> version of the command sets the default values for the respective thresholds.
<b>Default</b>	low - 90%  high - 95%
<b>Parameters</b>	<i>percentage</i> — Defines in percentage the threshold used to raise an alarm.  <b>Values</b> 1-100, integer

## host-tracking-policy

<b>Syntax</b>	<b>host-tracking-policy</b> <i>policy-name</i> [ <b>create</b> ] <b>no host-tracking-policy</b> <i>policy-name</i>
<b>Context</b>	config>subscr-mgmt config>subscr-mgmt>sub-prof
<b>Description</b>	This command configures a host tracking policy. IGMP host tracking is an option in the subscriber profile that allows the factoring in of a subscriber's (multicast) video traffic by reducing the unicast operational egress aggregate rate or the rate of the scheduler specified in the ANCP policy to account for a subscriber's multicast traffic. If no ANCP policy is defined, the egress aggregate rate configured in the subscriber profile is reduced. If an ANCP policy is defined, the "rate-modify" parameter in the policy specifies whether the egress aggregate rate or the rate of the egress policer specified in the policy is to be reduced to account for the subscriber's multicast traffic.
<b>Default</b>	disabled

## egress-rate-modify

<b>Syntax</b>	<b>egress-rate-modify</b> <b>agg-rate-limit</b> <b>egress-rate-modify</b> <b>scheduler</b> <i>scheduler-name</i> <b>no egress-rate-modify</b>
<b>Context</b>	config>subscr-mgmt>trk-plcy
<b>Description</b>	This command specifies the egress-rate modification that is to be applied.  <b>agg-rate-limit</b> — Specifies the egress rate limit.  <b>scheduler</b> <i>scheduler-name</i> — Specifies the scheduler name to use.

---

## PIM Policy Commands

### pim-policy

<b>Syntax</b>	<b>pim-policy</b> <i>pim-policy-name</i> [ <b>create</b> ] <b>no pim-policy</b> <i>pim-policy-name</i>
<b>Context</b>	config>subscr-mgmt
<b>Description</b>	This command creates a PIM policy or enables the context to configure a PIM policy. The <b>no</b> form of this command deletes the specified PIM policy.
<b>Default</b>	none
<b>Parameters</b>	<i>pim-policy-name</i> — Specifies the PIM policy name.  <b>Values</b> Valid names consist of any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.  <b>create</b> — Keyword used to create the PIM policy. The <b>create</b> keyword requirement can be enabled/disabled in the <b>environment&gt;create</b> context.

---

## Managed SAP Policy Commands

### msap-policy

<b>Syntax</b>	<b>msap-policy</b> <i>msap-policy-name</i> [ <b>create</b> ] <b>no msap-policy</b> <i>msap-policy-name</i>
<b>Context</b>	config>subscr-mgmt
<b>Description</b>	This command configures a managed SAP policy. Managed SAPs allow the use of policies and a SAP template for the creation of a SAP.
<b>Default</b>	none
<b>Parameters</b>	<i>msap-policy-name</i> — Specifies the managed SAP policy name.  <div> <b>Values</b> Valid names consist of any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.   <b>create</b> — Keyword used to create the managed SAP policy. The <b>create</b> keyword requirement can be enabled/disabled in the <b>environment&gt;create</b> context. </div>

### cpu-protection

<b>Syntax</b>	<b>cpu-protection</b>
<b>Context</b>	config>sys>security
<b>Description</b>	This command enables the context to configure CPU protection policies.

### cpu-protection

<b>Syntax</b>	<b>cpu-protection</b> <i>policy-id</i> [ <b>mac-monitoring</b> ] <b>no cpu-protection</b>
<b>Context</b>	config>subscr-mgmt>msap-policy [ <b>mac-monitoring</b> ] config>service>vpls>sap [ <b>mac-monitoring</b> ]
<b>Description</b>	This command assigns an existing CPU protection policy to the SAP or interface. CPU protection policies are configured in the <b>config&gt;sys&gt;security&gt;cpu-protection</b> context. The <b>no</b> form of the command removes the policy ID from the SAP or interface configuration.
<b>Default</b>	none
<b>Parameters</b>	<i>policy-id</i> — Specifies an existing CPU protection policy to assign to the SAP or interface. <b>mac-monitoring</b> — Specifies that the per-source rate limit be applied.

## cpu-protection

<b>Syntax</b>	<b>cpu-protection</b> <i>policy-id</i> <b>no cpu-protection</b>
<b>Context</b>	config>router>if config>service>ies>if
<b>Description</b>	This command assigns an existing CPU protection policy to the SAP or interface. CPU protection policies are configured in the <b>config&gt;sys&gt;security&gt;cpu-protection</b> context. The <b>no</b> form of the command removes the policy ID from the SAP or interface configuration.
<b>Default</b>	none
<b>Parameters</b>	<i>policy-id</i> — Specifies an existing CPU protection policy to assign to the SAP.

## default-host

<b>Syntax</b>	<b>default-host</b> <i>ip-address/mask next-hop next-hop-ip</i> <b>no default-host</b>
<b>Context</b>	config>service>ies>sub-if>grp-if>sap config>service>vprn>sub-if>grp-if>sap
<b>Description</b>	This command configures the default-host to be used. More than one default-host can be configured per SAP. The <b>no</b> form of the command removes the values from the configuration.
<b>Parameters</b>	<i>ip-address/mask</i> — Assigns an IP address/IP subnet format to the interface. <i>next-hop next-hop-ip</i> — Assigns the next hop IP address.

## dist-cpu-protection

<b>Syntax</b>	<b>dist-cpu-protection</b> <i>policy-name</i> <b>no dist-cpu-protection</b>
<b>Context</b>	config>subscriber-management>msap-policy
<b>Description</b>	This command assigns a Distributed CPU Protection (DCP) policy to the MSAP policy. The DCP policy will automatically get assigned to any MSAPs created with this policy. A non-existent DCP policy can be assigned to an msap-policy since an msap-policy is effectively a template that gets applied at some point in the future during msap creation. The existence of the DCP policy will be validated at the time that the msap is created, and the msap creation will be blocked (and an appropriate log event created) if the DCP policy does not exist. Note that for other types of objects (for example, normal non-msap SAPs and network interfaces) the DCP policy must exist before it can be assigned to the SAP.
<b>Default.</b>	no dist-cpu-protection

## ies-vprn-only-sap-parameters

<b>Syntax</b>	<b>ies-vprn-only-sap-parameters</b>
<b>Context</b>	config>subscr-mgmt>msap-policy
<b>Description</b>	This command configures Managed SAP IES properties.

## igmp-host-tracking

<b>Syntax</b>	<b>igmp-host-tracking</b>
<b>Context</b>	config>subscr-mgmt>msap-policy
<b>Description</b>	This command enables the context to configure IGMP host tracking parameters.

## expiry-time

<b>Syntax</b>	<b>expiry-time</b> <i>expiry-time</i> <b>no expiry-time</b>
<b>Context</b>	config>subscr-mgmt>msap-policy>igmp-host-tracking
<b>Description</b>	This command configures the time that the system continues to track inactive hosts. The <b>no</b> form of the command removes the values from the configuration.
<b>Default</b>	no expiry-time
<b>Parameters</b>	<i>expiry-time</i> — Specifies the time, in seconds, that this system continues to track an inactive host. <b>Values</b> 1 — 65535

## import

<b>Syntax</b>	<b>import</b> <i>policy-name</i> <b>no import</b>
<b>Context</b>	config>subscr-mgmt>msap-policy>igmp-host-tracking
<b>Description</b>	This command specifies the import routing policy to be used for IGMP packets to be used on this SAP or SDP. Only a single policy can be imported on a single SAP at any time. The <b>no</b> form of the command removes the policy association from the SAP or SDP.
<b>Default</b>	no import (No import policy is specified)
<b>Parameters</b>	<i>policy-name</i> — The routing policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters excluding double quotes. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

Routing policies are configured in the config>router>policy-options context. The router policy must be defined before it can be imported.

## max-num-group

<b>Syntax</b>	<b>max-num-groups</b> <i>max-num-groups</i> <b>no max-num-groups</b>
<b>Context</b>	config>subscr-mgmt>msap-policy>igmp-host-tracking
<b>Description</b>	This command configures the maximum number of multicast groups allowed to be tracked. The <b>no</b> form of the command removes the values from the configuration.
<b>Default</b>	no max-num-groups
<b>Parameters</b>	<i>max-num-groups</i> — Specifies the maximum number of multicast groups allowed to be tracked.
<b>Values</b>	1 — 196607

## max-num-sources

<b>Syntax</b>	<b>max-num-sources</b> <i>max-num-sources</i> <b>no max-num-sources</b>
<b>Context</b>	config>subscr-mgmt>msap-policy>igmp-host-tracking
<b>Description</b>	This command configures the maximum number of multicast sources allowed to be tracked per group. The no form of the command removes the value from the configuration.
<b>Parameters</b>	<i>max-num-sources</i> — Specifies the maximum number of multicast sources allowed to be tracked per group.
<b>Values</b>	1 — 1000

## max-num-grp-sources

<b>Syntax</b>	<b>max-num-grp-sources</b> [1..32000] <b>no max-num-grp-sources</b>
<b>Context</b>	config>subscr-mgmt>msap-policy>igmp-host-tracking
<b>Description</b>	This command configures the maximum number of group sources for which IGMP can have local receiver information based on received IGMP reports on this interface. When this configuration is changed dynamically to a value lower than currently accepted number of group sources, the group sources that are already accepted are not deleted. Only new group sources will not be allowed. When this object has a value of 0, there is no limit to the number of group sources.  The <b>no</b> form of the command removes the value from the configuration.
<b>Default</b>	no max-num-grp-sources

**Parameters**     **1..32000** — Specifies the maximum number of multicast sources allowed to be tracked per group

## lag-link-map-profile

**Syntax**     **lag-link-map-profile** *link-map-profile-id*  
**no lag-link-map-profile**

**Context**     config>subscr-mgmt>msap-policy

**Description**     This command assigns a pre-configured lag link map profile to a SAP/network interface configured on a LAG or a PW port that exists on a LAG. Once assigned/de-assigned, the SAP/network interface egress traffic will be re-hashed over LAG as required by the new configuration.

                    The **no** form of this command reverts the SAP/network interface to use per-flow, service or link hash as configured for the service/LAG.

**Default**     **no lag-link-map-profile**

**Parameters**     *link-map-profile-id* — An integer from 1 to 32 that defines a unique lag link map profile on which the LAG the SAP/network interface exist.

## sub-sla-mgmt

**Syntax**     **[no] sub-sla-mgmt**

**Context**     config>subscr-mgmt>msap-policy  
                 config>service>vpls>sap>sub-sla-mgmt

**Description**     This command enables the context to configure subscriber management parameters for an MSAP.

**Default**     no sub-sla-mgmt

## def-app-profile

**Syntax**     **def-app-profile** *app-profile-name*  
**no def-app-profile**

**Context**     config>service>vpls>sap>sub-sla-mgmt

**Description**     This command specifies the application profile to be used by a subscriber host.

                    The **no** form of the command removes the application profile name from the configuration.

**Default**     no def-app-profile

**Parameters**     *app-profile-name* — specifies an existing application profile to be mapped to the subscriber profile by default.



## def-inter-dest-id

<b>Syntax</b>	<b>def-inter-dest-id</b> { <b>string</b> <i>string</i>   <b>use-top-q</b> } <b>no def-inter-dest-id</b>
<b>Context</b>	config>subscr-mgmt>msap-policy>sub-sla-mgmt
<b>Description</b>	<p>This command specifies a default destination string for all subscribers associated with the SAP. The command also accepts the <b>use-top-q</b> flag that automatically derives the string based on the top most delineating Dot1Q tag from the SAP's encapsulation.</p> <p>The <b>no</b> form of the command removes the default subscriber identification string from the configuration.</p> <p>no def-sub-id</p>
<b>Default</b>	no def-inter-dest-id
<b>Parameters</b>	<p><b>use-top-q</b> — Derives the string based on the top most delineating Dot1Q tag from the SAP's encapsulation.</p> <p><b>string</b> <i>string</i> — Specifies the subscriber identification applicable for a subscriber host.</p>

## def-sub-id

<b>Syntax</b>	<b>def-sub-id use-auto-id</b> <b>def-sub-id use-sap-id</b> <b>def-sub-id string</b> <i>sub-id</i> <b>no def-sub-id</b>
<b>Context</b>	config>subscr-mgmt>msap-policy>sub-sla-mgmt config>service>vpls>sap>sub-sla-mgmt
<b>Description</b>	<p>This command specifies the explicit default sub-id for dynamic subscriber hosts (including ARP hosts) in case that the sub-id string is NOT supplied through RADIUS or LUDB.</p> <p>The sub-id is assigned to a new subscriber host in the following order of priority:</p> <ul style="list-style-type: none"> <li>• RADIUS</li> <li>• LUDB</li> <li>• Explicit default – with the def-sub-id command we explicitly set the sub-id name of the host to be one of the following: <ul style="list-style-type: none"> <li>→ The sap-id to which the new host is associated with</li> <li>→ Explicit string</li> <li>→ Auto-generated string consisting of the concatenated subscriber identification fields defined under the <b>subscr-mgmt&gt;auto-sub-id-key</b> node. The fields are taken in the order in which they are configured and are separated by a ' ' character. The subscriber host identification fields are separately defined for IPoE and PPPoE host types.</li> </ul> </li> <li>• Implicit default – in case that the sub-id string is not returned via RADIUS or LUDB and there is no def-sub-id configured, the sub-id name will be generated as a random 10 character encoded string based on the auto-sub-id-keys. This 10 characters encoded string will be unique per chassis</li> </ul>

as well as in dual-homed environment. It is generated based on auto-sub-id-keys. If auto-sub-id-keys are not explicitly configured, the default ones are:

→ <mac, sap-id, session-id> for PPP type hosts

→ <mac, sap-id> for IPoE type hosts.

This command does not apply to static subscribers.

**Parameters** **use-sap-id** — Specifies the sub-id name -id on which the original request for host creation arrived (DHCP Discover, or PADI or ARP Request).

**string sub-id** — Explicitly configured sub-id name.

**use-auto-id** — The sub-id name is the concatenated string of auto-sub-id-keys separated by a “|” character.

**Default** no def-sub-id

Implicit default – If the sub-id string is not supplied through RADIUS, LUDB or by configuration (def-sub-id), then a random 10 character encoded sub-id name will be generated. This random sub-id name will be based on the subscriber identification keys defined under the subscr-mgmt>auto-sub-id-key node. In case that the auto-sub-id-keys are not defined explicitly, the default ones are:

- <mac, sap-id, session-id>for PPPoE type hosts
- <mac, sap-id>for IPoE type hosts

## def-sla-profile

**Syntax** **def-sla-profile** *default-sla-profile-name*  
**no def-sla-profile**

**Context** config>subscr-mgmt>msap-policy>sub-sla-mgmt

**Description** This command specifies a default SLA profile for an MSAP.

An SLA profile is a named group of QoS parameters used to define per service QoS for all subscriber hosts common to the same subscriber within a provider service offering. A single SLA profile may define the QoS parameters for multiple subscriber hosts.

The **no** form of the command removes the default SLA profile from the MSAP configuration.

**Default** no def-sla-profile

**Parameters** *default-sla-profile-name* — Specifies a default SLA profile for an MSAP.

## def-sub-profile

**Syntax** **def-sub-profile** *default-subscriber-profile-name*

**Context** config>subscr-mgmt>msap-policy>sub-sla-mgmt

**Description** This command specifies a default subscriber profile for an MSAP.

A subscriber profile defines the aggregate QoS for all hosts within a subscriber context. This is done through the definition of the egress and ingress scheduler policies that govern the aggregate SLA for subscriber using the subscriber profile.

The **no** form of the command removes the default SLA profile from the SAP configuration.

**Parameters** *default-sub-profile* — Specifies a default subscriber profile for this SAP.

## multi-sub-sap

**Syntax** **multi-sub-sap** [*limit limit*]  
**no multi-sub-sap**

**Context** config>subscr-mgmt>msap-policy>sub-sla-mgmt

**Description** This command defines the maximum number of subscribers (dynamic + static) that can be simultaneously active on an MSAP.

If the limit is reached, a new host will be denied access and the corresponding DHCP ACK will be dropped.

The **no** form of the command reverts back to the default setting.

**Default** 1

**Parameters** **limit** *limit* — Specifies the maximum allowed.

**Values** 2 — 8000

## single-sub-parameters

**Syntax** **single-sub-parameters**

**Context** config>subscr-mgmt>msap-policy>sub-sla-mgmt

**Description** This command enables the context to configure single subscriber MSAP parameters.

## non-sub-traffic

**Syntax** **non-sub-traffic sub-profile** *sub-profile-name* **sla-profile** *sla-profile-name* [**subscriber** *sub-ident-string*] [**app-profile** *app-profile-name*]  
**no non-sub-traffic**

**Context** config>subscr-mgmt>msap-policy>sub-sla-mgmt>single-sub

**Description** This command configures traffic profiles for non-IP traffic such as PPPoE. It is used in conjunction with the profiled-traffic-only on single subscriber SAPs and creates a subscriber host which is used to forward non-IP traffic through the single subscriber SAP without the need for SAP queues.

The **no** form of the command removes any configured profile.

**Default** no non-sub-traffic

**Parameters**     *sub-profile-name* — Identifies the subscriber profile name.

**Values**            32 characters maximum

*sla-profile-name* — Identifies the SLA profile name.

**Values**            32 characters maximum

### profiled-traffic-only

**Syntax**            **[no] profiled-traffic-only**

**Context**            config>subscr-mgmt>msap-policy>sub-sla-mgmt>single-sub

**Description**        This command specifies whether only profiled traffic is applicable for an MSAP. When enabled, all queues will be deleted.

The **no** form of the command reverts to the default setting.

**Default**            no profiled-traffic-only

### sub-ident-policy

**Syntax**            **[no] sub-ident-policy** *sub-ident-policy-name*

**Context**            config>subscr-mgmt>msap-policy>sub-sla-mgmt

**Description**        This command specifies an existing subscriber identification policy. Each subscriber identification policy can have a default subscriber profile defined. The subscriber identification policy default subscriber profile overrides the system default and the subscriber SAP default subscriber profiles. Defining a subscriber identification policy default subscriber profile is optional.

Defining a subscriber profile as a subscriber identification policy default subscriber profile will cause all active subscribers currently associated with a subscriber SAP using the policy and associated with a subscriber policy through the system default or subscriber SAP default subscriber profiles to be reassigned to the subscriber policy defined as default on the subscriber identification policy.

Attempting to delete a subscriber profile that is currently defined as a default for a subscriber identification policy will fail.

When attempting to remove a subscriber identification policy default subscriber profile definition, the system will evaluate each active subscriber on all subscriber SAPs the subscriber identification policy is currently associated with that are using the default definition to determine whether the active subscriber can be either reassigned to a subscriber SAP default or the system default subscriber profile. If all active subscribers cannot be reassigned, the removal attempt will fail.

**Parameters**        *sub-ident-policy-name* — Specifies the name of the subscriber identification policy.

### vpls-only-sap-parameters

**Syntax**            **vpls-only-sap-parameters**

<b>Context</b>	config>subscr-mgmt>msap-policy
<b>Description</b>	This command enables the context to configure MSAP VPLS properties.

## arp-host

<b>Syntax</b>	<b>arp-host</b>
<b>Context</b>	config>subscr-mgmt>msap-policy>vpls-only config>service>vpls>sap>arp-host
<b>Description</b>	This command enables the context to configure ARP host parameters.

## host-limit

<b>Syntax</b>	<b>host-limit</b> <i>max-num-hosts</i> <b>no host-limit</b>
<b>Context</b>	config>subscr-mgmt>msap-policy>vpls-only config>service>vpls>sap>arp-host
<b>Description</b>	This command configures the maximum number of ARP hosts.
<b>Parameters</b>	<i>max-num-hosts</i> — Specifies the maximum number of ARP hosts. <b>Values</b> 1 — 32767

## min-auth-interval

<b>Syntax</b>	<b>min-auth-interval</b> <i>min-auth-interval</i> <b>no min-auth-interval</b>
<b>Context</b>	config>subscr-mgmt>msap-policy>vpls-only config>service>vpls>sap>arp-host
<b>Description</b>	This command configures the minimum authentication interval.
<b>Parameters</b>	<i>min-auth-interval</i> — Specifies the minimum authentication interval. <b>Values</b> 1 — 6000

## arp-reply-agent

<b>Syntax</b>	<b>arp-reply-agent</b> [sub-ident] <b>no arp-reply-agent</b>
<b>Context</b>	config>subscr-mgmt>msap-policy>vpls-only

<b>Description</b>	<p>This command enables a special ARP response mechanism in the system for ARP requests destined to static or dynamic hosts associated with the SAP. The system responds to each ARP request using the hosts MAC address as the both the source MAC address in the Ethernet header and the target hardware address in the ARP header.</p> <p>ARP replies and requests received on an MSAP with <b>arp-reply-agent</b> enabled will be evaluated by the system against the anti-spoof filter entries associated with the ingress SAP (if the SAP has anti-spoof filtering enabled). ARPs from unknown hosts on the SAP will be discarded when anti-spoof filtering is enabled.</p> <p>The ARP reply agent only responds if the ARP request enters an interface (SAP, spoke-SDP or mesh-SDP) associated with the VPLS instance of the MSAP.</p> <p>A received ARP request that is not in the ARP reply agent table is flooded to all forwarding interfaces of the VPLS capable of broadcast except the ingress interface while honoring split-horizon constraints.</p> <p>Static hosts can be defined using the <b>host</b> command. Dynamic hosts are enabled on the system by enabling the <b>lease-populate</b> command in the <b>dhcp</b> context. In the event that both a static host and a dynamic host share the same IP and MAC address, the VPLS ARP reply agent will retain the host information until both the static and dynamic information are removed. In the event that both a static and dynamic host share the same IP address, but different MAC addresses, the VPLS ARP reply agent is populated with the static host information.</p> <p>The <b>arp-reply-agent</b> command will fail if an existing static host does not have both MAC and IP addresses specified. Once the ARP reply agent is enabled, creating a static host on the MSAP without both an IP address and MAC address will fail.</p> <p>The ARP-reply-agent may only be enabled on SAPs supporting Ethernet encapsulation.</p> <p>The <b>no</b> form of the command disables ARP-reply-agent functions for static and dynamic hosts on the MSAP.</p>
<b>Default</b>	not enabled
<b>Parameters</b>	<p><b>sub-ident</b> — Configures the arp-reply-agent to discard ARP requests received on the MSAP that are targeted for a known host on the same MSAP with the same subscriber identification.</p> <p>Hosts are identified by their subscriber information. For DHCP subscriber hosts, the subscriber hosts, the subscriber information is configured using the optional subscriber parameter string.</p> <p>When arp-reply-agent is enabled with <b>sub-ident</b>:</p> <ul style="list-style-type: none"><li>• If the subscriber information for the destination host exactly matches the subscriber information for the originating host and the destination host is known on the same MSAP as the source, the ARP request is silently discarded.</li><li>• If the subscriber information for the destination host or originating host is unknown or undefined, the source and destination hosts are not considered to be the same subscriber. The ARP request is forwarded outside the MSAP's Split Horizon Group.</li><li>• When <b>sub-ident</b> is not configured, the arp-reply-agent does not attempt to identify the subscriber information for the destination or originating host and will not discard an ARP request based on subscriber information.</li></ul>

## dhcp

<b>Syntax</b>	<b>dhcp</b>
<b>Context</b>	config>subscr-mgmt>msap-policy>vpls-only
<b>Description</b>	This command enables the context to configure DHCP parameters.

## option

<b>Syntax</b>	<b>[no] option</b>
<b>Context</b>	config>subscr-mgmt>msap-policy>vpls-only>dhcp config>service>ies>sub-if>dhcp
<b>Description</b>	This command enables DHCP Option 82 (Relay Agent Information Option) parameters processing and enters the context for configuring Option 82 sub-options.  The <b>no</b> form of this command returns the system to the default.
<b>Default</b>	no option

## action

<b>Syntax</b>	<b>action {replace   drop   keep}</b> <b>no action</b>
<b>Context</b>	config>subscr-mgmt>msap-policy>vpls-only>dhcp>option
<b>Description</b>	This command configures the Relay Agent Information Option (Option 82) processing.  The <b>no</b> form of this command returns the system to the default value.
<b>Default</b>	The default is to keep the existing information intact.
<b>Parameters</b>	<b>replace</b> — In the upstream direction (from the user), the Option 82 field from the router is inserted in the packet (overwriting any existing Option 82 field). In the downstream direction (towards the user) the Option 82 field is stripped (in accordance with RFC 3046).  <b>drop</b> — The DHCP packet is dropped if an Option 82 field is present, and a counter is incremented.  <b>keep</b> — The existing information is kept in the packet and the router does not add any additional information. In the downstream direction the Option 82 field is not stripped and is forwarded towards the client.

## circuit-id

<b>Syntax</b>	<b>circuit-id [ascii-tuple   vlan-ascii-tuple]</b> <b>no circuit-id</b>
<b>Context</b>	config>subscr-mgmt>msap-policy>vpls-only>dhcp>option

## Managed SAP Policy Commands

<b>Description</b>	<p>When enabled, the router sends an ASCII-encoded tuple in the <b>circuit-id</b> sub-option of the DHCP packet. This ASCII-tuple consists of the access-node-identifier, service-id, and SAP-ID, separated by “ ”.</p> <p>If disabled, the <b>circuit-id</b> sub-option of the DHCP packet will be left empty.</p> <p>The <b>no</b> form of this command returns the system to the default.</p>
<b>Default</b>	circuit-id
<b>Parameters</b>	<p><b>ascii-tuple</b> — Specifies that the ASCII-encoded concatenated tuple consisting of the access-node-identifier, service-id, and interface-name is used.</p> <p><b>vlan-ascii-tuple</b> — Specifies that the format will include VLAN-id and dot1p bits in addition to what is included in ascii-tuple already. The format is supported on dot1q and qinq ports only. Thus, when the option 82 bits are stripped, dot1p bits will be copied to the Ethernet header of an outgoing packet.</p>

## vendor-specific-option

<b>Syntax</b>	<b>[no] vendor-specific-option</b>
<b>Context</b>	config>subscr-mgmt>msap-policy>vpls-only>dhcp>option config>service>ies>sub-if>dhcp
<b>Description</b>	This command configures the Alcatel-Lucent vendor specific sub-option of the DHCP relay packet.

## client-mac-address

<b>Syntax</b>	<b>[no] client-mac-address</b>
<b>Context</b>	config>subscr-mgmt>msap-policy>vpls-only>dhcp>option>vendor config>service>ies>sub-if>dhcp>option
<b>Description</b>	<p>This command enables the sending of the MAC address in the Alcatel-Lucent vendor specific sub-option of the DHCP relay packet.</p> <p>The <b>no</b> form of the command disables the sending of the MAC address in the Alcatel-Lucent vendor specific sub-option of the DHCP relay packet.</p>

## sap-id

<b>Syntax</b>	<b>[no] sap-id</b>
<b>Context</b>	config>subscr-mgmt>msap-policy>vpls-only>dhcp>option>vendor config>service>ies>sub-if>dhcp>option
<b>Description</b>	This command enables the sending of the SAP ID in the Alcatel-Lucent vendor specific sub-option of the DHCP relay packet.



The **no** form of the command disables the sending of the SAP ID in the Alcatel-Lucent vendor specific sub-option of the DHCP relay packet.

## service-id

<b>Syntax</b>	<b>[no] service-id</b>
<b>Context</b>	config>subscr-mgmt>msap-policy>vpls-only>dhcp>option>vendor config>service>ies>sub-if>dhcp>option
<b>Description</b>	This command enables the sending of the service ID in the Alcatel-Lucent vendor specific sub-option of the DHCP relay packet.  The <b>no</b> form of the command disables the sending of the service ID in the Alcatel-Lucent vendor specific sub-option of the DHCP relay packet.

## string

<b>Syntax</b>	<b>[no] string <i>text</i></b>
<b>Context</b>	config>subscr-mgmt>msap-policy>vpls-only>dhcp>option>vendor config>service>ies>sub-if>dhcp>option
<b>Description</b>	This command specifies the string in the Alcatel-Lucent vendor specific sub-option of the DHCP relay packet.  The <b>no</b> form of the command returns the default value.
<b>Parameters</b>	<i>text</i> — The string can be any combination of ASCII characters up to 32 characters in length. If spaces are used in the string, enclose the entire string in quotation marks (“ ”).

## system-id

<b>Syntax</b>	<b>[no] system-id</b>
<b>Context</b>	config>subscr-mgmt>msap-policy>vpls-only>dhcp>option>vendor config>service>ies>sub-if>dhcp>option
<b>Description</b>	This command specifies whether the system-id is encoded in the Alcatel-Lucent vendor specific sub-option of Option 82.

## emulated-server

<b>Syntax</b>	<b>emulated-server <i>ip-address</i></b> <b>no emulated-server</b>
<b>Context</b>	config>subscr-mgmt>msap-policy>vpls-only>dhcp>proxy config>service>ies>sub-if>dhcp

<b>Description</b>	<p>This command configures the IP address which will be used as the DHCP server address in the context of the MSAP. Typically, the configured address should be in the context of the subnet represented by the service.</p> <p>The <b>no</b> form of this command reverts to the default setting. The local proxy server will not become operational without the emulated-server address being specified.</p>
<b>Parameters</b>	<p><i>ip-address</i> — Specifies the emulated server's IP address. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).</p>

### lease-time

<b>Syntax</b>	<b>lease-time [days days] [hrs hours] [min minutes] [sec seconds] [override]</b> <b>no lease-time</b>
<b>Context</b>	config>subscr-mgmt>msap-policy>vpls-only>dhcp>proxy config>service>ies>sub-if>dhcp
<b>Description</b>	<p>This command defines the length of lease-time that will be provided to DHCP clients. By default the local-proxy-server will always make use of the lease-time information provide by either a RADIUS or DHCP server.</p> <p>The no form of this command disables the use of the lease-time command. The local-proxy-server will use the lease-time offered by either a RADIUS or DHCP server.</p>
<b>Default</b>	7 days 0 hours 0 seconds
<b>Parameters</b>	<p><b>override</b> — Specifies that the local-proxy-server will use the configured lease-time information to provide DHCP clients.</p> <p><i>days</i> — Specifies the number of days that the given IP address is valid.</p> <p><b>Values</b> 0 — 3650</p> <p><i>hours</i> — Specifies the number of hours that the given IP address is valid.</p> <p><b>Values</b> 0 — 23</p> <p><i>minutes</i> — Specifies the number of minutes that the given IP address is valid.</p> <p><b>Values</b> 0 — 59</p> <p><i>seconds</i> — Specifies the number of seconds that the given IP address is valid.</p> <p><b>Values</b> 0 — 59</p>

### egress

<b>Syntax</b>	<b>egress</b>
<b>Context</b>	config>subscr-mgmt>msap-policy>vpls-only
<b>Description</b>	This command configures egress policies for MSAPs.

## multicast-group

<b>Syntax</b>	<b>multicast-group</b> <i>group-name</i> <b>no multicast-group</b>
<b>Context</b>	config>subscr-mgmt>msap-policy>vpls-only>egress
<b>Description</b>	<p>This command specifies an existing egress multicast group (EMG). An EMG is created as an object used to group VPLS SAPs that are allowed to participate in efficient multicast replication (EMR). EMR is a method to increase the performance of egress multipoint forwarding by sacrificing some destination-based features. Eliminating the requirement to perform unique features for each destination allows the egress forwarding plane to chain together multiple destinations into a batch replication process. In order to perform this batch replication function, similar characteristics are required on each SAP within the EMG.</p> <p>Only SAPs defined on Ethernet access ports are allowed into an egress-multicast-group.</p> <p>In order to understand the purpose of an egress-multicast-group, an understanding of the system's use of flooding lists is required. A flooding list is maintained at the egress forwarding plane to define a set of destinations to which a packet must be replicated. Multipoint services make use of flooding lists to enable forwarding a single packet to many destinations. Examples of multipoint services that use flooding lists are VPLS, IGMP snooping and IP multicast routing. Currently, the egress forwarding plane will only use efficient multicast replication for VPLS and IGMP snooping flooding lists.</p> <p>In VPLS services, a unique flooding list is created for each VPLS context. The flooding list is used when a packet has a broadcast, multicast or unknown destination MAC address. From a system perspective, proper VPLS handling requires that a broadcast, multicast or unknown destined packet be sent to all destinations that are in the forwarding state. The ingress forwarding plane ensures the packet gets to all egress forwarding planes that include a destination in the VPLS context. It is the egress forwarding plane's job to replicate the packet to the subset of the destinations that are reached through its interfaces and each of these destinations are included in the VPLS context's flooding list.</p> <p>For IGMP snooping, a unique flooding list is created for each IP multicast (s,g) record. This (s,g) record is associated with an ingress VPLS context and may be associated with VPLS destinations in the source VPLS instance or other VPLS instances (in the case of MVR). Again, the ingress forwarding plane ensures that an ingress IP multicast packet matching the (s,g) record gets to all egress forwarding planes that have a VPLS destination associated with the (s,g) record. The egress forwarding plane uses the flooding list owned by the (s,g) record to replicate the packet to all VPLS destinations in the flooding list. The IGMP Snooping function identifies which VPLS destinations should be associated with the (s,g) record.</p> <p>With normal multicast replication, the egress forwarding plane examines which features are enabled for each destination. This includes ACL filtering, mirroring, encapsulation and queuing. The resources used to perform this per destination multicast processing are very expensive to the egress forwarding plane when high replication bandwidth is required. If destinations with similar egress functions can be grouped together, the egress forwarding plane can process them in a more efficient manner and maximize replication bandwidth.</p> <p>The egress-multicast-group object is designed to allow the identification of SAPs with similar egress characteristics. When a SAP is successfully provisioned into an egress-multicast-group, the system is ensured that it may be batched together with other SAPs in the same group at the egress forwarding plane for efficient multicast replication. A SAP that does not meet the common requirements is not allowed into the egress-multicast-group.</p>

At the forwarding plane level, a VPLS flooding list is categorized into chainable and non-chainable destinations. Currently, the only chainable destinations are SAPs within an egress-multicast-group. The chainable destinations are further separated by egress-multicast-group association. Chains are then created following the rules below:

- A replication batch chain may only contain SAPs from the same egress-multicast-group
- A replication batch chain length may not exceed the dest-chain-limit of the egress-multicast-group to which the SAPs are members

Further subcategories are created for an IGMP (s,g) flooding list. A Layer 2 (s,g) record is created in a specific VPLS instance (the instance the (s,g) flow ingresses). SAPs within that VPLS context that join the (s,g) record are considered native SAPs within the flooding list. SAPs that join the (s,g) flooding list through the multicast VPLS registration process (MVR) from another VPLS context using the **from-vpls** command are considered alien SAPs. The distinction between native and alien in the list is maintained to allow the forwarding plane to enforce or suspend split-horizon-group (SHG) squelching. When the source of the (s,g) matching packet is in the same SHG as a native SAP, the packet must not be replicated to that SAP. For a SAP in another VPLS context, the source SHG of the packet has no meaning and the forwarding plane must disregard SHG matching between the native source of the packet and the alien destination. Because the SHG squelch decision is done for the whole chain based on the first SAP in the chain, all SAPs in the chain must be all native or all alien SAPs. Chains for IGMP (s,g) flooding lists are created using the following rules:

1. A replication batch chain may only contain SAPs from the same egress-multicast-group.
2. A replication batch chain may only contain all alien or all native SAPs.
3. A replication batch chain length may not exceed the dest-chain-limit of the egress-multicast-group to which the SAPs are members

When a packet associated with a flooding list is received by the egress forwarding plane, it processes the packet by evaluating each destination on the list sequentially in a replication context. If the current entry being processed in the list is a non-chained destination, the forwarding plane processes the packet for that destination and then moves on to process other packets currently in the forwarding plane before returning to process the next destination in the list. If the current entry being processed is a chained destination, the forwarding plane remains in the replication context until it has forwarded to each entry in that chain. Once the replication context finishes with the last entry in the chain, it moves on to process other packets waiting for egress processing before returning to the replication context. Processing continues in this manner until the packet has been forwarded to all destinations in the list.

Batch chain processing of a chain of SAPs improves replication efficiency by bypassing the functions that perform egress mirroring decisions on SAPs within the chain and making a single ACL filtering decision for the whole chain. Each destination in the chain may have a unique egress QoS policy and per destination queuing is still performed for each destination in the chain. Also, while each SAP in the chain must be on access ports with the same encap-type, if the encap-type is dot1q, each SAP may have a unique dot1q tag.

One caveat to each SAP having a unique egress QoS policy in the chain is that only the Dot1P marking decisions for the first SAP in the list is enforced. If the first SAP's QoS policy forwarding class action states that the packet should not be remarked, none of the replicated packets in the chain will have the dot1P bits remarked. If the first SAP's QoS policy forwarding class action states that the packet should be remarked with a specific dot1P value, all the replicated packets for the remaining SAPs in the chain will have the same dot1P marking.

While the system supports 32 egress multicast groups, a single group would usually suffice. An instance where multiple groups would be needed is when all the SAPs requiring efficient multicast replication cannot share the same common requirements. In this case, an egress multicast group

would be created for each set of common requirements. An egress multicast group may contain SAPs from many different VPLS instances. It should be understood that an egress multicast group is not equivalent to an egress forwarding plane flooding list. An egress multicast group only identifies which SAPs may participate in efficient multicast replication. As stated above, entries in a flooding list are populated due to VPLS destination creation or IGMP snooping events.

The **no** form of the command removes a specific egress multicast group. Deleting an egress multicast group will only succeed when the group has no SAP members. To remove SAP members, use the **no multicast-group group-name** command under each SAP's egress context.

**Note:** Efficient multicast replication will only be performed on IOMs that support chassis mode b. If an IOM does not support mode b operation, egress-multicast-group membership is ignored on that IOM's egress forwarding planes. The chassis need not be placed into mode b for efficient multicast replication to be performed on the capable IOMs.

<b>Parameters</b>	<i>group-name</i> — Multiple egress multicast groups may be created on the system. Each must have a unique name. The egress-multicast-group-name is an ASCII string up to 16 characters in length and follows all the naming rules as other named policies in the system. The group's name is used throughout the system to uniquely identify the Egress Multicast Group and is used to provision a SAP into the group.
<b>Default</b>	None, each egress multicast group must be explicitly configured.
<b>Values</b>	Up to 32 egress multicast groups may be created on the system.

## igmp-snooping

<b>Syntax</b>	<b>igmp-snooping</b>
<b>Context</b>	config>subscr-mgmt>msap-policy>vpls-only
<b>Description</b>	This command enables the Internet Group Management Protocol (IGMP) snooping context.
<b>Default</b>	none

## fast-leave

<b>Syntax</b>	<b>[no] fast-leave</b>
<b>Context</b>	config>subscr-mgmt>msap-policy>vpls-only>igmp-snp
<b>Description</b>	<p>This command enables fast leave.</p> <p>When IGMP fast leave processing is enabled, the 7450 ESS will immediately remove a SAP or SDP from the IP multicast group when it detects an IGMP 'leave' on that SAP or SDP. Fast leave processing allows the switch to remove a SAP or SDP that sends a 'leave' from the forwarding table without first sending out group-specific queries to the SAP or SDP, and thus speeds up the process of changing channels ('zapping').</p> <p>Fast leave should only be enabled when there is a single receiver present on the SAP or SDP.</p> <p>When fast leave is enabled, the configured last-member-query-interval value is ignored.</p>
<b>Default</b>	no fast-leave

## import

<b>Syntax</b>	<b>import</b> <i>policy-name</i> <b>no import</b>
<b>Context</b>	config>subscr-mgmt>msap-policy>vpls-only>igmp-snp
<b>Description</b>	This command specifies the import routing policy to be used for IGMP packets to be used on this SAP or SDP. Only a single policy can be imported on a single SAP at any time.  The <b>no</b> form of the command removes the policy association from the SAP or SDP.
<b>Default</b>	no import (No import policy is specified)
<b>Parameters</b>	<i>policy-name</i> — The routing policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters excluding double quotes. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes. Routing policies are configured in the config>router>policy-options context. The router policy must be defined before it can be imported.

## last-member-query-interval

<b>Syntax</b>	<b>last-member-query-interval</b> <i>tenths-of-seconds</i> <b>no last-member-query-interval</b>
<b>Context</b>	config>subscr-mgmt>msap-policy>vpls-only>igmp-snp
<b>Description</b>	This command configures the maximum response time used in group-specific queries sent in response to 'leave' messages, and is also the amount of time between 2 consecutive group-specific queries. This value may be tuned to modify the leave latency of the network. A reduced value results in reduced time to detect the loss of the last member of a group.  The configured last-member-query-interval is ignored when fast-leave is enabled on the SAP or SDP.
<b>Default</b>	10
<b>Parameters</b>	<i>seconds</i> — Specifies the frequency, in tenths of seconds, at which query messages are sent.  <b>Values</b> 1 — 50

## max-num-groups

<b>Syntax</b>	<b>max-num-groups</b> <i>max-num-groups</i> <b>no max-num-groups</b>
<b>Context</b>	config>subscr-mgmt>msap-policy>vpls-only>igmp-snp
<b>Description</b>	This command defines the maximum number of multicast groups that can be joined on an MSAP or SDP. If the router receives an IGMP join message that would exceed the configured number of groups, the request is ignored.
<b>Default</b>	no max-num-groups

**Parameters** *max-num-groups* — Specifies the maximum number of groups that can be joined on an MSAP or SDP.

**Values** 1 — 1000

## mcac

**Syntax** **mcac**

**Context** config>subscr-mgmt>msap-policy>vpls-only>igmp-snp>

**Description** This command enables the context to configure multicast CAC parameters.

**Default** none

## mc-constraints

**Syntax** **mc-constraints**

**Context** config>subscr-mgmt>msap-policy>vpls-only>igmp-snp>mcac

**Description** This command enables the context to configure the level and its associated bandwidth for a bundle or a logical interface.

**Default** none

## level

**Syntax** **level** *level-id* **bw** *bandwidth*  
**no level** *level-id*

**Context** config>subscr-mgmt>msap-policy>vpls-only>igmp-snp>mcac

**Description** This command configures levels and their associated bandwidth for multicast CAC policy on an interface.

**Parameters** *level-id* — Specifies has an entry for each multicast CAC policy constraint level configured on a system.

**Values** 1 — 8

*bandwidth* — Specifies the bandwidth in kilobits per second (kbps) for the level.

**Values** 1 — 2147483647

## number-down

**Syntax** **number-down** *number-lag-port-down* **level** *level-id*  
**no number-down** *number-lag-port-down*

<b>Context</b>	config>subscr-mgmt>msap-policy>vpls-only>igmp-snp>mcac
<b>Description</b>	This command configures the number of ports down along with level for multicast CAC policy on an MSAP
<b>Parameters</b>	<p><i>number-lag-port-down</i> — If the number of ports available in the LAG is reduced by the number of ports configured in this command here then bandwidth allowed for bundle and/or interface will be as per the levels configured in this context.</p> <p><b>Values</b>      1 — 64 (for 64-link LAG)                            1 — 32 (for other LAGs)</p> <p><b>level</b> <i>level-id</i> — Specifies the amount of bandwidth available within a given bundle for MC traffic for a specified level.</p>

### policy

<b>Syntax</b>	<b>policy</b> <i>policy-name</i> <b>no policy</b>
<b>Context</b>	config>subscr-mgmt>msap-policy>vpls-only>igmp-snp>mcac
<b>Description</b>	This command configures the multicast CAC policy name.
<b>Parameters</b>	<i>policy-name</i> — The multicast CAC policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

### unconstrained-bw

<b>Syntax</b>	<b>unconstrained-bw</b> <i>bandwidth</i> <b>mandatory-bw</b> <i>mandatory-bw</i> <b>no unconstrained-bw</b>
<b>Context</b>	config>subscr-mgmt>msap-policy>vpls-only>igmp-snp>mcac
<b>Description</b>	This command configures the bandwidth for the interface's multicast CAC policy traffic. When disabled ( <b>no unconstrained-bw</b> ) there will be no checking of bandwidth constraints on the interface level. When enabled and a policy is defined, enforcement is performed. The allocated bandwidth for optional channels should not exceed the <b>unconstrained-bw</b> minus the <b>mandatory-bw</b> and the mandatory channels have to stay below the specified value for the <b>mandatory-bw</b> . After this interface check, the bundle checks are performed.
<b>Parameters</b>	<p><i>bandwidth</i> — The bandwidth assigned for interface's MCAC policy traffic, in kilo-bits per second (kbps).</p> <p><b>Values</b>      0 — 2147483647</p> <p><b>mandatory-bw</b> <i>mandatory-bw</i> — Specifies the bandwidth pre-reserved for all the mandatory channels on a given interface in kilo-bits per second (kbps).</p> <p>If the <i>bandwidth</i> value is 0, no mandatory channels are allowed. If the value of <i>bandwidth</i> is '-1', then all mandatory and optional channels are allowed.</p>



If the value of *mandatory-bw* is equal to the value of *bandwidth*, then all the unconstrained bandwidth on a given interface is allocated to mandatory channels configured through multicast CAC policy on that interface and no optional groups (channels) are allowed.

The value of *mandatory-bw* should always be less than or equal to that of *bandwidth*. An attempt to set the value of *mandatory-bw* greater than that of *bandwidth*, will result in inconsistent value error.

**Values** 0 — 2147483647

## sub-mcac-policy

<b>Syntax</b>	<b>sub-mcac-policy</b> <i>sub-mcac-policy-name</i> [ <b>create</b> ] <b>no sub-mcac-policy</b> <i>b</i>
<b>Context</b>	config>subscr-mgmt
<b>Description</b>	<p>This command will create a policy template with mcac bandwidth limits that will be applied to the subscriber.</p> <p>Per interface mcac bandwidth limits will be set directly under the interface (regular interface or group-interface) and no such policy templates are needed.</p> <p>The need for a separate policy template for subscribers is due to the fact that groups of subscribers under the same group-interface can share certain settings that can be configured via this template.</p> <p>To summarize, the mcac bandwidth constraints for subscribers are defined in the sub-mcac-policy while the mcac bandwidth constraints for the interface are configured directly under the <b>igmp&gt;interface&gt;mcac</b> or <b>igmp&gt;group-interface&gt;mcac</b> context without the need for policy templates.</p> <p>Note that the sub-mcac-policy only deals with the mcac bandwidth limits and not the channel bandwidth definitions. Channels bandwidth is defined in a different policy (under the config&gt;router&gt;mcac hierarchy) and that policy is applied on the interface level as follows:</p> <ul style="list-style-type: none"> <li>• For group-interface: under the <b>configure&gt;service&gt;vprn&gt;igmp&gt;group-interface&gt;mcac</b> context</li> <li>• For regular interface: under the <b>configure&gt;service/router&gt;igmp&gt;interface&gt;mcac</b> context.</li> </ul> <p>In case of HQoS Adjustment, it is mandatory that the sub-mcac-policy be created and applied to the subscriber. The sub-mac-policy does not have to contain any bandwidth constraints, but it has to be in a no shutdown state in order for HQoS Adjustment to work.</p>
<b>Parameters</b>	<i>policy-name</i> — Specifies the name of the policy.

## mvr

<b>Syntax</b>	<b>mvr</b>
<b>Context</b>	config>subscr-mgmt>msap-policy>vpls-only>igmp-snp
<b>Description</b>	This command enables the context to configure Multicast VPLS Registration (MVR) parameters.

## from-vpls

<b>Syntax</b>	<b>from-vpls</b> <i>service-id</i> <b>no from-vpls</b>
<b>Context</b>	config>subscr-mgmt>msap-policy>vpls-only>igmp-snp>mvr
<b>Description</b>	This command configures the VPLS from which multicast traffic is copied upon receipt of an IGMP join request.  IGMP snooping must be enabled on the MVR VPLS.
<b>Default</b>	no from-vpls
<b>Parameters</b>	<i>service-id</i> — Specifies the MVR VPLS from which multicast channels should be copied into an MSAP.
<b>Values</b>	<i>service-id:</i> 1 — 2147483647

## query-interval

<b>Syntax</b>	<b>query-interval</b> <i>seconds</i> <b>no query-interval</b>
<b>Context</b>	config>subscr-mgmt>msap-policy>vpls-only>igmp-snp
<b>Description</b>	This command configures the IGMP query interval. If the <b>send-queries</b> command is enabled, this parameter specifies the interval between two consecutive general queries sent by the system on an MSAP or SDP.  The configured query-interval must be greater than the configured query-response-interval.  If send-queries is not enabled on an MSAP or SDP, the configured query-interval value is ignored.
<b>Default</b>	125
<b>Parameters</b>	<i>seconds</i> — The time interval, in seconds, that the router transmits general host-query messages.
<b>Values</b>	2 — 1024

## query-response-interval

<b>Syntax</b>	<b>query-response-interval</b> <i>seconds</i>
<b>Context</b>	config>subscr-mgmt>msap-policy>vpls-only>igmp-snp
<b>Description</b>	This command configures the IGMP query response interval. If the <b>send-queries</b> command is enabled, this parameter specifies the maximum response time advertised in IGMPv2/v3 queries.  The configured query-response-interval must be smaller than the configured query-interval.  If send-queries is not enabled on an MSAP or SDP, the configured query-response-interval value is ignored.
<b>Default</b>	10

**Parameters** *seconds* — Specifies the length of time to wait to receive a response to the host-query message from the host.

**Values** 1 — 1023

## robust-count

**Syntax** **robust-count** *robust-count*  
**no robust-count**

**Context** config>subscr-mgmt>msap-policy>vpls-only>igmp-snp

**Description** This command configures the IGMP robustness variable. If the **send-queries** command is enabled, this parameter allows tuning for the expected packet loss on a SAP or SDP. The robust-count variable allows tuning for the expected packet loss on a subnet and is comparable to a retry count. If an MSAP or SDP is expected to be “lossy”, this parameter may be increased. IGMP snooping on an MSAP or SDP is robust to (robust-count-1) packet losses.

If send-queries is not enabled, this parameter will be ignored.

**Default** 2

**Parameters** *robust-count* — Specifies the robust count for the SAP or SDP.

**Values** 2 — 7

## send-queries

**Syntax** **[no] send-queries**

**Context** config>subscr-mgmt>msap-policy>vpls-only>igmp-snp

**Description** This command specifies whether to send IGMP general query messages on the managed SAP. When send-queries is configured, all type of queries generate ourselves are of the configured version. If a report of a version higher than the configured version is received, the report will get dropped and a new wrong version counter will get incremented.

If send-queries is not configured, the version command has no effect. The version used on that SAP/SDP will be the version of the querier. This implies that, for example, when there is a v2 querier, a v3 group or group-source specific query is never sent when a host wants to leave a certain group.

**Default** no send-queries

## version

**Syntax** **version** *version*  
**no version**

**Context** config>subscr-mgmt>msap-policy>vpls-only>igmp-snp

<b>Description</b>	<p>This command specifies the version of IGMP which is running on an MSAP. This object can be used to configure a router capable of running either value. For IGMP to function correctly, all routers on a LAN must be configured to run the same version of IGMP on that LAN.</p> <p>When the <b>send-query</b> command is configured, all type of queries generate ourselves are of the configured <b>version</b>. If a report of a version higher than the configured version is received, the report gets dropped and a new “wrong version” counter is incremented.</p> <p>If the <b>send-query</b> command is not configured, the <b>version</b> command has no effect. The version used on that SAP or SDP will be the version of the querier. This implies that, for example, when there is a v2 querier, a v3 group or group-source specific query when a host wants to leave a certain group will never be sent.</p>		
<b>Parameters</b>	<p><i>version</i> — Specify the IGMP version.</p> <table><tr><td><b>Values</b></td><td>1, 2, 3</td></tr></table>	<b>Values</b>	1, 2, 3
<b>Values</b>	1, 2, 3		

### mac-da-hashing

**[no] mac-da-hashing**

<b>Context</b>	config>subscr-mgmt>msap-policy>vpls-only>igmp-snp
<b>Description</b>	<p>This command specifies whether subscriber traffic egressing a LAG SAP has its egress LAG link selected by a function of the MAC destination address instead of the subscriber ID.</p> <p>This command is only meaningful if subscriber management is enabled and can be configured for a VPLS service.</p>

### split-horizon-group

<b>Syntax</b>	<b>split-horizon-group</b> <i>group-name</i>
<b>Context</b>	config>subscr-mgmt>msap-policy>vpls-only>igmp-snp
<b>Description</b>	This command specifies the name of the split horizon group to which the MSAP belongs.

### default-msap-policy

<b>Syntax</b>	<b>default-msap-policy</b> <i>policy-name</i> <b>no default-msap-policy</b>
<b>Context</b>	config>service>vpls>sap
<b>Description</b>	<p>This command specifies the default managed SAP policy to use to create MSAPs when the response from the RADIUS server does not specify a managed SAP policy.</p> <p>The <i>policy-name</i> parameter is only valid for a SAP with the keywords <b>capture-sap</b> specified in the SAP's configuration. The <b>capture-sap</b> keyword in the SAP configuration captures the SAP where triggering packets will be sent to the CPM. Non-triggering packets captured by the capture SAP will be dropped.</p>

The managed SAP policy must already be defined in the **config>subscr-mgmt>msap-policy** context

The **no** form of the command removes the policy-name from the configuration.

<b>Default</b>	no default-msap-policy
<b>Parameters</b>	<i>policy-name</i> — /Specifies an existing default managed SAP policy.

## trigger-packet

<b>Syntax</b>	<b>trigger-packet [dhcp] [pppoe] [arp] [dhcp6] [ppp]</b> <b>no trigger-packet</b>
<b>Context</b>	config>service>vpls>sap
<b>Description</b>	This command enables triggering packet to initiate RADIUS authentication that provides a service context. The authentication, together with the service context for this request, creates a managed SAP. The VLAN is the same as the triggering packet. This SAP behaves as a regular SAP but the configuration is not user-editable and not maintained in the configuration file. The managed SAP remains active as long as the session is active.
<b>Default</b>	none
<b>Parameters</b>	<p><b>dhcp</b> — Specifies whether the receipt of DHCP trigger packets on this VPLS SAP when the keyword <b>capture-sap</b> is specified in the <b>sap</b> command creation string, will result in a RADIUS authentication that will provide a service context and the creation of a SAP with a value of 'managed'.</p> <p><b>pppoe</b> — Specifies whether the receipt of PPPoE trigger packets on this VPLS SAP when the keyword <b>capture-sap</b> is specified in the <b>sap</b> command creation string, will result in a RADIUS authentication that will provide a service context and the creation of a SAP with a value of 'managed'.</p> <p><b>arp</b> — Indicates that ARP is the type of trigger packets for this entry.</p> <p><b>dhcp6</b> — Indicates that DHCP6 is the type of trigger packets for this entry.</p> <p><b>ppp</b> — Indicates that PPP is the type of trigger packets for this entry.</p>

## eval-msap

<b>Syntax</b>	<b>eval-msap {policy <i>msap-policy-name</i>   msap <i>sap-id</i>}</b>
<b>Context</b>	tools>perform>subscr-mgmt
<b>Description</b>	This command evaluates managed SAP policies.
<b>Parameters</b>	<b>policy <i>msap-policy-name</i></b> — Specifies an existing MSAP policy.

**msap** *sap-id* — Specifies an MSAP sap-id.

**Values**      *[port-id|lag-id]:qtag1*  
                 *[port-id|lag-id]:qtag1.qtag2*

## Multi-Chassis Redundancy Commands

### redundancy

<b>Syntax</b>	<b>redundancy</b>
<b>Context</b>	config
<b>Description</b>	This command allows the user to perform redundancy operations.
<b>Parameters</b>	<b>force-switchover</b> — Forces a switchover to the standby CPM card
<b>Values</b>	<b>now</b> keyword - switch to standby CPM)
	<b>NOTE:</b> Switching to the standby displays the following message.
	WARNING: Configuration and/or Boot options may have changed since the last save. Are you sure you want to switchover (y/n)?
	<b>synchronize</b> — Synchronizes the secondary CPM.
<b>Values</b>	<b>boot-env config</b> : keywords

### synchronize

<b>Syntax</b>	<b>synchronize {boot-env   config}</b>
<b>Context</b>	config>redundancy
<b>Description</b>	<p>This command performs a synchronization of the standby CPM images and/or config files to the active CPM. Either the <b>boot-env</b> or <b>config</b> parameter must be specified.</p> <p>In the <b>config&gt;redundancy</b> context, this command performs an automatically triggered standby CPM synchronization.</p> <p>When the standby CPM takes over operation following a failure or reset of the active CPM, it is important to ensure that the active and standby CPM have identical operational parameters. This includes the saved configuration, CPM and IOM images.</p> <p>The active CPM ensures that the active configuration is maintained on the standby CPM. However, to ensure smooth operation under all circumstances, runtime images and system initialization configurations must also be automatically synchronized between the active and standby CPM.</p> <p>If synchronization fails, alarms and log messages that indicate the type of error that caused the failure of the synchronization operation are generated. When the error condition ceases to exist, the alarm is cleared.</p> <p>Only files stored on the router are synchronized. If a configuration file or image is stored in a location other than on a local compact flash, the file is not synchronized (for example, storing a configuration file on an FTP server).</p>
<b>Default</b>	enabled

## Multi-Chassis Redundancy Commands

<b>Parameters</b>	<b>boot-env</b> — Synchronizes all files required for the boot process (loader, BOF, images, and configuration files).
	<b>config</b> — Synchronize only the primary, secondary, and tertiary configuration files.
	<b>Default</b> config

### multi-chassis

<b>Syntax</b>	<b>multi-chassis</b>
<b>Context</b>	config>redundancy
<b>Description</b>	This command enables the context to configure multi-chassis parameters.

### peer

<b>Syntax</b>	<b>[no] peer <i>ip-address</i></b>
<b>Context</b>	config>redundancy>multi-chassis
<b>Description</b>	This command configures a multi-chassis redundancy peer.
<b>Parameters</b>	<i>ip-address</i> — Specifies a peer IP address. Multicast address are not allowed.

### authentication-key

<b>Syntax</b>	<b>authentication-key [<i>authentication-key</i>   <i>hash-key</i>] [<b>hash</b>   <b>hash2</b>] <b>no authentication-key</b></b>
<b>Context</b>	config>redundancy>multi-chassis>peer
<b>Description</b>	This command configures the authentication key used between this node and the multi-chassis peer. The authentication key can be any combination of letters or numbers.
<b>Parameters</b>	<i>authentication-key</i> — Specifies the authentication key. Allowed values are any string up to 20 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.
	<i>hash-key</i> — The hash key. The key can be any combination of ASCII characters up to 33 (hash1-key) or 55 (hash2-key) characters in length (encrypted). If spaces are used in the string, enclose the entire string in quotation marks (“ ”).
	<b>hash</b> — Specifies the key is entered in an encrypted form. If the hash or hash2 parameter is not used, the key is assumed to be in a non-encrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the hash or hash2 parameter specified.
	<b>hash2</b> — Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, this means that hash2 encrypted variable cannot be copied and pasted. If the hash or hash2 parameter is not used, the key is assumed to be in a non-encrypted, clear text



form. For security, all keys are stored in encrypted form in the configuration file with the hash or hash2 parameter specified.

## mc-ipsec

<b>Syntax</b>	<b>mc-ipsec</b>
<b>Context</b>	config>redundancy>multi-chassis>peer
<b>Description</b>	This command enters the configuration context of multi-chassis IPsec.

## discovery-interval

<b>Syntax</b>	<b>discovery-interval</b> <i>interval-1</i> [ <b>boot</b> <i>interval-2</i> ] <b>no discovery-interval</b>
<b>Context</b>	config>redundancy>multi-chassis>peer>mc-ipsec
<b>Description</b>	This command specifies the time interval of 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 at system boot-up. It is optional and when it is not specified, interval-1 will be used.
<b>Default</b>	300
<b>Parameters</b>	<i>interval-1/2</i> — Specifies the interval in seconds.
<b>Values</b>	1..1800 seconds

## keep-alive-interval

<b>Syntax</b>	<b>keep-alive-interval</b> <i>time-interval</i> <b>no keep-alive-interval</b>
<b>Context</b>	config>redundancy>multi-chassis>peer>mc-ipsec
<b>Description</b>	This command specifies the time interval of the mastership election protocol sending the keep-alive packet.
<b>Default</b>	10
<b>Parameters</b>	<i>time-interval</i> — Specifies the time interval in tenths of a second.
<b>Values</b>	5..500

## hold-on-neighbor-failure

<b>Syntax</b>	<b>hold-on-neighbor-failure</b> <i>multiplier</i>
---------------	---

### **no hold-on-neighbor-failure**

<b>Context</b>	config>redundancy>multi-chassis>peer>mc-ipsec
<b>Description</b>	This command specifies the number of keep-alive failures before the peer is considered down.
<b>Default</b>	3
<b>Parameters</b>	<i>multiplier</i> — Specifies the multiplier.
<b>Values</b>	2..25

## bfd-enable

<b>Syntax</b>	<b>bfd-enable service <i>service-id</i> interface <i>interface-name</i> dst-ip <i>ip-address</i></b> <b>no bfd-enable</b>
<b>Context</b>	config>redundancy>multi-chassis>peer>mc-ipsec
<b>Description</b>	<p>This command enables tracking a central BFD session. If the BFD session goes down, then the system considers the peer down and changes the mc-ipsec status of the configured tunnel-group accordingly.</p> <p>The BFD session uses the specified loopback interface (in the specified service) address as the source address and uses the specified dst-ip as the destination address. Other BFD parameters are configured with the “bfd” command on the specified interface.</p>
<b>Parameters</b>	<i>interface-name</i> — Specifies the name of the loopback interface. <i>service-id</i> — Specifies the ID of the service. <i>dst-id</i> — Specifies the destination address of the BFD packet.

## tunnel-group

<b>Syntax</b>	<b>tunnel-group <i>group-id</i> [create]</b> <b>no tunnel-group <i>group-id</i></b>
<b>Context</b>	config>redundancy>multi-chassis>peer>mc-ipsec
<b>Description</b>	This command enables multi-chassis redundancy for the specified tunnel-group or enters an already configured tunnel-group context. The configured tunnel-group could failover independently.
<b>Parameters</b>	<i>group-id</i> — Specifies the tunnel-group ID.
<b>Values</b>	1..16
	<b>create</b> — Enables multi-chassis redundancy for the specified tunnel-group.

## peer-group

<b>Syntax</b>	<b>peer-group <i>group-id</i></b>
---------------	-----------------------------------

**no peer-group**

<b>Context</b>	config>redundancy>multi-chassis>peer>mc-ipsec>tunnel-group
<b>Description</b>	This command specifies the corresponding tunnel-group ID on the peer node. The peer tunnel-group ID does not necessarily equal the local tunnel-group ID.
<b>Parameters</b>	<i>group-id</i> — Specifies the tunnel-group ID.
<b>Values</b>	1..16

## priority

<b>Syntax</b>	<b>priority</b> <i>priority</i> <b>no priority</b>
<b>Context</b>	config>redundancy>multi-chassis>peer>mc-ipsec>tunnel-group
<b>Description</b>	This command specifies the local priority of the tunnel-group. This is used to elect the master (higher number is the master). If priorities are the same, then the peer with the more active ISA becomes the master. If the priority and the number of active ISAs are the same, then the peer with the higher IP address is the master.
<b>Parameters</b>	<i>priority</i> — Specifies the priority of the tunnel-group.
<b>Values</b>	0..255

## preempt

<b>Syntax</b>	<b>[no] preempt</b>
<b>Context</b>	config>redundancy>multi-chassis>peer>mc-ipsec>tunnel-group
<b>Description</b>	This command enables the preempt behavior of local node.

## mc-lag

<b>Syntax</b>	<b>[no] mc-lag</b>
<b>Context</b>	config>redundancy>multi-chassis>peer>mc-lag
<b>Description</b>	This command enables the context to configure multi-chassis LAG operations and related parameters. The <b>no</b> form of this command administratively disables multi-chassis LAG. MC-LAG can only be issued only when mc-lag is shutdown.

## hold-on-neighbor-failure

<b>Syntax</b>	<b>hold-on-neighbor-failure</b> <i>multiplier</i>
---------------	---

### **no hold-on-neighbor-failure**

<b>Context</b>	config>redundancy>multi-chassis>peer>mc-lag
<b>Description</b>	<p>This command specifies the interval that the standby node will wait for packets from the active node before assuming a redundant-neighbor node failure. This delay in switch-over operation is required to accommodate different factors influencing node failure detection rate, such as IGP convergence, or HA switch-over times and to prevent the standby node to take action prematurely.</p> <p>The <b>no</b> form of this command sets this parameter to default value.</p>
<b>Default</b>	3
<b>Parameters</b>	<p><i>multiplier</i> — The time interval that the standby node will wait for packets from the active node before assuming a redundant-neighbor node failure.</p> <p><b>Values</b> 2 — 25</p>

## keep-alive-interval

<b>Syntax</b>	<b>keep-alive-interval</b> <i>interval</i> <b>no keep-alive-interval</b>
<b>Context</b>	config>redundancy>multi-chassis>peer>mc-lag
<b>Description</b>	<p>This command sets the interval at which keep-alive messages are exchanged between two systems participating in MC-LAG. These keep-alive messages are used to determine remote-node failure and the interval is set in deci-seconds.</p> <p>The no form of this command sets the interval to default value</p>
<b>Default</b>	1s (10 hundreds of milliseconds means interval value of 10)
<b>Parameters</b>	<p><i>interval</i> — The time interval expressed in deci-seconds</p> <p><b>Values</b> 5 — 500</p>

## lag

<b>Syntax</b>	<b>lag</b> <i>lag-id</i> <b>lACP-key</b> <i>admin-key</i> <b>system-id</b> <i>system-id</i> [ <b>remote-lag</b> <i>lag-id</i> ] <b>system-priority</b> <i>system-priority</i> <b>no lag</b> <i>lag-id</i>
<b>Context</b>	config>redundancy>multi-chassis>peer>mc-lag
<b>Description</b>	<p>This command defines a LAG which is forming a redundant-pair for MC-LAG with a LAG configured on the given peer. The same LAG group can be defined only in the scope of 1 peer.</p> <p>The same <b>lACP-key</b>, <b>system-id</b>, and <b>system-priority</b> must be configured on both nodes of the redundant pair in order to MC-LAG to become operational. In order MC-LAG to become operational, all parameters (<b>lACP-key</b>, <b>system-id</b>, <b>system-priority</b>) must be configured the same on both nodes of the same redundant pair.</p>

The partner system (the system connected to all links forming MC-LAG) will consider all ports using the same **lACP-key**, **system-id**, **system-priority** as the part of the same LAG. In order to achieve this in MC operation, both redundant-pair nodes have to be configured with the same values. In case of the mismatch, MC-LAG is kept in oper-down status.

<b>Default</b>	none
<b>Parameters</b>	<p><i>lag-id</i> — The LAG identifier, expressed as a decimal integer. Specifying the <i>lag-id</i> allows the mismatch between lag-id on redundant-pair. If no <b>lag-id</b> is specified it is assumed that neighbor system uses the same <i>lag-id</i> as a part of the given MC-LAG. If no matching MC-LAG group can be found between neighbor systems, the individual LAGs will operate as usual (no MC-LAG operation is established.).</p> <p><b>Values</b> 1 — 800</p> <p><b>lACP-key</b> <i>admin-key</i> — Specifies a 16 bit key that needs to be configured in the same manner on both sides of the MC-LAG in order for the MC-LAG to come up.</p> <p><b>Values</b> 1 — 65535</p> <p><b>system-id</b> <i>system-id</i> — Specifies a 6 byte value expressed in the same notation as MAC address</p> <p><b>Values</b> xx:xx:xx:xx:xx:xx - xx [00..FF]</p> <p><b>remote-lag</b> <i>lag-id</i> — Specifies the LAG ID on the remote system.</p> <p><b>Values</b> 1 — 200</p> <p><b>system-priority</b> <i>system-priority</i> — Specifies the system priority to be used in the context of the MC-LAG. The partner system will consider all ports using the same <b>lACP-key</b>, <b>system-id</b>, and <b>system-priority</b> as part of the same LAG.</p> <p><b>Values</b> 1 — 65535</p>

## source-address

<b>Syntax</b>	<b>source-address</b> <i>ip-address</i> <b>no source-address</b>
<b>Context</b>	config>redundancy>multi-chassis>peer
<b>Description</b>	This command specifies the source address used to communicate with the multi-chassis peer.
<b>Parameters</b>	<i>ip-address</i> — Specifies the source address used to communicate with the multi-chassis peer.

## sync

<b>Syntax</b>	[no] <b>sync</b>
<b>Context</b>	config>redundancy>multi-chassis>peer
<b>Description</b>	This command enables the context to configure synchronization parameters.

### igmp

<b>Syntax</b>	<b>[no] igmp</b>
<b>Context</b>	config>redundancy>multi-chassis>peer>sync
<b>Description</b>	This command specifies whether IGMP protocol information should be synchronized with the multi-chassis peer.
<b>Default</b>	no igmp

### igmp-snooping

<b>Syntax</b>	<b>[no] igmp-snooping</b>
<b>Context</b>	config>redundancy>multi-chassis>peer>sync
<b>Description</b>	This command specifies whether IGMP snooping information should be synchronized with the multi-chassis peer.
<b>Default</b>	no igmp-snooping

### local-dhcp-server

<b>Syntax</b>	<b>[no] local-dhcp-server</b>
<b>Context</b>	config>redundancy>multi-chassis>peer>sync
<b>Description</b>	This command synchronizes DHCP server information.

### mc-ring

<b>Syntax</b>	<b>[no] mc-ring</b>
<b>Context</b>	config>redundancy>multi-chassis>peer>sync
<b>Description</b>	This command synchronizes mc-ring information.

### mld-snooping

<b>Syntax</b>	<b>[no] mld-snooping</b>
<b>Context</b>	config>redundancy>multi-chassis>peer>sync
<b>Description</b>	This command synchronizes MLD snooping information.

## port

<b>Syntax</b>	<b>port</b> [ <i>port-id</i>   <i>lag-id</i> ] [ <b>sync-tag</b> <i>sync-tag</i> ] [ <b>create</b> ] <b>no port</b> [ <i>port-id</i>   <i>lag-id</i> ]
<b>Context</b>	config>redundancy>multi-chassis>peer>sync
<b>Description</b>	This command specifies the port to be synchronized with the multi-chassis peer and a synchronization tag to be used while synchronizing this port with the multi-chassis peer.
<b>Parameters</b>	<i>port-id</i> — Specifies the port to be synchronized with the multi-chassis peer. <i>lag-id</i> — Specifies the LAG ID to be synchronized with the multi-chassis peer. <b>sync-tag</b> <i>sync-tag</i> — Specifies a synchronization tag to be used while synchronizing this port with the multi-chassis peer.

## range

<b>Syntax</b>	<b>range</b> <i>encap-range</i> <b>sync-tag</b> <i>sync-tag</i> <b>no range</b> <i>encap-range</i>						
<b>Context</b>	config>redundancy>multi-chassis>peer>sync>port						
<b>Description</b>	This command configures a range of encapsulation values.						
<b>Parameters</b>	<i>encap-range</i> — Specifies a range of encapsulation values on a port to be synchronized with a multi-chassis peer.  <table><tr><td><b>Values</b></td><td>Dot1Q</td><td><i>start-vlan-end-vlan</i></td></tr><tr><td></td><td>QinQ</td><td><i>Q1.start-vlan-Q1.end-vlan</i></td></tr></table> <b>sync-tag</b> <i>sync-tag</i> — specifies a synchronization tag up to 32 characters in length to be used while synchronizing this encapsulation value range with the multi-chassis peer.	<b>Values</b>	Dot1Q	<i>start-vlan-end-vlan</i>		QinQ	<i>Q1.start-vlan-Q1.end-vlan</i>
<b>Values</b>	Dot1Q	<i>start-vlan-end-vlan</i>					
	QinQ	<i>Q1.start-vlan-Q1.end-vlan</i>					

## srrp

<b>Syntax</b>	[ <b>no</b> ] <b>srrp</b>
<b>Context</b>	config>redundancy>multi-chassis>peer>sync
<b>Description</b>	This command specifies whether subscriber routed redundancy protocol (SRRP) information should be synchronized with the multi-chassis peer.
<b>Default</b>	no srrp

## sub-host-trk

<b>Syntax</b>	[ <b>no</b> ] <b>sub-host-trk</b>
<b>Context</b>	config>redundancy>multi-chassis>peer>sync

## Multi-Chassis Redundancy Commands

**Description** This command synchronizes subscriber host tracking information.

### sub-mgmt

**Syntax** **sub-mgmt** [ipoe | pppoe]  
**no sub-mgmt**

**Context** config>redundancy>multi-chassis>peer>sync

**Description** This command will enable synchronization of subscriber states between chassis. Synchronization will be enabled per protocol type (IPoE or PPPoE).  
The keywords (**ipoe**, **pppoe**) must match on both nodes. If not, subscriber synchronization will fail.  
For example if one node is configured with:  
config>multi-chassis>peer>sync>sub-mgmt ipoe  
but the other node is configured with:  
config>multi-chassis>peer>sync>sub-mgmt ipoe pppoe  
synchronization will fail even for ipoe application.

**Default** no sub-mgmt

**Parameters** **ipoe** — ipoe subscribers will be synchronized  
**pppoe** — pppoe subscribers will be synchronized

### tunnel-group

**Syntax** **tunnel-group** *tunnel-group-id* **sync-tag** *tag-name* [create]  
**no tunnel-group**

**Context** config>redundancy>multi-chassis>peer>sync

**Description** This command enables multi-chassis synchronization of IPsec states of a specified tunnel-group with its peer. Sync-tag is used to match corresponding tunnel-groups on both peers. IPsec states will be synchronized between tunnel-groups with the same sync-tag.

**Parameters** *tunnel-group-id* — Specifies the ID of the tunnel-group  
*tag-name* — Specifies the name of sync-tag.

### ipsec

**Syntax** [no] ipsec

**Context** config>redundancy>multi-chassis>peer>sync

**Description** This command enables multi-chassis synchronization of IPsec states on system level.



## mc-ring

<b>Syntax</b>	<b>mc-ring</b>
<b>Context</b>	config>redundancy>multi-chassis>peer
<b>Description</b>	This command enables the context to configure the multi-chassis ring parameters.
<b>Default</b>	mc-ring

## ring

<b>Syntax</b>	<b>[no] ring sync-tag [create]</b>
<b>Context</b>	config>redundancy>multi-chassis>peer>mcr
<b>Description</b>	This command configures a multi-chassis ring. The <b>no</b> form of the command removes the sync-tag from the configuration.
<b>Default</b>	none

## l3-ring

<b>Syntax</b>	<b>[no] l3-ring name [create]</b>
<b>Context</b>	config>redundancy>multi-chassis>peer>mcr
<b>Parameters</b>	This command configures a layer 3 multi-chassis ring.

## in-band-control-path

<b>Syntax</b>	<b>in-band-control-path</b>
<b>Context</b>	config>redundancy>multi-chassis>peer>mcr>ring config>redundancy>multi-chassis>peer>mc>l3-ring
<b>Description</b>	This command enables the context to configure control path parameters.
<b>Default</b>	none

## debounce

<b>Syntax</b>	<b>[no] debounce</b>
<b>Context</b>	config>redundancy>multi-chassis>peer>mcr>ring>in-band-control-path config>redundancy>multi-chassis>peer>mc>l3-ring>in-band-control-path

## Multi-Chassis Redundancy Commands

**Description** This command enables the inband control path debouncing. The **no** form of the command disables inband control path debouncing.

### dst-ip

**Syntax** **dst-ip** *ip-address*  
**no dst-ip**

**Context** config>redundancy>multi-chassis>peer>mcr>ring>in-band-control-path  
config>redundancy>multi-chassis>peer>mc>l3-ring>in-band-control-path

**Description** This command specifies the destination IP address used in the inband control connection.  
If the destination IP address is not configured, the ring cannot become operational.

**Default** none

**Parameters** *ip-address* — The destination IP address.

### interface

**Syntax** **interface** *ip-int-name*  
**no interface**

**Context** config>redundancy>multi-chassis>peer>mcr>ring>in-band-control-path  
config>redundancy>multi-chassis>peer>mc>l3-ring>in-band-control-path

**Description** This command specifies the name of the IP interface used for the inband control connection.  
If an interface name is not configured, the ring cannot become operational.

**Parameters** *ip-int-name* — Specifies an interface name up to 32 characters in length.

### max-debounce-time

**Syntax** **max-debounce-time** *max-debounce-time*  
**no max-debounce-time**

**Context** config>redundancy>multi-chassis>peer>mcr>ring>in-band-control-path  
config>redundancy>multi-chassis>peer>mc>l3-ring>in-band-control-path

**Description** This command configures the inband control path maximum debounce time.

**Parameters** *max-debounce-time* — Specifies the maximum debounce time on the transition of the operational state of the inband control connection.

**Values** 5 — 200 seconds

## service-id

<b>Syntax</b>	<b>service-id</b> <i>service-id</i> <b>no service-id</b>
<b>Context</b>	config>redundancy>multi-chassis>peer>mcr>ring>in-band-control-path config>redundancy>multi-chassis>peer>mc>l3-ring>in-band-control-path
<b>Description</b>	This command configures the service ID of the SAP used for the Ring-Node Connectivity Verification of this ring node.
<b>Parameters</b>	<i>service-id</i> — [Specifies an existing service ID or service name.
<b>Values</b>	service-id: 1 — 214748364 svc-name: A string up to 64 characters in length.

## path-b

<b>Syntax</b>	<b>[no] path-b</b>
<b>Context</b>	config>redundancy>multi-chassis>peer>mcr>ring
<b>Description</b>	This command specifies the set of upper-VLAN IDs associated with the SAPs that belong to path B with respect to load-sharing. All other SAPs belong to path A.
<b>Default</b>	If not specified, the default is an empty set.

## range

<b>Syntax</b>	<b>[no] range</b> <i>vlan-range</i>
<b>Context</b>	config>redundancy>multi-chassis>peer>mcr>ring>path-b config>redundancy>multi-chassis>peer>mcr>ring>path-excl
<b>Description</b>	This command specifies the set of VLAN IDs associated with the SAPs that are controlled by the remote peer. It is a bitmap that associates bit <i>i</i> with VLAN ID <i>i</i> , with <i>i</i> in [0..4094]. Setting the value to the empty string is equivalent to setting it to 512 zeroes.

## ring-node

<b>Syntax</b>	<b>[no] ring-node</b> <i>ring-node-name</i>
<b>Context</b>	config>redundancy>mc>peer>mcr>ring
<b>Description</b>	This command specifies the unique name of a multi-chassis ring access node.

## path-excl

## Multi-Chassis Redundancy Commands

<b>Syntax</b>	<b>[no] path-excl</b>
<b>Context</b>	config>redundancy>multi-chassis>peer>mcr>ring
<b>Description</b>	This command specifies the set of upper-VLAN IDs associated with the SAPs that are to be excluded from control by the multi-chassis ring.
<b>Default</b>	If not specified, the default is an empty set.

### connectivity-verify

<b>Syntax</b>	<b>connectivity-verify</b>
<b>Context</b>	config>redundancy>multi-chassis>peer>mcr>ring config>redundancy>multi-chassis>peer>mc>l3-ring
<b>Description</b>	This command configures the node connectivity check.

### interval

<b>Syntax</b>	<b>interval</b> <i>interval</i>
<b>Context</b>	config>redundancy>multi-chassis>peer>mcr>ring>>connectivity-verify config>redundancy>multi-chassis>peer>mc>l3-ring>connectivity-verify
<b>Description</b>	This command specifies the polling interval of the ring-node connectivity verification of this ring node.
<b>Parameters</b>	<i>interval</i> — Specifies the polling interval of the ring-node connectivity verification of this ring node.
<b>Values</b>	1 — 6000

### service-id

<b>Syntax</b>	<b>service-id</b> <i>service-id</i> <b>no service-id</b>
<b>Context</b>	config>redundancy>mc>peer>mcr>ring-node>connect-verify config>redundancy>multi-chassis>peer>mc>l3-ring>connectivity-verify
<b>Description</b>	This command specifies the service ID of the SAP used for ring-node connectivity verification of this ring node.
<b>Parameters</b>	<i>service-id</i> — Specifies the service ID or service name.
<b>Values</b>	service-id: 1 — 214748364 svc-name: A string up to 64 characters in length.

### src-ip

<b>Syntax</b>	<b>src-ip</b> <i>ip-address</i> <b>no src-ip</b>
<b>Context</b>	config>redundancy>mc>peer>mcr>ring-node>connect-verify config>redundancy>multi-chassis>peer>mc>l3-ring>connectivity-verify
<b>Description</b>	This command specifies the source IP address used in ring-node connectivity verification of this ring node.
<b>Parameters</b>	<i>ip-address</i> — Specifies the source IP address used in ring-node connectivity verification of this ring node.

## src-mac

<b>Syntax</b>	<b>src-mac</b> <i>ieee-address</i> <b>no src-mac</b>
<b>Context</b>	config>redundancy>mc>peer>mcr>ring-node>connect-verify config>redundancy>multi-chassis>peer>mc>l3-ring>connectivity-verify
<b>Description</b>	This command specifies the source MAC address used for the Ring-Node Connectivity Verification of this ring node.  If all zeros are specified, then the MAC address of the system management processor (CPM) is used.
<b>Parameters</b>	<i>ieee-address</i> — Specifies the 48-bit MAC address for the static ARP in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee, and ff are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

## vlan

<b>Syntax</b>	<b>vlan</b> [0..4094]
<b>Context</b>	config>redundancy>mc>peer>mcr>ring-node>connect-verify config>redundancy>mc>peer>mcr>l3ring>node>cv
<b>Description</b>	This command specifies the VLAN tag of the SAP used for ring-node connectivity verification of this ring node. It is only meaningful if the value of is not zero.

## srrp-instance

<b>Syntax</b>	<b>[no] srrp-instance</b> <i>srrp-id</i>
<b>Context</b>	config>redundancy>multi-chassis>peer>mc>l3-ring
<b>Description</b>	This command configures an SRRP instance for layer 3 ring.
<b>Parameters</b>	<i>srrp-id</i> — Specifies the SRRP ID of this SRRP instance.  <b>Values</b> 1 — 4294967295

---

## SLA Profile Commands

### sla-profile

<b>Syntax</b>	<b>sla-profile</b> <i>sla-profile-name</i>
<b>Context</b>	config>subscr-mgmt
<b>Description</b>	<p>This command configures an SLA profile mapping. Hosts associated with a subscriber are subdivided into Service Level Agreement (SLA) profiles. For each subscriber host an SLA profile can be specified. For a subscriber host, the SLA profile determines:</p> <ul style="list-style-type: none"><li>• The QoS-policies to use<ul style="list-style-type: none"><li>–The classification</li><li>–The queues</li><li>–The queue mapping</li></ul></li><li>• The IP filters to use</li></ul> <p>The SLA profile also has the attribute host-limit which limits the total number of hosts (belonging to the same subscriber) on a certain SAP that can be using this SLA profile.</p>
<b>Default</b>	none
<b>Parameters</b>	<i>sla-profile-name</i> — Specifies the name of the SLA profile.

### egress

<b>Syntax</b>	<b>egress</b>
<b>Context</b>	config>subscr-mgmt>sla-profile
<b>Description</b>	This command configures egress parameters for the SLA profile.

### ingress

<b>Syntax</b>	<b>ingress</b>
<b>Context</b>	config>subscr-mgmt>sla-profile
<b>Description</b>	This command configures ingress parameters for the SLA profile.

### host-limits

<b>Syntax</b>	<b>[no] no host-limits</b>
---------------	----------------------------

<b>Context</b>	config>subscr-mgmt>sla-profile
<b>Description</b>	This command specifies the maximum number of hosts for this SLA profile.
<b>Parameters</b>	<i>max-number-of-hosts</i> — Specifies the host limit for this SLA profile. <b>Values</b> 1 — 100 <b>remove-oldest</b> — When the keywords <b>remove-oldest</b> are specified, the oldest subscriber host will be removed when the host limit is reached.

## ipv4-arp

<b>Syntax</b>	<b>ipv4-arp</b> <i>max-nr-of-hosts</i> <b>no ipv4-arp</b>
<b>Context</b>	config>subscr-mgmt>sla-profile>host-limits
<b>Description</b>	This command configures the maximum number of IPv4 ARP hosts. The <b>no</b> form of the command removes the number of IPv4 ARP hosts from the SLA profile.
<b>Default</b>	no ipv4-arp
<b>Parameters</b>	<i>max-nr-of-hosts</i> — Specifies the maximum number of IPv4 ARP hosts. <b>Values</b> 0 — 32767

## ipv4-dhcp

<b>Syntax</b>	<b>ipv4-dhcp</b> <i>max-nr-of-hosts</i> <b>no ipv4-dhcp</b>
<b>Context</b>	config>subscr-mgmt>sla-profile>host-limits
<b>Description</b>	This command limits the number of IPv4 DHCP hosts. The <b>no</b> form of the command removes the number of IPv4 DHCP hosts from the SLA profile.
<b>Default</b>	no ipv4-dhcp
<b>Parameters</b>	<i>max-nr-of-hosts</i> — Specifies the maximum number of IPv4 DHCP hosts. <b>Values</b> 0 — 32767

## ipv4-overall

<b>Syntax</b>	<b>ipv4-overall</b> <i>max-nr-of-hosts</i> <b>no ipv4-overall</b>
<b>Context</b>	config>subscr-mgmt>sla-profile>host-limits
<b>Description</b>	This command limits the total number of IPv4 hosts.

The **no** form of the command removes the number of IPv4 hosts from the SLA profile.

<b>Default</b>	no ipv4-overall
<b>Parameters</b>	<i>max-nr-of-hosts</i> — Specifies the maximum number of IPv4 hosts.
<b>Values</b>	0 — 32767

### ipv4-ppp

<b>Syntax</b>	<b>ipv4-ppp</b> <i>max-nr-of-hosts</i> <b>no ipv4-ppp</b>
<b>Context</b>	config>subscr-mgmt>sla-profile>host-limits
<b>Description</b>	This command limits the total number of IPv4 PPP hosts. The <b>no</b> form of the command removes the number of IPv4 PPP hosts from the SLA profile.
<b>Default</b>	no ipv4-ppp
<b>Parameters</b>	<i>max-nr-of-hosts</i> — Specifies the maximum number of IPv4 PPP hosts.
<b>Values</b>	0 — 32767

### ipv6-overall

<b>Syntax</b>	<b>ipv6-overall</b> <i>max-nr-of-hosts</i> <b>no ipv6-overall</b>
<b>Context</b>	config>subscr-mgmt>sla-profile>host-limits
<b>Description</b>	This command limits the total number of IPv6 hosts. The <b>no</b> form of the command removes the number of IPv6 hosts from the SLA profile.
<b>Default</b>	no ipv6-overall
<b>Parameters</b>	<i>max-nr-of-hosts</i> — Specifies the maximum number of IPv6 hosts.
<b>Values</b>	0 — 32767

### ipv6-pd-ipoe-dhcp

<b>Syntax</b>	<b>ipv6-pd-ipoe-dhcp</b> <i>max-nr-of-hosts</i> <b>no ipv6-pd-ipoe-dhcp</b>
<b>Context</b>	config>subscr-mgmt>sla-profile>host-limits
<b>Description</b>	This command configures the total number of IPv6 DHCP PD hosts. The <b>no</b> form of the command removes the number of IPv6 DHCP hosts from the SLA profile.



<b>Default</b>	no ipv6-dhcp
<b>Parameters</b>	<i>max-nr-of-hosts</i> — Specifies the total number of IPv6 DHCP PD hosts.
<b>Values</b>	0 — 32767

## ipv6-pd-overall

<b>Syntax</b>	<b>ipv6-pd-overall</b> <i>max-nr-of-hosts</i> <b>no ipv6-pd-overall</b>
<b>Context</b>	config>subscr-mgmt>sla-profile>host-limits
<b>Description</b>	This command limits the total number of IPv6-PD hosts. The <b>no</b> form of the command removes the number of IPv6-PD hosts from the SLA profile.
<b>Default</b>	no ipv6-pd-overall
<b>Parameters</b>	<i>max-nr-of-hosts</i> — Specifies the maximum number of IPv6-PD hosts overall.
<b>Values</b>	0 — 32767

## ipv6-pd-ppp-dhcp

<b>Syntax</b>	<b>ipv6-pd-ppp-dhcp</b> <i>max-nr-of-hosts</i> <b>no ipv6-pd-ppp-dhcp</b>
<b>Context</b>	config>subscr-mgmt>sla-profile>host-limits
<b>Description</b>	This command configures the maximum number of IPv6-WAN PPP DHCP hosts. The <b>no</b> form of the command removes the number of IPv6-WAN PPP DHCP hosts from the SLA profile.
<b>Default</b>	no ipv6-pd-ppp-dhcp
<b>Parameters</b>	<i>max-nr-of-hosts</i> — Specifies the maximum number of IPv6-WAN PPP DHCP hosts.
<b>Values</b>	0 — 32767

## ipv6-wan-ipoe-dhcp

<b>Syntax</b>	<b>ipv6-wan-ipoe-dhcp</b> <i>max-nr-of-hosts</i> <b>no ipv6-wan-ipoe-dhcp</b>
<b>Context</b>	config>subscr-mgmt>sla-profile>host-limits
<b>Description</b>	This command configures the maximum number of IPv6-WAN PPP DHCP hosts. The <b>no</b> form of the command removes the number of IPv6-WAN PPP DHCP hosts from the SLA profile.

<b>Default</b>	no ipv6-wan-ipoe-dhcp
<b>Parameters</b>	<i>max-nr-of-hosts</i> — Specifies the maximum number of IPv6-WAN PPP DHCP hosts.
<b>Values</b>	0 — 32767

### ipv6-wan-ipoe-slaac

<b>Syntax</b>	<b>ipv6-wan-ipoe-slaac</b> <i>max-nr-of-hosts</i> <b>no ipv6-wan-ipoe-slaac</b>
<b>Context</b>	config>subscr-mgmt>sla-profile>host-limits
<b>Description</b>	This command configures the maximum number of IPv6-WAN IPOE SLAAC hosts. The <b>no</b> form of the command removes the number of IPv6-WAN IPOE SLAAC hosts from the SLA profile.
<b>Default</b>	no ipv6-wan-ipoe-slaac
<b>Parameters</b>	<i>max-nr-of-hosts</i> — Specifies the maximum number of IPv6-WAN IPOE SLAAC hosts.
<b>Values</b>	0 — 32767

### ipv6-wan-overall

<b>Syntax</b>	<b>ipv6-wan-overall</b> <i>max-nr-of-hosts</i> <b>no ipv6-wan-overall</b>
<b>Context</b>	config>subscr-mgmt>sla-profile>host-limits
<b>Description</b>	This command configures the total number of IPV6 WAN hosts. The <b>no</b> form of the command removes the number of IPV6 WAN hosts from the SLA profile.
<b>Default</b>	no ipv6-wan-overall
<b>Parameters</b>	<i>max-nr-of-hosts</i> — Specifies the maximum number of IPV6 WAN hosts.
<b>Values</b>	0 — 32767

### ipv6-wan-ppp-dhcp

<b>Syntax</b>	<b>ipv6-wan-ppp-dhcp</b> <i>max-nr-of-hosts</i> <b>no ipv6-wan-ppp-dhcp</b>
<b>Context</b>	config>subscr-mgmt>sla-profile>host-limits
<b>Description</b>	This command configures the total number of IPV6 PPP DHCP WAN hosts. The <b>no</b> form of the command removes the number of IPV6 PPP DHCP WAN hosts from the SLA profile.

<b>Default</b>	no ipv6-wan-ppp-dhcp
<b>Parameters</b>	<i>max-nr-of-hosts</i> — Specifies the maximum number of IPV6 PPP DHCP WAN hosts.
<b>Values</b>	0 — 32767

## ipv6-wan-ppp-slaac

<b>Syntax</b>	<b>ipv6-wan-ppp-slaac</b> <i>max-nr-of-hosts</i> <b>no ipv6-wan-ppp-slaac</b>
<b>Context</b>	config>subscr-mgmt>sla-profile>host-limits
<b>Description</b>	This command configures the total number of SLAAC hosts. The <b>no</b> form of the command removes the number of SLAAC hosts from the SLA profile.
<b>Default</b>	no ipv6-wan-ppp-slaac
<b>Parameters</b>	<i>max-nr-of-hosts</i> — Specifies the maximum number of SLAAC hosts.
<b>Values</b>	0 — 32767

## lac-overall

<b>Syntax</b>	<b>lac-overall</b> <i>max-nr-of-hosts</i> <b>no lac-overall</b>
<b>Context</b>	config>subscr-mgmt>sla-profile>host-limits
<b>Description</b>	This command configures the total number of L2TP LAC hosts The <b>no</b> form of the command removes the number of L2TP LAC from the SLA profile.
<b>Default</b>	no lac-overall
<b>Parameters</b>	<i>max-nr-of-hosts</i> — Specifies the maximum number of L2TP LAC hosts.
<b>Values</b>	0 — 32767

## overall

<b>Syntax</b>	<b>overall</b> <i>max-nr-of-hosts</i> <b>no overall</b>
<b>Context</b>	config>subscr-mgmt>sla-profile>host-limits
<b>Description</b>	This command configures the total number of hosts. The <b>no</b> form of the command reverts to the default.
<b>Default</b>	no overall

## SLA Profile Commands

**Parameters**    *max-nr-of-hosts* — Specifies the maximum number of hosts.

**Values**        0 — 32767

### remove-oldest

**Syntax**        [no] **remove-oldest**

**Context**        config>subscr-mgmt>sla-profile>host-limits

**Description**    This command removes the oldest subscriber host when the host limit is reached.  
The **no** form of the command maintains the oldest subscriber host when the host limit is reached.

**Default**        no remove-oldest

### ip-filter

**Syntax**        [no] **ip-filter** *filter-id*

**Context**        config>subscr-mgmt>sla-profile>egress  
config>subscr-mgmt>sla-profile>ingress

**Description**    This command configures an egress or ingress IP filter.

**Parameters**    *filter-id* — Specify an existing IP filter policy ID.

**Values**        1 — 65535

## SLA Profile QoS Commands

### qos

<b>Syntax</b>	<b>qos</b> <i>sap-egress-policy-id</i> [ <i>vport-scheduler</i>   <i>port-scheduler</i> ] [ <b>force</b> ] <b>no qos</b>
<b>Context</b>	config>subscr-mgmt>sla-prof>egress
<b>Description</b>	This command specifies the egress QoS policy applicable to this SLA profile. The policy must already be defined in the <b>configure&gt;qos&gt;sap-egress</b> context.
<b>Default</b>	1
<b>Parameters</b>	<i>sap-egress-policy-id</i> — Specifies the egress policy to be applied to the egress SLA profile. <b>Values</b> 1 — 65535 <i>vport-scheduler</i>   <i>port-scheduler</i> — Specifies if a host queue with the port-parent option enabled should be scheduled within the context of a vport port scheduler policy or a the port's port scheduler policy. <b>force</b> — Forces a policy change.

### qos

<b>Syntax</b>	<b>qos</b> <i>policy-id</i> [ <b>shared-queuing</b>   <b>multipoint-shared</b>   <b>service-queuing</b> ] [ <b>force</b> ] <b>no qos</b>
<b>Context</b>	config>subscr-mgmt>sla-prof>ingress
<b>Description</b>	This command specifies the ingress QoS policy applicable to this SLA profile. The policy must already be defined in the <b>configure&gt;qos&gt;sap-ingress</b> context.
<b>Default</b>	qos 1
<b>Parameters</b>	<i>sap-ingress-policy-id</i> — Specifies the policy to be applied to the ingress SLA profile. <b>Values</b> 1 — 65535 <b>shared-queuing</b> — This keyword is mutually exclusive with the <b>multipoint-shared</b> and <b>service-queuing</b> keywords to specify the policy used by this SAP. When the value of this object is null it means that the SAP will use individual ingress QoS queues, instead of the shared ones. <b>multipoint-shared</b> — This keyword is mutually exclusive with the <b>shared-queuing</b> and <b>service-queuing</b> keywords. When multipoint-shared is specified, the ingress forwarding plane will conserve hardware queues by performing two tier queuing on ingress unicast and multipoint packets through the SAP. Unicast service queues defined in the SAP ingress QoS policy are created for the SAP on the ingress forwarding plane without regard for the switch fabric destinations to which the SAP may need to forward (other destinations in the VPLS context). The multipoint queues defined in the SAP ingress QoS policy are not created for the SAP. Instead, all multipoint traffic is mapped to the unicast queues based on forwarding class in the first pass. In

the second pass the unicast packets will be mapped to the unicast shared queues while the multipoint traffic will be mapped to the multipoint shared queues.

**service-queueing** — This keyword is mutually exclusive with the **multipoint-shared** and **shared-queueing** keywords to state that service queueing is needed.

**force** — Forces a policy change.

## queue

<b>Syntax</b>	<b>[no] queue</b> <i>queue-id</i>
<b>Context</b>	config>subscr-mgmt>sla-prof>egress>qos config>subscr-mgmt>sla-prof>ingress>qos
<b>Description</b>	<p>This command configures the context to configure egress or ingress queue parameters. Parameters defined in the <b>config&gt;qos&gt;sap-egress</b> <i>policy-id</i> or the <b>config&gt;qos&gt;sap-ingress</b> <i>policy-id</i> context are overridden by parameters specified in the subscriber management SLA profile context.</p> <p>The classification and the queue mapping are shared by all the hosts on the same complex that use the same QoS policy (specified in the <b>sla-profile</b> SAP egress and SAP ingress policy IDs).</p> <p>The queues are shared by all the hosts (of the same subscriber) on the same SAP that are using the same SLA profile. Queues are instantiated when, on a given SAP, a host of a subscriber is the first to use a certain SLA profile. This instantiation is referred to as an SLA profile instance.</p> <p>The <b>no</b> form of the command removes the <i>queue-id</i> from the SLA profile.</p>
<b>Default</b>	none
<b>Parameters</b>	<p><i>queue-id</i> — Specifies the <i>queue-id</i> for the SAP egress or ingress queue, expressed as a decimal integer. The <i>queue-id</i> uniquely identifies the queue within the profile.</p> <p><b>Default</b>      none</p>

## avg-frame-overhead

<b>Syntax</b>	<b>avg-frame-overhead</b> <i>percent</i> <b>no avg-frame-overhead</b>
<b>Context</b>	config>subscr-mgmt>sla-prof>egress>qos>queue
<b>Description</b>	<p>This command configures the average frame overhead to define the average percentage that the offered load to a queue will expand during the frame encapsulation process before sending traffic on-the-wire. While the avg-frame-overhead value may be defined on any queue, it is only used by the system for queues that egress a SONET or SDH port or channel. Queues operating on egress Ethernet ports automatically calculate the frame encapsulation overhead based on a 20 byte per packet rule (8 bytes for preamble and 12 bytes for Inter-Frame Gap).</p> <p>When calculating the frame encapsulation overhead for port scheduling purposes, the system determines the following values:</p>

- **Offered-load** — The offered-load of a queue is calculated by starting with the queue depth in octets, adding the received octets at the queue and subtracting queue discard octets. The result is the number of octets the queue has available to transmit. This is the packet based offered-load.
- **Frame encapsulation overhead** — Using the avg-frame-overhead parameter, the frame encapsulation overhead is simply the queues current offered-load (how much has been received by the queue) multiplied by the avg-frame-overhead. If a queue had an offered load of 10000 octets and the avg-frame-overhead equals 10%, the frame encapsulation overhead would be  $10000 \times 0.1$  or 1000 octets.

For egress Ethernet queues, the frame encapsulation overhead is calculated by multiplying the number of offered-packets for the queue by 20 bytes. If a queue was offered 50 packets then the frame encapsulation overhead would be  $50 \times 20$  or 1000 octets.

- **Frame based offered-load** — The frame based offered-load is calculated by adding the offered-load to the frame encapsulation overhead. If the offered-load is 10000 octets and the encapsulation overhead was 1000 octets, the frame based offered-load would equal 11000 octets.
- **Packet to frame factor** — The packet to frame factor is calculated by dividing the frame encapsulation overhead by the queues offered-load (packet based). If the frame encapsulation overhead is 1000 octets and the offered-load is 10000 octets then the packet to frame factor would be  $1000 / 10000$  or 0.1. When in use, the avg-frame-overhead will be the same as the packet to frame factor making this calculation unnecessary.
- **Frame based CIR** — The frame based CIR is calculated by multiplying the packet to frame factor with the queues configured CIR and then adding that result to that CIR. If the queue CIR is set at 500 octets and the packet to frame factor equals 0.1, the frame based CIR would be  $500 \times 1.1$  or 550 octets.
- **Frame based within-cir offered-load** — The frame based within-cir offered-load is the portion of the frame based offered-load considered to be within the frame-based CIR. The frame based within-cir offered-load is the lesser of the frame based offered-load and the frame based CIR. If the frame based offered-load equaled 11000 octets and the frame based CIR equaled 550 octets, the frame based within-cir offered-load would be limited to 550 octets. If the frame based offered-load equaled 450 octets and the frame based CIR equaled 550 octets, the frame based within-cir offered-load would equal 450 octets (or the entire frame based offered-load).

As a special case, when a queue or associated intermediate scheduler is configured with a CIR-weight equal to 0, the system automatically sets the queue's frame based within-cir offered-load to 0, preventing it from receiving bandwidth during the port scheduler's within-cir pass.

- **Frame based PIR** — The frame based PIR is calculated by multiplying the packet to frame factor with the queue's configured PIR and then adding the result to that PIR. If the queue PIR is set to 7500 octets and the packet to frame factor equals 0.1, the frame based PIR would be  $7500 \times 1.1$  or 8250 octets.
- **Frame based within-pir offered-load** — The frame based within-pir offered-load is the portion of the frame based offered-load considered to be within the frame based PIR. The frame based within-pir offered-load is the lesser of the frame based offered-load and the frame based PIR. If the frame based offered-load equaled 11000 octets and the frame based PIR equaled 8250 octets, the frame based within-pir offered-load would be limited to 8250 octets. If the frame based offered-load equaled 7000 octets and the frame based PIR equaled 8250 octets, the frame based within-pir offered load would equal 7000 octets.

**Port scheduler operation using frame transformed rates** — The port scheduler uses the frame based rates to determine the maximum rates that each queue may receive during the within-cir and above-cir

bandwidth allocation passes. During the within-cir pass, a queue may receive up to its frame based within-cir offered-load. The maximum it may receive during the above-cir pass is the difference between the frame based within-pir offered load and the amount of actual bandwidth allocated during the within-cir pass.

SAP and subscriber SLA-profile average frame overhead override — The average frame overhead parameter on a sap-egress may be overridden at an individual egress queue basis. On each SAP and within the sla-profile policy used by subscribers an avg-frame-overhead command may be defined under the queue-override context for each queue. When overridden, the queue instance will use its local value for the average frame overhead instead of the sap-egress defined overhead.

The **no** form of this command restores the average frame overhead parameter for the queue to the default value of 0 percent. When set to 0, the system uses the packet based queue statistics for calculating port scheduler priority bandwidth allocation. If the no avg-frame-overhead command is executed in a queue-override queue id context, the avg-frame-overhead setting for the queue within the sap-egress QoS policy takes effect.

**Default** 0

**Parameters** *percent* — This parameter sets the average amount of packet-to-frame encapsulation overhead expected for the queue. This value is not used by the system for egress Ethernet queues.

**Values** 0 — 100

burst-limit

**Syntax** **burst-limit {default | size [byte | kilobyte]}**  
**no burst-limit**

**Context** config>subscr-mgmt>sla-prof>egress>qos>queue  
config>subscr-mgmt>sla-prof>ingress>qos>queue

**Description** The `queue burst-limit` command is used to define an explicit shaping burst size for a queue. The configured size defines the shaping leaky bucket threshold level that indicates the maximum burst over the queue’s shaping rate.

The `burst-limit` command is supported under the sap-ingress and sap-egress QoS policy queues. The command is also supported under the ingress and egress queue-group-templates queues.

The **no** form of this command is used to restore the default burst limit to the specified queue. This is equivalent to specifying burst-limit default within the QoS policies or queue group templates. When specified within a queue-override queue context, any current burst limit override for the queue will be removed and the queue’s burst limit will be controlled by its defining policy or template.

**Parameters** **default** — The default parameter is mutually exclusive to specifying an explicit size value. When burst-limit default is executed, the queue is returned to the system default value.

*size* — When a numeric value is specified (size), the system interprets the value as an explicit burst limit size. The value is expressed as an integer and by default is interpreted as the burst limit in Kilobytes. If the value is intended to be interpreted in bytes, the byte qualifier must be added following size.

**Values** 1 to 14,000 (14,000 or 14,000,000 depending on bytes or kilobytes)

**Default** No default for size, use the default keyword to specify default burst limit



**byte** — The **bytes** qualifier is used to specify that the value given for size must be interpreted as the burst limit in bytes. The byte qualifier is optional and mutually exclusive with the kilobytes qualifier.

**kilobyte** — The **kilobyte** qualifier is used to specify that the value given for size must be interpreted as the burst limit in Kilobytes. The kilobyte qualifier is optional and mutually exclusive with the bytes qualifier. If neither bytes nor kilobytes is specified, the default qualifier is kilobytes.

## cbs

<b>Syntax</b>	<b>cbs</b> <i>size-in-kbytes</i> <b>no cbs</b>
<b>Context</b>	config>subscr-mgmt>sla-prof>egress>qos>queue config>subscr-mgmt>sla-prof>ingress>qos>queue
<b>Description</b>	<p>This command can be used to override specific attributes of the specified queue's CBS parameters. It is permissible, and possibly desirable, to oversubscribe the total CBS reserved buffers for a given access port egress buffer pool. Oversubscription may be desirable due to the potential large number of service queues and the economy of statistical multiplexing the individual queues' CBS settings into the defined reserved total.</p> <p>When oversubscribing the reserved total, it is possible for a queue depth to be lower than its CBS setting and still not receive a buffer from the buffer pool for an ingress frame. As more queues are using their CBS buffers and the total in use exceeds the defined reserved total, essentially the buffers are being removed from the shared portion of the pool without the shared in use average and total counts being decremented. This can affect the operation of the high and low priority RED slopes on the pool, causing them to miscalculate when to start randomly drop packets.</p> <p>If the CBS value is larger than the MBS value, an error will occur, preventing the CBS change.</p> <p>The <b>no</b> form of this command returns the CBS size to the size as configured in the QoS policy.</p>
<b>Default</b>	no cbs
<b>Parameters</b>	<p><i>size-in-kbytes</i> — The size parameter is an integer expression of the number of kilobytes reserved for the queue. If a value of 10KBytes is desired, enter the value 10. A value of 0 specifies that no reserved buffers are required by the queue (a minimal reserved size can still be applied for scheduling purposes).</p> <p><b>Values</b>      0 — 131072 or default</p>

## high-prio-only

<b>Syntax</b>	<b>high-prio-only</b> <i>percent</i> <b>no high-prio-only</b>
<b>Context</b>	config>subscr-mgmt>sla-prof>egress>qos>queue config>subscr-mgmt>sla-prof>ingress>qos>queue
<b>Description</b>	This command configures the value of the percentage of buffer space for the queue, used exclusively by high priority packets. The specified value overrides the default value for the context.

The priority of a packet can only be set in the SAP ingress QoS policy and is only applicable on the ingress queues for a SAP. The **high-prio-only** parameter is used to override the default value derived from the **network-queue** command.

The defined **high-prio-only** value cannot be greater than the MBS size of the queue. Attempting to change the MBS to a value smaller than the high priority reserve will generate an error and fail execution. Attempting to set the **high-prio-only** value larger than the current MBS size will also result in an error and fail execution.

The **no** form of this command returns high-prio-only to the size as configured in the QoS policy.

<b>Default</b>	no high-prio-only
<b>Parameters</b>	<i>percent</i> — The <i>percent</i> parameter is the percentage reserved for high priority traffic on the queue. If a value of 10KBytes is desired, enter the value 10. A value of 0 specifies that none of the MBS of the queue will be reserved for high priority traffic. This does not affect RED slope operation for packets attempting to be queued.
<b>Values</b>	0 — 100   default

## mbs

<b>Syntax</b>	<b>mbs</b> <i>size-in-kbytes</i> <b>no mbs</b>
<b>Context</b>	config>subscr-mgmt>sla-prof>egress>qos>queue
<b>Description</b>	<p>This command configures the maximum size for the queue.</p> <p>The sum of the MBS for all queues on an egress access port can oversubscribe the total amount of buffering available. When congestion occurs and buffers become scarce, access to buffers is controlled by the RED slope a packet is associated with. A queue that has not exceeded its MBS size is not guaranteed that a buffer will be available when needed or that the packet's RED slope will not force the discard of the packet. Setting proper CBS parameters and controlling CBS oversubscription is one major safeguard to queue starvation (when a queue does not receive its fair share of buffers). Another is properly setting the RED slope parameters for the needs of services on this port or channel.</p> <p>If the CBS value is larger than the MBS value, an error will occur, preventing the MBS change.</p> <p>The <b>no</b> form of this command returns the MBS size to the size as configured in the QoS policy.</p>
<b>Default</b>	no mbs
<b>Parameters</b>	<i>size-in-kbytes</i> — The size parameter is an integer expression of the maximum number of kilobytes of buffering allowed for the queue. For a value of 100 kbps, enter the value 100. A value of 0 causes the queue to discard all packets.
<b>Values</b>	0 — 131072 or default

## mbs

<b>Syntax</b>	<b>mbs</b> <i>size</i> [bytes   kilobytes] <b>no mbs</b>
---------------	---

<b>Context</b>	config>subscr-mgmt>sla-prof>ingress>qos>queue
<b>Description</b>	<p>The Maximum Burst Size (MBS) command configures the explicit definition of the maximum amount of buffers allowed for a specific queue.</p> <p>The MBS value is used by a queue to determine whether it has exhausted all of its buffers while enqueueing packets. Once the queue has exceeded the amount of buffers allowed by MBS, all packets are discarded until packets have been drained from the queue.</p> <p>The sap-ingress context for mbs provides a mechanism for overriding the default maximum size for the queue.</p> <p>The sum of the MBS for all queues on an ingress access port can oversubscribe the total amount of buffering available. When congestion occurs and buffers become scarce, access to buffers is controlled by the RED slope a packet is associated with. A queue that has not exceeded its MBS size is not guaranteed that a buffer will be available when needed or that the packet's RED slope will not force the discard of the packet. Setting proper CBS parameters and controlling CBS oversubscription is one major safeguard to queue starvation (when a queue does not receive its fair share of buffers). Another is properly setting the RED slope parameters for the needs of services on this port or channel.</p> <p>If the CBS value is larger than the MBS value, an error will occur, preventing the MBS change.</p> <p>The defined high-prio-only value cannot be greater than the MBS size of the queue. Attempting to change the MBS to a value smaller than the high priority reserve will generate an error and fail execution. Attempting to set the high-prio-only value larger than the current MBS size will also result in an error and fail execution.</p> <p>The <b>no</b> form of this command returns the MBS size to the size as configured in the QoS policy.</p>
<b>Default</b>	no mbs
<b>Parameters</b>	<p><i>size</i> [bytes   kilobytes] — The size parameter is an integer expression of the maximum number of bytes or kilobytes of buffering allowed for the queue. For a value of 100 kbps enter the value 100 and specify the <b>kilobytes</b> parameter. A value of 0 causes the queue to discard all packets.</p> <p><b>Values</b>      0 — 131072 or default</p>

## rate

<b>Syntax</b>	<b>rate</b> <i>pir-rate</i> [ <i>cir cir-rate</i> ] <b>no rate</b>
<b>Context</b>	config>subscr-mgmt>sla-prof>egress>qos>queue
<b>Description</b>	<p>This command defines the administrative Peak Information Rate (PIR) and the administrative Committed Information Rate (CIR) parameters for the queue. The PIR defines the maximum rate that the queue can transmit packets out an egress interface (for SAP egress queues). Defining a PIR does not necessarily guarantee that the queue can transmit at the intended rate. The actual rate sustained by the queue can be limited by oversubscription factors or available egress bandwidth.</p> <p>The CIR defines the rate at which the system prioritizes the queue over other queues competing for the same bandwidth. In-profile packets are preferentially queued by the system at egress and at subsequent next hop nodes where the packet can traverse. To be properly handled as in- or out-of-profile throughout the network, the packets must be marked accordingly for profiling at each hop.</p>

The CIR can be used by the queue's parent command's *cir-level* and *cir-weight* parameters to define the amount of bandwidth considered to be committed for the child queue during bandwidth allocation by the parent scheduler.

The **rate** command can be executed at anytime, altering the PIR and CIR rates for all queues created through the association of the SAP egress QoS policy with the *queue-id*.

The **no** form of the command returns all queues created with the *queue-id* by association with the QoS policy to the default PIR and CIR parameters (**max**, 0).

<b>Default</b>	no rate
<b>Parameters</b>	<p><i>pir-rate</i> — Defines the administrative PIR rate, in kilobits, for the queue. When the <b>rate</b> command is executed, a valid PIR setting must be explicitly defined. When the <b>rate</b> command has not been executed, the default PIR of <b>max</b> is assumed.</p> <p>Fractional values are not allowed and must be given as a positive integer.</p> <p>The actual PIR rate is dependent on the queues <b>adaptation-rule</b> parameters and the actual hardware where the queue is provisioned.</p> <p><b>Values</b>      1 — 100000000, max</p> <p><b>Default</b>      max</p> <p><i>cir-rate</i> — The <b>cir</b> parameter overrides the default administrative CIR used by the queue. When the <b>rate</b> command is executed, a CIR setting is optional. When the <b>rate</b> command has not been executed or the <b>cir</b> parameter is not explicitly specified, the default CIR (0) is assumed.</p> <p>Fractional values are not allowed and must be given as a positive integer.</p> <p><b>Values</b>      0 — 100000000, max</p> <p><b>Default</b>      0</p>

## qos-marking-from-sap

<b>Syntax</b>	[no] qos-marking-from-sap
<b>Context</b>	configure>subscr-mgmt>sla-profile>egress
<b>Description</b>	<p>This command sets the QoS policy from which the egress QoS marking rules are applied. Note that if applied to a managed SAP, the default SAP-egress qos-policy (sap-egress 1) cannot be changed.</p> <p>The <b>no</b> form of the command reverts to the egress QoS marking defined in SAP-egress policy defined at sla-profile level.</p>
<b>Default</b>	qos-marking-from-sap

## report-rate

<b>Syntax</b>	<b>report-rate agg-rate-limit</b> <b>report-rate scheduler scheduler-name</b> <b>report-rate pppoe-actual-rate</b> <b>report-rate rfc5515-actual-rate</b> <b>no report-rate</b>
---------------	---

<b>Context</b>	config>subscr-mgmt>sla-prof>ingress config>subscr-mgmt>sla-prof>egress
<b>Description</b>	This command configures the source for Tx and Rx connect speeds in AVP 38 (Rx Connect Speed) and AVP 24 (Tx Connect Speed) of an L2TP session established on a LAC.
<b>Default</b>	no report-rate – Rates takes from the physical port speed.
<b>Parameters</b>	<p><b>agg-rate-limit</b> — (egress only) rate taken from:</p> <ol style="list-style-type: none"> <li>1. The agg-rate RADIUS override (RADIUS VSA “Alc-Subscriber-QoS-Override” in a RADIUS Access-Accept message) if present.</li> <li>2. The configured agg-rate-limit in the <b>config&gt;subscr-mgmt&gt;sub-prof&gt;egr</b> context.</li> <li>3. Fall back to the default (no report-rate).</li> </ol> <p><b>scheduler scheduler-name</b> — Specifies the rate taken from the <b>scheduler scheduler-name</b>. If the <b>scheduler scheduler-name</b> is not present in the scheduler-policy configured in the <b>config&gt;subscr-mgmt&gt;sub-prof&gt;egr</b> context, fall back to the default (no report-rate)</p> <p><b>pppoe-actual-rate</b> — Specifies rates taken from the “DSL Line characteristics” PPPoE tags (Actual Data Rate Upstream/Downstream) if present; otherwise fall back to the default (no report-rate).</p> <p><b>report-rate rfc5515-actual-rate</b> — Puts 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.</p>

## scheduler-policy

<b>Syntax</b>	<b>scheduler-policy scheduler-policy-name</b> <b>no scheduler-policy</b>
<b>Context</b>	config>subscr-mgmt>sla-prof>egress
<b>Description</b>	<p>This command specifies a scheduler policy to associate to the sla profile. Scheduler policies are configured in the <b>configure&gt;qos&gt;scheduler&gt;policy</b> context. Each scheduler policy is divided up into groups of schedulers based on the tier each scheduler is created under. A tier is used to give structure to the schedulers within a policy and define rules for parent scheduler associations. The policy defines the hierarchy and operating parameters for virtual schedulers.</p> <p>The <b>no</b> form of the command removes the scheduler-policy-name from the configuration.</p>
<b>Default</b>	no scheduler-policy
<b>Parameters</b>	<i>scheduler-policy-name</i> — Specify an existing scheduler policy name.

## scheduler

<b>Syntax</b>	<b>scheduler scheduler-name rate pir-rate [cir cir-rate]</b> <b>no scheduler scheduler-name</b>
<b>Context</b>	config>subscr-mgmt>sla-prof>egress>sched

<b>Description</b>	This command provides a way to override parameters of the existing scheduler associated with the egress scheduler policy. A scheduler defines bandwidth controls that limit each child (other schedulers and queues) associated with the scheduler. Scheduler objects are created within the hierarchical tiers of the policy. It is assumed that each scheduler created will have queues or other schedulers defined as child associations. The scheduler can be a child (take bandwidth from a scheduler in a higher tier).
<b>Parameters</b>	<p><b>scheduler</b> <i>scheduler-name</i> — Specify an existing scheduler policy name.</p> <p><b>pir-rate</b> — The pir-rate parameter, in kilobits, overrides the administrative PIR used by the scheduler. When the rate command is executed, a valid PIR setting must be explicitly defined. Fractional values are not allowed and must be given as a positive integer.</p> <p><b>Values</b> 1 — 3200000000, max</p> <p><b>Default</b> none</p> <p><b>cir-rate</b> — The cir parameter, in kilobits, overrides the administrative CIR used by the scheduler. When the rate command is executed, a CIR setting is optional. The sum keyword specifies that the CIR be used as the summed CIR values of the children schedulers or queues. Fractional values are not allowed and must be given as a positive integer.</p> <p><b>Values</b> 0 — 3200000000, sum, max</p> <p><b>Default</b> sum</p>

## use-ingress-l2tp-dscp

<b>Syntax</b>	<b>[no] use-ingress-l2tp-dscp</b>
<b>Context</b>	config>subscr-mgmt>sla-prof>egress
<b>Description</b>	<p>This command enables the use of the DSCP marking taken from the L2TP header received on an L2TP Access Concentrator (LAC) for egress classification for the subscriber host using the associated sla-profile.</p> <p>This command is ignored if the ingress packet is not identified as an L2TP packet.</p>
<b>Default</b>	no use-ingress-l2tp-dscp

## one-time-http-redirection

<b>Syntax</b>	<b>one-time-http-redirection</b> <i>filter-id</i> <b>one-time-http-redirection</b>
<b>Context</b>	config>subscr-mgmt>sla-prof
<b>Description</b>	This command specify the one-time http redirection filter id. This filter will apply to the host when host is created, and will be replaced by the sla-profile ingress filter (configured in the <b>config&gt;subscr-mgmt&gt;sla-prof&gt;ingress</b> context) after first HTTP request from host has been redirected.

**Note:** system does not check if the configured filter include http-redirection entry. If the filter does not include the http-redirection then it will not be replaced in future.

If 7750 receives filter insertion via CoA or access-accept when one-time redirection filter is still active then the received filter entries will only be applied to the sla-profile ingress filter. And after 1st http redirection, the original sla-profile ingress filter + received filter will replace the redirection filter.

<b>Default</b>	no
<b>Parameters</b>	<i>filter-id</i> — Specifies the id of filter that is used for HTTP redirection.

## rate

<b>Syntax</b>	<b>rate</b> <i>pir-rate</i> [ <i>cir</i> <i>cir-rate</i> ] <b>no rate</b>
<b>Context</b>	config>subscr-mgmt>sla-prof>ingress>qos>queue
<b>Description</b>	<p>This command defines the administrative Peak Information Rate (PIR) and the administrative Committed Information Rate (CIR) parameters for the queue. The PIR defines the maximum rate that the queue can transmit packets through the switch fabric (for SAP ingress queues). Defining a PIR does not necessarily guarantee that the queue can transmit at the intended rate. The actual rate sustained by the queue can be limited by oversubscription factors or available egress bandwidth.</p> <p>The CIR defines the rate at which the system prioritizes the queue over other queues competing for the same bandwidth. For SAP ingress, the CIR also defines the rate that packets are considered in-profile by the system. In-profile packets are preferentially queued by the system at egress and at subsequent next hop nodes where the packet can traverse. To be properly handled as in- or out-of-profile throughout the network, the packets must be marked accordingly for profiling at each hop.</p> <p>The CIR can be used by the queue's parent command's <i>cir-level</i> and <i>cir-weight</i> parameters to define the amount of bandwidth considered to be committed for the child queue during bandwidth allocation by the parent scheduler.</p> <p>The <b>rate</b> command can be executed at anytime, altering the PIR and CIR rates for all queues created through the association of the SAP ingress or SAP egress QoS policy with the <i>queue-id</i>.</p> <p>The <b>no</b> form of the command returns all queues created with the <i>queue-id</i> by association with the QoS policy to the default PIR and CIR parameters (<b>max</b>, 0).</p>
<b>Default</b>	no rate
<b>Parameters</b>	<p><i>pir-rate</i> — Defines the administrative PIR rate, in kilobits, for the queue. When the <b>rate</b> command is executed, a valid PIR setting must be explicitly defined. When the <b>rate</b> command has not been executed, the default PIR of <b>max</b> is assumed.</p> <p>Fractional values are not allowed and must be given as a positive integer.</p> <p>The actual PIR rate is dependent on the queues <b>adaptation-rule</b> parameters and the actual hardware where the queue is provisioned.</p> <p><b>Values</b>      1 — 100000000, max</p> <p><b>Default</b>      <b>max</b></p> <p><i>cir-rate</i> — Specifies the <b>cir</b> parameter used by the queue. When the <b>rate</b> command is executed, a CIR setting is optional. When the <b>rate</b> command has not been executed or the <b>cir</b> parameter is not</p>

explicitly specified, the default CIR (0) is assumed.  
Fractional values are not allowed and must be given as a positive integer.

**Values** 0 — 1000000000, **max**

**Default** 0

## policer

<b>Syntax</b>	<b>policer</b> <i>policer-id</i> [ <b>create</b> ] <b>no policer</b> <i>policer-id</i>
<b>Context</b>	config>subscr-mgmt>sla-prof>ingress>qos config>subscr-mgmt>sla-prof>egress>qos config>subscr-mgmt>sub-profile>hsmda>ingress-qos>qos>policer
<b>Description</b>	<p>This command is used in the sap-ingress and sap-egress QoS policies to create, modify or delete a policer. Policers are created and used in a similar manner to queues. The policer ID space is separate from the queue ID space, allowing both a queue and a policer to share the same ID. The sap-ingress policy may have up to 32 policers (numbered 1 through 32) may be defined while the sap-egress QoS policy supports a maximum of 8 (numbered 1 through 8). While a policer may be defined within a QoS policy, it is not actually created on SAPs or subscribers associated with the policy until a forwarding class is mapped to the policer's ID.</p> <p>All policers must be created within the QoS policies. A default policer is not created when a sap-ingress or sap-egress QoS policy is created.</p> <p>Once a policer is created, the policer's metering rate and profiling rates may be defined as well as the policer's maximum and committed burst sizes (MBS and CBS respectively). Unlike queues which have dedicated counters, policers allow various stat-mode settings that define the counters that will be associated with the policer. Another supported feature—packet-byte-offset—provides a policer with the ability to modify the size of each packet based on a defined number of bytes.</p> <p>Once a policer is created, it cannot be deleted from the QoS policy unless any forwarding classes that are mapped to the policer are first moved to other policers or queues.</p> <p>The system will allow a policer to be created on a SAP QoS policy regardless of the ability to support policers on objects where the policy is currently applied. The system only scans the current objects for policer support and sufficient resources to create the policer when a forwarding class is first mapped to the policer ID. If the policer cannot be created due to one or more instances of the policy not supporting policing or having insufficient resources to create the policer, the forwarding class mapping will fail.</p> <p>The <b>no</b> form of this command is used to delete a policer from a sap-ingress or sap-egress QoS policy. The specified policer cannot currently have any forwarding class mappings for the removal of the policer to succeed. It is not necessary to actually delete the policer ID for the policer instances to be removed from SAPs or subscribers associated with the QoS policy once all forwarding classes have been moved away from the policer. It is automatically deleted from each policing instance although it still appears in the QoS policy.</p>
<b>Parameters</b>	<i>policer-id</i> — The <i>policer-id</i> must be specified when executing the policer command. If the specified ID already exists, the system enters that policer's context to allow the policer's parameters to be modified. If the ID does not exist and is within the allowed range for the QoS policy type, a context for the policer ID will be created (depending on the system's current create keyword



requirements which may require the create keyword to actually add the new policer ID to the QoS policy) and the system will enter that new policer's context for possible parameter modification.

**Values** 1—63

## cbs

<b>Syntax</b>	<b>cbs</b> { <i>size</i> [ <b>bytes</b>   <b>kilobytes</b> ]   <b>default</b> } <b>no cbs</b>
<b>Context</b>	config>subscr-mgmt>sla-prof>ingress>qos>policer config>subscr-mgmt>sla-prof>egress>qos>policer config>subscr-mgmt>sub-profile>hsmda>ingress-qos>qos>policer
<b>Description</b>	<p>This command is used to configure the policer's CIR leaky bucket's exceed threshold. The CIR bucket's exceed threshold represents the committed burst tolerance allowed by the policer. If the policer's forwarding rate is equal to or less than the policer's defined CIR, the CIR bucket depth hovers around the 0 depth with spikes up to the maximum packet size in the offered load. If the forwarding rate increases beyond the profiling rate, the amount of data allowed to be in-profile above the rate is capped by the threshold.</p> <p>The policer's <b>cbs</b> size defined in the QoS policy may be overridden on an <b>sla-profile</b> or SAP where the policy is applied.</p> <p>The <b>no</b> form of this command returns the policer to its default CBS size.</p>
<b>Default</b>	<b>none</b>
<b>Parameters</b>	<p><i>size</i> [<b>bytes</b>   <b>kilobytes</b>] — The <i>size</i> parameter is required when specifying <b>cbs</b> and is expressed as an integer representing the required size in either bytes or kilobytes. The default is kilobytes. The optional <b>byte</b> and <b>kilobyte</b> keywords are mutually exclusive and are used to explicitly define whether size represents bytes or kilobytes.</p> <p><b>byte</b> — When <b>byte</b> is defined, the value given for size is interpreted as the queue's MBS value given in bytes.</p> <p><b>kilobyte</b> — When <b>kilobytes</b> is defined, the value is interpreted as the queue's MBS value given in kilobytes.</p> <p><b>Values</b> 0 — 16777216</p> <p><b>Default</b> <b>kilobyte</b></p>

## cbs

<b>Syntax</b>	<b>cbs</b> { <i>size</i> [ <b>bytes</b>   <b>kilobytes</b> ]   <b>default</b> } <b>no cbs</b>
<b>Context</b>	config>subscr-mgmt>sub-profile>hsmda>egress-qos>qos>queue config>subscr-mgmt>sub-profile>hsmda>ingress-qos>qos>queue

<b>Description</b>	<p>This command is used to configure the policer's CIR leaky bucket's exceed threshold. The CIR bucket's exceed threshold represents the committed burst tolerance allowed by the policer. If the policer's forwarding rate is equal to or less than the policer's defined CIR, the CIR bucket depth hovers around the 0 depth with spikes up to the maximum packet size in the offered load. If the forwarding rate increases beyond the profiling rate, the amount of data allowed to be in-profile above the rate is capped by the threshold.</p> <p>The policer's <b>cbs</b> size defined in the QoS policy may be overridden on an <b>sla-profile</b> or SAP where the policy is applied.</p> <p>The <b>no</b> form of this command returns the policer to its default CBS size.</p>
<b>Default</b>	<b>none</b>
<b>Parameters</b>	<p><i>size</i> [<b>bytes</b>   <b>kilobytes</b>] — The <i>size</i> parameter is required when specifying <b>cbs</b> and is expressed as an integer representing the required size in either bytes or kilobytes. The default is kilobytes. The optional <b>byte</b> and <b>kilobyte</b> keywords are mutually exclusive and are used to explicitly define whether size represents bytes or kilobytes.</p> <p><b>byte</b> — When <b>byte</b> is defined, the value given for size is interpreted as the queue's MBS value given in bytes.</p> <p><b>kilobyte</b> — When <b>kilobytes</b> is defined, the value is interpreted as the queue's MBS value given in kilobytes.</p> <p><b>Values</b>      1 — 4194304</p> <p><b>Default</b>      <b>kilobyte</b></p>

## mbs

<b>Syntax</b>	<b>mbs {size [bytes   kilobytes]   default}</b> <b>no mbs</b>
<b>Context</b>	config>subscr-mgmt>sla-prof>ingress>qos>policer config>subscr-mgmt>sla-prof>egress>qos>policer config>subscr-mgmt>sub-profile>hsmda>ingress-qos>qos>policer
<b>Description</b>	<p>This command is used to configure the policer's PIR leaky bucket's high priority violate threshold. The <b>high-prio-only</b> command is applied to the MBS value to derive the bucket's low priority violate threshold. For ingress, trusted in-profile packets and un-trusted high priority packets use the policer's high priority violate threshold while trusted out-of-profile and un-trusted low priority packets use the policer's low priority violate threshold. At egress, in-profile packets use the policer's high priority violate threshold and out-of-profile packets use the policer's low priority violate threshold.</p> <p>The PIR bucket's violate threshold represent the maximum burst tolerance allowed by the policer. If the policer's offered rate is equal to or less than the policer's defined rate, the PIR bucket depth hovers around the 0 depth with spikes up to the maximum packet size in the offered load. If the offered rate increases beyond the metering rate, the amount of data allowed above the rate is capped by the threshold. The low priority violate threshold provides a smaller burst size for the lower priority traffic associated with the policer. Since all lower priority traffic is discarded at the lower burst tolerance size, the remaining burst tolerance defined by <b>high-prio-only</b> is available for the higher priority traffic.</p> <p>The policer's mbs size defined in the QoS policy may be overridden on an sla-profile or SAP where the policy is applied.</p>

The no form of this command returns the policer to its default MBS size.

<b>Default</b>	None
<b>Parameters</b>	<p><i>size</i> [<b>bytes</b>   <b>kilobytes</b>] — The <i>size</i> parameter is required when specifying <b>mbs</b> and is expressed as an integer representing the required size in either bytes or kilobytes. The default is kilobytes. The optional <b>byte</b> and <b>kilobyte</b> keywords are mutually exclusive and are used to explicitly define whether size represents bytes or kilobytes.</p> <p><b>byte</b> — When <b>byte</b> is defined, the value given for size is interpreted as the queue's MBS value given in bytes.</p> <p><b>kilobyte</b> — When <b>kilobytes</b> is defined, the value is interpreted as the queue's MBS value given in kilobytes.</p> <p><b>Values</b>      0 — 16777216</p> <p><b>Default</b>      <b>kilobyte</b></p>

## mbs

<b>Syntax</b>	<b>mbs {size [bytes   kilobytes]   default}</b> <b>no mbs</b>
<b>Context</b>	config>subscr-mgmt>sub-profile>hsmda>egress-qos>qos>queue config>subscr-mgmt>sub-profile>hsmda>ingress-qos>qos>queue
<b>Description</b>	<p>This command is used to configure the policer's PIR leaky bucket's high priority violate threshold. The <b>high-prio-only</b> command is applied to the MBS value to derive the bucket's low priority violate threshold. For ingress, trusted in-profile packets and un-trusted high priority packets use the policer's high priority violate threshold while trusted out-of-profile and un-trusted low priority packets use the policer's low priority violate threshold. At egress, in-profile packets use the policer's high priority violate threshold and out-of-profile packets use the policer's low priority violate threshold.</p> <p>The PIR bucket's violate threshold represent the maximum burst tolerance allowed by the policer. If the policer's offered rate is equal to or less than the policer's defined rate, the PIR bucket depth hovers around the 0 depth with spikes up to the maximum packet size in the offered load. If the offered rate increases beyond the metering rate, the amount of data allowed above the rate is capped by the threshold. The low priority violate threshold provides a smaller burst size for the lower priority traffic associated with the policer. Since all lower priority traffic is discarded at the lower burst tolerance size, the remaining burst tolerance defined by <b>high-prio-only</b> is available for the higher priority traffic.</p> <p>The policer's mbs size defined in the QoS policy may be overridden on an sla-profile or SAP where the policy is applied.</p> <p>The no form of this command returns the policer to its default MBS size.</p>
<b>Default</b>	None
<b>Parameters</b>	<p><i>size</i> [<b>bytes</b>   <b>kilobytes</b>] — The <i>size</i> parameter is required when specifying <b>mbs</b> and is expressed as an integer representing the required size in either bytes or kilobytes. The default is kilobytes. The optional <b>byte</b> and <b>kilobyte</b> keywords are mutually exclusive and are used to explicitly define whether size represents bytes or kilobytes.</p>

**byte** — When **byte** is defined, the value given for size is interpreted as the queue's MBS value given in bytes.

**kilobyte** — When **kilobytes** is defined, the value is interpreted as the queue's MBS value given in kilobytes.

**Values** 1 — 4194304

**Default** kilobyte

## packet-byte-offset

<b>Syntax</b>	<b>packet-byte-offset</b> { <b>add bytes</b>   <b>subtract bytes</b> } <b>no packet-byte-offset</b>
<b>Context</b>	config>subscr-mgmt>sla-prof>ingress>qos>policer config>subscr-mgmt>sla-prof>egress>qos>policer
<b>Description</b>	<p>This command is used to modify the size of each packet handled by the policer by adding or subtracting a number of bytes. The actual packet size is not modified; only the size used to determine the bucket depth impact is changed. The <b>packet-byte-offset</b> command is meant to be an arbitrary mechanism that can be used to either add downstream frame encapsulation or remove portions of packet headers. Both the policing metering and profiling throughput is affected by the offset as well as the stats associated with the policer.</p> <p>When child policers are adding to or subtracting from the size of each packet, the parent policer's <b>min-thresh-separation</b> value should also need to be modified by the same amount.</p> <p>The policer's <b>packet-byte-offset</b> defined in the QoS policy may be overridden on an <b>sla-profile</b> or SAP where the policy is applied.</p> <p>The <b>no</b> version of this command is used to remove per packet size modifications from the policer.</p>
<b>Parameters</b>	<p><b>add bytes</b> — The <b>add</b> keyword is mutually exclusive to the <b>subtract</b> keyword. Either <b>add</b> or <b>subtract</b> must be specified. When <b>add</b> is defined the corresponding bytes parameter specifies the number of bytes that is added to the size each packet associated with the policer for rate metering, profiling and accounting purposes. From the policer's perspective, the maximum packet size is increased by the amount being added to the size of each packet.</p> <p><b>Values</b> 0 — 31</p> <p><b>Default</b> None</p> <p><b>subtract bytes</b> — The <b>subtract</b> keyword is mutually exclusive to the <b>add</b> keyword. Either <b>add</b> or <b>subtract</b> must be specified. When <b>b</b> is defined the corresponding bytes parameter specifies the number of bytes that is subtracted from the size of each packet associated with the policer for rate metering, profiling and accounting purposes. From the policer's perspective, the maximum packet size is reduced by the amount being subtracted from the size of each packet. Note that the minimum resulting packet size used by the system is 1 byte.</p> <p><b>Values</b> ingress 1—32 egress: 1—64</p> <p><b>Default</b> None</p>

## rate

<b>Syntax</b>	<b>rate {max   kilobits-per-second} [cir {max   kilobits-per-second}]</b> <b>no rate</b>
<b>Context</b>	config>subscr-mgmt>sla-prof>ingress>qos>policer config>subscr-mgmt>sla-prof>egress>qos>policer
<b>Description</b>	<p>This command is used to configure the policer's metering and optional profiling rates. The metering rate is used by the system to configure the policer's PIR leaky bucket's decrement rate while the profiling rate configures the policer's CIR leaky bucket's decrement rate. The decrement function empties the bucket while packets applied to the bucket attempt to fill it based on the each packets size. If the bucket fills faster than how much is decremented per packet, the bucket's depth eventually reaches it's exceed (CIR) or violate (PIR) threshold. The <b>cbs</b>, <b>mbs</b>, and <b>high-prio-only</b> commands are used to configure the policer's PIR and CIR thresholds.</p> <p>If a packet arrives at the policer while the bucket's depth is less than the threshold associated with the packet, the packet is considered to be conforming to the bucket's rate. If the bucket depth is equal to or greater than the threshold, the packet is considered to be in the exception state. For the CIR bucket, the exception state is exceeding the CIR rate while the PIR bucket's exception state is violating the PIR bucket rate. If the packet is violating the PIR, the packet is marked red and will be discarded. If the packet is not red, it may be green or yellow based on the conforming or exceeding state from the CIR bucket.</p> <p>When a packet is red neither the PIR or CIR bucket depths are incremented by the packets size. When the packet is yellow the PIR bucket is incremented by the packet size, but the CIR bucket is not. When the packet is green, both the PIR and CIR buckets are incremented by the packet size. This ensures that conforming packets impact the bucket depth while exceeding or violating packets do not.</p> <p>The policer's <b>adaptation-rule</b> command settings are used by the system to convert the specified rates into hardware timers and decrement values for the policer's buckets.</p> <p>By default, the policer's metering rate is <b>max</b> and the profiling rate is 0 Kbps (all packets out-of-profile).</p> <p>The <b>rate</b> settings defined for the policer in the QoS policy may be overridden on an <b>sla-profile</b> or SAP where the policy is applied.</p> <p>The <b>no</b> form of this command is used to restore the default metering and profiling rate to a policer.</p>
<b>Parameters</b>	<p>{<b>max</b>   <i>kilobits-per-second</i>} — Specifying the keyword <b>max</b> or an explicit <i>kilobits-per-second</i> parameter directly following the rate command is required and identifies the policer's metering rate for the PIR leaky bucket. When the policer is first created, the metering rate defaults to max. The <i>kilobits-per-second</i> value must be expressed as an integer and defines the rate in kilobits-per-second. The integer value is multiplied by 1,000 to derive the actual rate in bits-per-second. When max is specified, the maximum policer rate used will be equal to the maximum capacity of the card on which the policer is configured. If the policer rate is set to a value larger than the maximum rate possible for the card, then the PIR used is equivalent to max.</p> <p><b>Values</b>      <b>max</b> or 1—2000000000</p> <p><b>cir</b> {<b>max</b>   <i>kilobits-per-second</i>} — The optional <b>cir</b> keyword is used to override the default CIR rate of the policer. Specifying the keyword max or an explicit <i>kilobits-per-second</i> parameter directly following the cir keyword is required and identifies the policer's profiling rate for the CIR leaky bucket. When the policer is first created, the profiling rate defaults to 0 Kbps. The <i>kilobits-per-second</i> value must be expressed as an integer and defines the rate in kilobits-per-second. The</p>

integer value is multiplied by 1,000 to derive the actual rate in bits-per-second. When max is specified, the maximum policer rate used will be equal to the maximum capacity of the card on which the policer is configured. If the policer rate is set to a value larger than the maximum rate possible for the card, then the CIR used is equivalent to max.

**Values** max or 0—2000000000

## stat-mode

<b>Syntax</b>	<b>stat-mode</b> <i>stat-mode</i> <b>no stat mode</b>
<b>Context</b>	config>subscr-mgmt>sla-prof>ingress>qos>policer config>subscr-mgmt>sla-prof>ingress>qos>queue config>subscr-mgmt>sla-prof>egress>qos>policer config>subscr-mgmt>sla-prof>egress>qos>queue
<b>Description</b>	<p>This command is used to configure the forwarding plane octet and packet counters of a policer or queue to count packets of a specific type or state. For example separate counters for IPv4/IPv6 or separate counters for offered high and low priority policed traffic.</p> <p>For policers, this command overrides the policer stat-mode configuration as defined in the sap-ingress or sap-egress qos policy. For details on sap-ingress and sap-egress policer stat-mode, refer to the 7750 SR OS Quality of Service Guide. For use in Enhanced Subscriber Management (ESM) context only, an additional stat-mode enables separate counters for IPv4 and IPv6 packets.</p> <p>When a policer's stat-mode is changed while the sla profile is in use, any previous counter values are lost and any new counters are set to zero.</p> <p>For queues, this command sets the stat-mode. Queue stat-mode is only available for use in Enhanced Subscriber Management (ESM) context to enable separate IPv4/IPv6 counters.</p> <p>A queue's stat-mode cannot be changed while the SLA profile is in use.</p>
<b>Default</b>	<p>no stat-mode</p> <p>For policers, the default is <b>no stat-mode override</b>. The <b>sap-ingress</b> or <b>sap-egress stat-mode</b> is used instead.</p> <p>For queues, the default is to count in-/out-of-profile octets and packets.</p>
<b>Parameters</b>	<p>For ingress and egress qos queue stat-mode overrides:</p> <p><i>statmode</i> — {v4-v6}</p> <p>For ingress qos policer stat-mode overrides:</p> <p><i>stat-mode</i> — <b>Values</b>no-stats, minimal, offered-profile-no-cir, offered-priority-no-cir, offered-profile-cir, offered-priority-cir, offered-total-cir, offered-limited-profile-cir, offered-profile-capped-cir, offered-limited-capped-cir, v4-v6</p> <p>For egress qos policer stat-mode overrides:</p> <p><i>stat-mode</i> — <b>Values</b>no-stats, minimal, offered-profile-no-cir, offered-profile-cir, offered-total-cir, offered-limited-capped-cir, offered-profile-capped-cir, v4-v6</p>

Refer to the 7750 SR OS Quality of Service Guide for details on the **sap-ingress** and **sap-egress** **policer stat-mode** parameters:

- no-stats
- minimal
- offered-profile-no-cir
- offered-priority-no-cir
- offered-limited-profile-cir
- offered-profile-cir
- offered-priority-cir
- offered-total-cir
- offered-limited-capped-cir
- offered-profile-capped-cir

For use in Enhanced Subscriber Management (ESM) context only:

**v4-v6** — Count IPv4 and IPv6 forwarded/dropped Octets and Packets separately

---

## Subscriber Identification Policy Commands

### sub-ident-policy

<b>Syntax</b>	<b>[no] sub-ident-policy</b> <i>sub-ident-policy-name</i>
<b>Context</b>	config>subscr-mgmt
<b>Description</b>	<p>This command configures a subscriber identification policy. Each subscriber identification policy can have a default subscriber profile defined. The subscriber identification policy default subscriber profile overrides the system default and the subscriber SAP default subscriber profiles. Defining a subscriber identification policy default subscriber profile is optional.</p> <p>The subscriber identification policy default subscriber profile cannot be defined with the subscriber profile name default.</p> <p>Defining a subscriber profile as a subscriber identification policy default subscriber profile will cause all active subscribers currently associated with a subscriber SAP using the policy and associated with a subscriber policy through the system default or subscriber SAP default subscriber profiles to be reassigned to the subscriber policy defined as default on the subscriber identification policy.</p> <p>Attempting to delete a subscriber profile that is currently defined as a default for a subscriber identification policy will fail.</p> <p>When attempting to remove a subscriber identification policy default subscriber profile definition, the system will evaluate each active subscriber on all subscriber SAPs the subscriber identification policy is currently associated with that are using the default definition to determine whether the active subscriber can be either reassigned to a subscriber SAP default or the system default subscriber profile. If all active subscribers cannot be reassigned, the removal attempt will fail.</p>
<b>Parameters</b>	<i>sub-ident-policy-name</i> — Specifies the name of the subscriber identification policy.

### app-profile-map

<b>Syntax</b>	<b>app-profile-map</b>
<b>Context</b>	config>subscr-mgmt>sub-ident-pol
<b>Description</b>	This command enables the context to configure an application profile mapping.

### entry

<b>Syntax</b>	<b>entry key</b> <i>app-profile-string</i> <b>app-profile</b> <i>app-profile-name</i> <b>no entry key</b> <i>app-profile-string</i>
<b>Context</b>	config>subscr-mgmt>sub-ident-pol>app-profile-map
<b>Description</b>	<p>This command configures an application profile string.</p> <p>The <b>no</b> form of the command removes the values from the configuration.</p>



**Parameters**    *app-profile-string* — Specifies the application profile string.  
                       *app-profile-name* — Specifies the application profile name.

## use-direct-map-as-default

**Syntax**        **[no] use-direct-map-as-default**

**Context**        config>subscr-mgmt>sub-ident-pol>app-profile-map  
                       config>subscr-mgmt>sub-ident-pol>sla-profile-map

**Description**    This command enables direct mapping of application profile as default. With this flag, a script returned string will be used as the named profile. If the named profile cannot be found, the default profile will be used.  
                       The **no** form of the command disables the direct mapping.

**Default**        no use-direct-map-as-default

## primary

**Syntax**        **primary**

**Context**        config>subscr-mgmt>sub-ident-pol

**Description**    This command configures a primary identification script.

## script-url

**Syntax**        **script-url** *dhcp-script-url*

**Context**        config>subscr-mgmt>sub-ident-pol>primary  
                       config>subscr-mgmt>sub-ident-pol>secondary  
                       config>subscr-mgmt>sub-ident-pol>tertiary

**Description**    This command specifies the URL of the identification scripts.

**Parameters**    *dhcp-primary-script-url* — Specifies the URL of the primary identification script.  
                       *dhcp-secondary-script-url* — Specifies the URL of the secondary identification script.  
                       *dhcp-tertiary-script-url* — Specifies the URL of the tertiary identification script.

## secondary

**Syntax**        **secondary**

**Context**        config>subscr-mgmt>sub-ident-pol

**Description** This command configures a secondary identification script.

### sla-profile-map

**Syntax** **sla-profile-map**

**Context** config>subscr-mgmt>sub-ident-pol

**Description** This command configures an SLA profile mapping.

### sub-profile-map

**Syntax** **sla-profile-map**

**Context** config>subscr-mgmt>sub-ident-pol

**Description** This command configures a subscriber profile mapping.

### entry

**Syntax** **entry key** *sla-profile-string* **sla-profile** *sla-profile-name*  
**no entry key** *sla-profile-string*

**Context** config>subscr-mgmt>sub-ident-pol>sla-profile-map

**Description** This command configures an SLA profile string. Each subscriber identification string can be provisioned into a subscriber mapping table providing an explicit mapping of the string to a specific subscriber profile. This allows certain subscribers to be directly mapped to the appropriate subscriber profile in the event that the default mappings are not desired for the subscriber.

An explicit mapping of a subscriber identification string to a subscriber profile cannot be defined with the subscriber profile name default. It is possible for the subscriber identification string to be entered in the mapping table without a defined subscriber profile which can result in the explicitly defined subscriber to be associated with the subscriber profile named default.

Explicitly mapping a subscriber identification string to a subscriber profile will cause an existing active subscriber associated with the string to be reassigned to the newly mapped subscriber profile. An explicit mapping overrides all default subscriber profile definitions.

Attempting to delete a subscriber profile that is currently defined as in an explicit subscriber identification string mapping will fail.

The system will fail the removal attempt of an explicit subscriber identification string mapping to a subscriber profile definition when an active subscriber is using the mapping and cannot be reassigned to a defined default non-provisioned subscriber profile.

**Parameters** *sla-profile-string* — Identifies the SLA profile string.

**Values** 16 characters maximum

*sla-profile-name* — Identifies the SLA profile name.

**Values** 32 characters maximum

## entry

**Syntax** **entry key** *sub-profile-string* **sub-profile** *sub-profile-name*  
**no entry key** *sub-profile-string*

**Context** config>subscr-mgmt>sub-ident-pol>sub-profile-map

**Description** This command configures a subscriber profile string.

**Parameters** *sub-profile-string* — Specifies the subscriber profile string.

**Values** 16 characters maximum

*sub-profile-name* — Specifies the subscriber profile name.

**Values** 32 characters maximum

## tertiary

**Syntax** **tertiary**

**Context** config>subscr-mgmt>sub-ident-pol

**Description** This command configures a tertiary identification script.

---

## Auto-Generated Subscriber Identification Key Commands

### auto-sub-id-key

**Syntax**     **auto-sub-id-key**

**Context**    config>subscr-mgmt

### ipoe-sub-id-key

**Syntax**     **ipoe-sub-id-key** *sub-id-key* [*sub-id-key*...(up to 4 max)]  
**no ipoe-sub-id-key**

**Context**    config>subscr-mgmt>>auto-sub-id-key

**Description**    This command enables certain fields to become the base for auto-generation of the default sub-id name. The sub-id name will be auto generated if there is not a more specific method available. Such more specific methods would be a default sub-id name as a sap-id, a preconfigured static string or explicit mappings based on RADIUS/LUDB returned strings.

In case that a more specific sub-id name generation method is not available AND the auto-id keyword is defined under the def-sub-id hierarchy, the sub-id name will be generated by concatenating fields defined in this command separated by a “|” character.

The maximum sub-id name length is 32 characters while the concatenation of subscriber identification fields can easily exceed 32 characters. Subscriber host instantiation will fail in case that the sub-id name is based on subscriber identification fields whose concatenated length exceeds 32 characters. Failing the host creation rather than truncating sub-id name on a 32 character boundary will prevent collision of sub-ids (subscriber name duplication).

In case that a more specific sub-id name generation method is not available AND the auto-id keyword is NOT defined under the def-sub-id hierarchy, the sub-id name will be a random 10 character encoded string based on the fields defined under this command.

There is only one set of identification fields allowed per host type (IPoE or PPP) per chassis.

**Parameters**    *sub-id-key* — Specifies the auto-generated sub-id keys for IPoE hosts.

**Values**        **mac** — The MAC address can be combined with other subscriber host identification fields (circuit-id, remote-id, session-id or sap-id) to form a sub-id name in a user readable format or as a random 10 character encoded value.

In case that the mac address is used as a concatenation field in the sub-id name, then its format becomes a string xx:xx:xx:xx:xx:xx with the length 17B.

The MAC address as the subscriber host identification field is not applicable to PPPoA hosts or static hosts.

**circuit-id** — The circuit-id can be combined with other subscriber host identification fields (mac, remote-id, session-id or sap-id) to form a sub-id name in a user readable format or as a random 10 character encoded value.

In case that the circuit-id is used as a concatenation field in the sub-id name, then its format becomes access-node-id eth slot/port:[vlan-id] or access-node-id atm slot/port:vpi.vci with a variable length.

Note that if circuit-id contains any non printable ASCII characters, the entire circuit-id string will be formatted in hex in the sub-id name output. Otherwise all characters in circuit-id will be converted to ASCII. ASCII printable characters contain bytes in range 0x20..0x7E.

The circuit-id as the subscriber identification field is not applicable to PPPoA hosts, ARP hosts or static hosts.

**remote-id** — The remote-id can be combined with other subscriber host identification fields (mac, circuit-id, session-id or sap-id) to form a sub-id name in a user readable format or as a random 10 character encoded value.

In case that the remote-id is used as a concatenation field in the sub-id name, then its format becomes a remote-id string with a variable length.

Note that if remote-id contains any non printable ASCII characters, the entire remote-id string will be formatted in hex in the sub-id name output. Otherwise all characters in remote-id will be converted to ASCII. ASCII printable characters contain bytes in range 0x20..0x7E.

The remote-id as the subscriber identification field is not applicable to PPPoA hosts, ARP hosts or static hosts.

**sap-id** — The sap-id can be combined with other subscriber host identification fields (mac, circuit-id, remote-id, or session-id) to form a sub-id name in a user readable format or as a random 10 character encoded value.

In case that the circuit-id is used as a concatenation field in the sub-id name, then its format becomes : slot/mda:[outer-vlan].[inner-vlan] with a variable length.

The sap-id as the subscriber identification field is applicable to all hosts types with exception of static hosts.

**Default** ipoe-sub-id-key mac sap-id

## ppp-sub-id-key

<b>Syntax</b>	<b>ppp-sub-id-key</b> <i>sub-id-key</i> [ <i>sub-id-key</i> ...(up to 5 max)] <b>no ppp-sub-id-key</b>
<b>Context</b>	config>subscr-mgmt>>auto-sub-id-key
<b>Description</b>	<p>This command enable certain fields to become the base for auto-generation of default sub-id name. The sub-id name will be auto-generated if there is not a more specific method available. Examples of these specific methods would be a default sub-id name as a sap-id, a preconfigured static string or explicit mappings based on RADIUS/LUDB returned strings.</p> <p>In case that a more specific sub-id name generation method is not available and the <b>auto-id</b> keyword is defined under the def-sub-id hierarchy, the sub-id name will be generated by concatenating fields defined in this command separated by a “ ” character.</p>

The maximum sub-id name length is 32 characters while the concatenation of subscriber identification fields can easily exceed 32 characters. The subscriber host instantiation will fail if the sub-id name is based on subscriber identification fields whose concatenated length exceeds 32 characters. Failing the host creation rather than truncating sub-id name on a 32 character boundary will prevent collision of sub-ids (subscriber name duplication).

In case that a more specific sub-id name generation method is not available and the **auto-id** keyword is not defined under the def-sub-id hierarchy, the sub-id name will be a random 10 character encoded string based on the fields defined under this command.

There is only one set of identification fields allowed per host type (IPoE or PPP) per chassis.

**Parameters**     *sub-id-key* — Specifies the auto-generated sub-id keys for PPP hosts.

**Values**

**mac** — The MAC address can be combined with other subscriber host identification fields (circuit-id, remote-id, session-id or sap-id) to form a sub-id name in a user readable format or as a random 10 character encoded value.

In case that the mac address is used as a concatenation field in the sub-id name, then its format becomes a string xx:xx:xx:xx:xx:xx with the length 17B.

The MAC address as the subscriber host identification field is not applicable to PPPoA hosts or static hosts.

**circuit-id** — The circuit-id can be combined with other subscriber host identification fields (mac, remote-id, session-id or sap-id) to form a sub-id name in a user readable format or as a random 10 character encoded value.

In case that the circuit-id is used as a concatenation field in the sub-id name, then its format becomes access-node-id eth slot/port:[vlan-id] or access-node-id atm slot/port:vpi.vci with a variable length.

Note that if circuit-id contains any non printable ASCII characters, the entire circuit-id string will be formatted in hex in the sub-id name output. Otherwise all characters in circuit-id will be converted to ASCII. ASCII printable characters contain bytes in range 0x20..0x7E.

.The circuit-id as the subscriber identification field is not applicable to PPPoA hosts, ARP hosts or static hosts.

**remote-id** — The remote-id can be combined with other subscriber host identification fields (mac, circuit-id, session-id or sap-id) to form a sub-id name in a user readable format or as a random 10 character encoded value.

In case that the remote-id is used as a concatenation field in the sub-id name, then its format becomes a remote-id string with a variable length.

Please note that if remote-id contains any non printable ASCII characters, the entire remote-id string will be formatted in hex in the sub-id name output. Otherwise all characters in remote-id will be converted to ASCII. ASCII printable characters contain bytes in range 0x20..0x7E.

The remote-id as the subscriber identification field is not applicable to PPPoA hosts, ARP hosts or static hosts.

**sap-id** — The sap-id can be combined with other subscriber host identification fields (mac, circuit-id, remote-id, or session-id) to form a sub-id name in a user readable format or as a random 10 character encoded value.

In case that the circuit-id is used as a concatenation field in the sub-id name, then its format becomes : slot/mda:[outer-vlan].[inner-vlan] with a variable length.

The sap-id as the subscriber identification field is applicable to all hosts types with exception of static hosts.

**session-id** — The session-id can be combined with other subscriber host identification fields (mac, circuit-id, remote-id, or sap-id) to form a sub-id name in a user readable format or as a random 10 character encoded value.

In case that the circuit-id is used as a concatenation field in the sub-id name, then its format becomes a decimal number with variable length.

The session-id as the subscriber identification field is applicable only to PPPoE/PPPoEoA type hosts.

**Default**      ppp-sub-id-key mac sap-id session-id

---

## Subscriber Profile Commands

### sub-profile

<b>Syntax</b>	<b>[no] sub-profile</b> <i>subscriber-profile-name</i>
<b>Context</b>	config>subscr-mgmt
<b>Description</b>	<p>This command enables the context to configure a subscriber profile. A subscriber profile is a template used to define the aggregate QoS for all hosts within a subscriber context. This is done through the definition of the egress and ingress scheduler policies that govern the aggregate SLA for subscribers using the subscriber profile. Subscriber profiles also allow for specific SLA profile definitions when the default definitions from the subscriber identification policy must be overridden.</p> <p>Subscribers are either explicitly mapped to a subscriber profile template or are dynamically associated by one of various non-provisioned subscriber profile definitions.</p> <p>A subscriber host can be associated with a subscriber profile in the following ways, listed from lowest to highest precedence:</p> <ol style="list-style-type: none"><li>1. The subscriber profile named default.</li><li>2. The subscriber profile defined as the subscriber SAP default.</li><li>3. The subscriber profile found by the subscriber identification policy sub-profile-map.</li><li>4. The subscriber profile found by the subscriber identification policy explicit map.</li></ol> <p>In the event that no defaults are defined and the subscriber identification string is not explicitly provisioned to map to a subscriber profile, either the static subscriber host creation will fail or the dynamic subscriber host DHCP ACK will be discarded.</p> <p>Default Subscriber profile:</p> <p>When a subscriber profile is created with the <i>subscriber-profile-name</i> default, it will be used when no other subscriber profile is associated with the subscriber host by the system. Creating a subscriber profile with the <i>subscriber-profile-name</i> default is optional. If a default subscriber profile is not created, all subscriber hosts subscriber identification strings must match either a non-provisioned default or be provisioned as an explicit match to a subscriber profile.</p> <p>The default profile has no effect on existing active subscriber on the system as they exist due to higher precedence mappings.</p> <p>Attempting to delete any subscriber profile (including the profile named default) while in use by existing active subscribers will fail.</p>
<b>Parameters</b>	<i>subscriber-profile-name</i> — Specify the name of the subscriber profile.
<b>Values</b>	32 characters maximum, default



## accounting-policy

<b>Syntax</b>	<b>accounting-policy</b> <i>acct-policy-id</i> <b>no accounting-policy</b>
<b>Context</b>	config>subscr-mgmt>sub-prof
<b>Description</b>	This command specifies the policy to use to collect accounting statistics on this subscriber profile.  A maximum of one accounting policy can be associated with a profile at one time. Accounting policies are configured in the <b>config&gt;log</b> context.  The <b>no</b> form of this command removes the accounting policy association.
<b>Default</b>	no accounting policy
<b>Parameters</b>	<i>acct-policy-id</i> — Enter the accounting <i>policy-id</i> as configured in the <b>config&gt;log&gt;accounting-policy</b> context.
<b>Values</b>	1 — 99

## collect-stats

<b>Syntax</b>	<b>[no] collect-stats</b>
<b>Context</b>	config>subscr-mgmt>sub-prof
<b>Description</b>	When enabled, the agent collects non-RADIUS accounting statistics.  When the <b>no collect-stats</b> command is issued the statistics are still accumulated by the IOM cards. However, the CPU will not obtain the results and write them to the billing file. If a subsequent <b>collect-stats</b> command is issued then the counters written to the billing file include all the traffic while the <b>no collect-stats</b> command was in effect.
<b>Default</b>	collect-stats

## agg-rate-limit

<b>Syntax</b>	<b>agg-rate-limit</b> { <b>max</b>   <i>kilobits-per-second</i> } [ <b>queue-frame-based-accounting</b> ] <b>no agg-rate-limit</b>
<b>Context</b>	config>subscr-mgmt>sub-prof>egress
<b>Description</b>	This command define a subscriber aggregate limit when the subscriber profile is directly associated with an egress port based scheduler instead of a scheduler policy. The optional queue-frame-based-accounting keyword allows the subscriber queues to operate in the frame based accounting mode.  Once egress frame based accounting is enabled on the subscriber profile, all queues associated with the subscriber (created through the sla-profile associated with each subscriber host) will have their rate and CIR values interpreted as frame based values. When shaping, the queues will include the 12 byte Inter-Frame Gap (IFG) and 8 byte preamble for each packet scheduled out the queue. The profiling CIR threshold will also include the 20 byte frame encapsulation overhead. Statistics associated with the queue do not include the frame encapsulation overhead.

The `queue-frame-based-accounting` keyword does not change the behavior of the `egress-agg-rate-limit` rate value. Since `egress-agg-rate-limit` is always associated with egress port based scheduling and egress port based scheduling is dependant on frame based operation, the `egress-agg-rate-limit` rate is always interpreted as a frame based value.

Enabling `queue-frame-based-accounting` will not cause statistics for queues associated with the subscriber to be cleared.

The **no** form of the command removes both an egress aggregate rate limit and egress frame based accounting for all subscribers associated with the sub-profile. If a subscriber's accounting mode is changed, the subscriber's queue statistics are cleared.

**Parameters**     `{max | kilobits-per-second}` — The **max** keyword and *kilobits-per-second* parameter are mutually exclusive. Either **max** or a value for *kilobits-per-second* must follow the `egress-agg-rate-limit` command.

**max** — The **max** keyword specifies that the egress aggregate rate limit for the subscriber is unlimited. Scheduling for the subscriber queues will only be governed by the individual queue parameters and any congestion on the port relative to each queues scheduling priority.

*kilobits-per-second* — The *kilobits-per-second* parameter defines an actual egress aggregate rate to which all queues associated with the sub-profile will be limited. The limit will be managed per subscriber associated with the sub-profile. The value must be defined as an integer and is representative of increments of 1000 bits per second.

**Values**            1 to 40000000

**Default**            max

*queue-frame-based-accounting* — The optional `queue-frame-based-accounting` keyword enables frame based accounting on all queues associated with the subscriber profile. If frame based accounting is required when a subscriber aggregate limit is not necessary, the **max** keyword should precede the `queue-frame-based-accounting` keyword. If frame based accounting must be disabled, execute `egress-agg-rate-limit` without the `queue-frame-based-accounting` keyword present.

**Default**            Frame based accounting is disabled by default

**queue-frame-based-accounting** — Specifies whether to use frame-based accounting when evaluating the aggregation rate limit for the egress queues for this SAP.

avg-frame-size

**Syntax**            **avg-frame-size** *bytes*  
                      **no avg-frame-size**

**Context**            config>subscriber-managemet>sub-profile>egress

**Description**       This command specifies the average frame size used in the calculation of the fixed and variable encapsulation offset when the command `encap-offset` is enabled in the egress context of a subscriber profile.

If the user does not explicitly configure a value for the `avg-frame-size` parameter, then it will also be assumed the offset is zero.

The **no** form of the command removes the `avg-frame-size` parameter from the subscriber profile.

<b>Default</b>	0
<b>Parameters</b>	<i>bytes</i> — specifies the average frame size value to be used in the adjustment of the subscriber aggregate rate to account for the per packet variable expansion of the last mile for the specific session used by the subscriber host.
<b>Values</b>	64 — 4096

## encap-offset

<b>Syntax</b>	<b>encap-offset [type type]</b> <b>no encap-offset</b>
<b>Context</b>	config>subscriber-managemet>sub-profile>egress
<b>Description</b>	<p>This command enables the adjustment of the queue and subscriber aggregate rate based on the last mile Ethernet or ATM encapsulation.</p> <p>In R9.0, the data path computes the adjusted frame size real-time for each serviced packet from a queue by adding the actual packet size to the fixed offset provided by CPM for this queue and variable AAL5 padding.</p> <p>When this command is enabled, the fixed packet offset is derived from the encapsulation type value signaled in the Access-loop-encapsulation sub-TLV in the Vendor-Specific PPPoE Tags or DHCP Relay Options as per RFC 4679. If the user specifies an encapsulation type with the command, this value is used as the default value for all hosts of this subscriber until a host session signaled a valid value. The signaled value is applied to this host only and the remaining hosts of this subscriber continue to use the user entered default type value if configured, or no offset is applied. Note that however, hosts of the same subscriber using the same SLA profile and which are on the same SAP will share the same instance of FC queues. In this case, the last valid encapsulation value signaled by a host of that same instance of the SAP egress QoS policy will override any previous signaled or configured value.</p> <p>If the user manually applied a constant byte offset to each packet serviced by the queue by configuring the packet-byte-offset, it will have no effect on the net offset computed for the packet. This net offset is stored in the subscriber host table.</p> <p>The procedures for handling signaling changes or configuration changes affecting the subscriber profile are as follows:</p> <ol style="list-style-type: none"> <li>1. The avg-frame-size parameter in the subscriber profile is ignored.</li> <li>2. If the user specifies an encapsulation type with the command, this value is used as the default value for all hosts of this subscriber until a host session signaled a valid value. The signaled value is applied to this host and other hosts of the same subscriber sharing the same SLA profile and which are on the same SAP. The remaining hosts of this subscriber continue to use the user entered default type value if configured, or no offset is applied.</li> <li>3. If the user enables/disables the encap-offset option, or changes the parameter value of the encap-offset option, CPM immediately triggers a re-evaluation of subscribers hosts using the corresponding subscriber profile and an update the IOM with the new fixed offset value.</li> <li>4. If a subscriber has a static host or an ARP host, the subscriber host continues to use the user-configured default encapsulation type value or the last valid encapsulation value signaled in the</li> </ol>

PPPoE tags or DHCP relay options by other hosts of the same subscriber which use the same SLA profile instance. If none was signaled or configured, then no rate adjustment is applied.

When the encap-offset option is configured in the subscriber profile, the subscriber host queue rates, that is, CLI and operational PIR and CIR as well as queue bucket updates, the queue statistics, that is, forwarded, dropped, and HQoS offered counters use the last-mile frame-over-the-wire format. The scheduler policy CLI and operational rates also use LM-FoW format. The port scheduler max-rate and the priority level rates and weights, if a Weighted Scheduler Group is used, are always entered in CLI and interpreted as local port frame-over-the-wire rates. The same is true for an agg-rate-limit applied to a vport. Finally the subscriber agg-rate-limit is entered in CLI as last-mile frame-over-the-wire rate. The system maintains a running average frame expansion ratio for each queue to convert queue rates between these two formats.

<b>Parameters</b>	<p><b>type</b> <i>type</i> — The name of the default encapsulation used for all host queues of a subscriber in the absence of a valid value signaled in the PPPoE tags.</p> <p><b>Values</b> pppoa-llc pppoa-null pppoeoa-llc pppoeoa-llc-fcs pppoeoa-llc-tagged pppoeoa-llc-tagged-fcs pppoeoa-null pppoeoa-null-fcs pppoeoa-null-tagged pppoeoa-null-tagged-fcs ipoa-llc ipoa-null ipoeoa-llc ipoeoa-llc-fcs ipoeoa-llc-tagged ipoeoa-llc-tagged-fcs ipoeoa-null ipoeoa-null-fcs ipoeoa-null-tagged ipoeoa-null-tagged-fcs pppoe pppoe-tagged ipoe ipoe-tagged</p>
-------------------	--

## scheduler

<b>Syntax</b>	<p><b>scheduler</b> <i>scheduler-name</i> <b>rate</b> <i>pir-rate</i> [<b>cir</b> <i>cir-rate</i>]</p> <p><b>no scheduler</b> <i>scheduler-name</i></p>
<b>Context</b>	<p>config&gt;subscr-mgmt&gt;sub-prof&gt;egress&gt;sched</p> <p>config&gt;subscr-mgmt&gt;sub-prof&gt;ingress&gt;sched</p>
<b>Description</b>	<p>This command provides a way to override parameters of the existing scheduler associated with the egress or ingress scheduler policy. A scheduler defines bandwidth controls that limit each child (other schedulers and queues) associated with the scheduler. Scheduler objects are created within the hierarchical tiers of the policy. It is assumed that each scheduler created will have queues or other schedulers defined as child associations. The scheduler can be a child (take bandwidth from a scheduler in a higher tier).</p>
<b>Parameters</b>	<p><b>scheduler</b> <i>scheduler-policy-name</i> — Specify an existing scheduler policy name.</p> <p><i>pir-rate</i> — Specify the pir-rate, in kilobits, to override the administrative PIR used by the scheduler. When the <b>rate</b> command is executed, a valid PIR setting must be explicitly defined. Fractional values are not allowed and must be given as a positive integer.</p> <p><b>Values</b> 1 — 3200000000, max</p> <p><b>Default</b> none</p> <p><i>cir-rate</i> — The <b>cir</b> parameter overrides the administrative CIR used by the scheduler. When the <b>rate</b> command is executed, a CIR setting is optional. The sum keyword specifies that the CIR be used as the summed CIR values of the children schedulers or queues. Fractional values are not allowed and must be given as a positive integer.</p> <p><b>Values</b> 0 — 3200000000, <b>sum</b>, <b>max</b></p> <p><b>Default</b> sum</p>

## scheduler-policy

<b>Syntax</b>	<b>scheduler-policy</b> <i>scheduler-policy-name</i> <b>no scheduler-policy</b>
<b>Context</b>	config>subscriber-mgmt>sub-profile>egress config>subscriber-mgmt>sub-profile>ingress
<b>Description</b>	This command specifies a scheduler policy to associate to the subscriber profile. Scheduler policies are configured in the <b>configure&gt;qos&gt;scheduler&gt;policy</b> context. Each scheduler policy is divided up into groups of schedulers based on the tier each scheduler is created under. A tier is used to give structure to the schedulers within a policy and define rules for parent scheduler associations. The policy defines the hierarchy and operating parameters for virtual schedulers.
<b>Parameters</b>	<i>scheduler-policy-name</i> — Specify an existing scheduler policy name.

## lag-per-link-hash

<b>Syntax</b>	<b>lag-per-link-hash class</b> {1   2   3} <b>weight</b> 1..1024 <b>no lag-per-link-hash</b>
<b>Special Cases</b>	config>subscr-mgmt>sub-profile>egress
<b>Description</b>	This command configures weight and class to be used on LAG egress when the LAG uses weighted per-link-hash by subscribers with the profile assigned. Subscribers using profile with lag-per-link-hash default configuration, inherit weight and class from the SAP configuration (1 and 1 respectively if none configured under SAP).  The no form of this command restores default configuration.
<b>Default</b>	no lag-per-link-hash

## policer-control-policy

<b>Syntax</b>	<b>policer-control-policy</b> <i>policy-name</i> [create] <b>no policer-control-policy</b>
<b>Context</b>	config>subscr-mgmt>sub-prof>ingress config>subscr-mgmt>sub-prof>egress
<b>Description</b>	This command is used to create, delete, or modify policer control policies. The <b>policer-control-policy</b> controls the aggregate bandwidth available to a set of child policers. Once created, the policy can be applied to ingress or egress SAPs. The policy can also be applied to the ingress or egress context of a sub-profile.
<b>Default</b>	no policer-control-policy
<b>Parameters</b>	<i>policy-name</i> — Each policer-control-policy must be created with a unique policy name. The name must given as <i>policy-name</i> must adhere to the system policy ASCII naming requirements. If the defined policy-name already exists, the system will enter that policy's context for editing

purposes. If policy-name does not exist, the system will attempt to create a policy with the specified name. Creating a policy may require use of the create parameter when the system is configured for explicit object creation mode.

**Default**      None

**create** — The **create** keyword is required when a new policy is being created and the system is configured for explicit object creation mode.

max-rate

<b>Syntax</b>	<b>max-rate</b> { <i>kilobits-per-second</i>   <b>max</b> } <b>no max-rate</b>
<b>Context</b>	config>subscr-mgmt>sub-prof>ingress>policer-control-policy config>subscr-mgmt>sub-prof>egress>policer-control-policy
<b>Description</b>	<p>The <b>max-rate</b> command defines the parent policer’s PIR leaky bucket’s decrement rate. A parent policer is created for each time the policer-control-policy is applied to either a SAP or subscriber instance. Packets that are not discarded by the child policers associated with the SAP or subscriber instance are evaluated against the parent policer’s PIR leaky bucket.</p> <p>For each packet, the bucket is first decremented by the correct amount based on the decrement rate to derive the current bucket depth. The current depth is then compared to one of two discard thresholds associated with the packet. The first discard threshold (discard-unfair) is applied if the FIR (Fair Information Rate) leaky bucket in the packet’s child policer is in the confirming state. The second discard threshold (discard-all) is applied if the child policer's FIR leaky bucket is in the exceed state. Only one of the two thresholds is applied per packet. If the current depth of the parent policer PIR bucket is less than the threshold value, the parent PIR bucket is in the conform state for that particular packet. If the depth is equal to or greater than the applied threshold, the bucket is in the violate state for the packet.</p> <p>If the result is “conform,” the bucket depth is increased by the size of the packet (plus or minus the per-packet-offset setting in the child policer) and the packet is not discarded by the parent policer. If the result is “violate,” the bucket depth is not increased and the packet is discarded by the parent policer. When the parent policer discards a packet, any bucket depth increases (PIR, CIR and FIR) in the parent policer caused by the packet are canceled. This prevents packets that are discarded by the parent policer from consuming the child policers PIR, CIR and FIR bandwidth.</p> <p>The <b>policer-control-policy root max-rate</b> setting may be overridden on each SAP or sub-profile where the policy is applied.</p>
<b>Default</b>	max
<b>Parameters</b>	<p><i>kilobits-per-second</i> — Defining a kilobits-per-second value is mutually exclusive with the max parameter. The kilobits-per-second value must be defined as an integer that represents the number of kilobytes that the parent policer will be decremented per second. The actual decrement is performed per packet based on the time that has elapsed since the last packet associated with the parent policer.</p> <p><b>Values</b>      Integer 0 – 2000000000</p> <p><i>max</i> — The <b>max</b> parameter is mutually exclusive with defining a <b>kilobits-per-second</b> value. When max is specified, the parent policer does not enforce a maximum rate on the aggregate throughput</p>

of the child policers. This is the default setting when the **policer-control-policy** is first created and is the value that the parent policer returns to when no max-rate is executed. In order for the parent policer to be effective, a kilobits-per-second value should be specified.

*no max-rate* — The **no max-rate** command returns the policer-control-policy's parent policer maximum rate to max.

## priority-mbs-thresholds

<b>Syntax</b>	<b>priority-mbs-thresholds</b>
<b>Context</b>	config>subscr-mgmt>sub-prof>ingress>policer-control-policy config>subscr-mgmt>sub-prof>egress>policer-control-policy
<b>Description</b>	<p>The <b>priority-mbs-thresholds</b> command contains the root arbiter parent policer's <b>min-thresh-separation</b> command and each priority level's <b>mbs-contribution</b> command that is used to internally derive each priority level's shared-portion and fair-portion values. The system uses each priority level's shared-portion and fair-portion value to calculate each priority level's discard-unfair and discard-all MBS thresholds that enforce priority sensitive rate-based discards within the root arbiter's parent policer.</p> <p>The <b>priority-mbs-thresholds</b> CLI node always exists and does not need to be created.</p>
<b>Default</b>	None.

## min-thresh-separation

<b>Syntax</b>	<b>min-thresh-separation</b> <i>size</i> [bytes   kilobytes] <b>no min-thresh-separation</b>
<b>Context</b>	config>subscr-mgmt>sub-prof>ingress>policer-control-policy>priority-mbs-thresholds config>subscr-mgmt>sub-prof>egress>policer-control-policy>priority-mbs-thresholds
<b>Description</b>	<p>The <b>min-thresh-separation</b> command defines the minimum required separation between each in-use discard threshold maintained for each parent policer context associated with the policer-control-policy. The min-thresh-separation value may be overridden on each SAP or sub-profile to which the policy is applied.</p> <p>The system uses the default or specified min-thresh-separation value in order to determine the minimum separation required between each of the of the parent policer discard thresholds. The system enforces the minimum separation based on the following behavior in two ways. The first is determining the size of the shared-portion for each priority level (when the <b>mbs-contribution</b> command's optional fixed keyword is not specified):</p> <ul style="list-style-type: none"> <li>• When a parent policer instance's priority level has less than two child policers associated, the shared-portion for the level will be zero.</li> <li>• When a parent policer instance's priority level has two or more child policers associated, the shared-portion for the level will be equal to the current value of <b>min-thresh-separation</b>.</li> </ul> <p>The second function the system uses the <b>min-thresh-separation</b> value for is determining the value per priority level for the fair-portion:</p>

- When a parent policer instance's priority level has no child policers associated, the fair-portion for the level will be zero.
- When a parent policer instance's priority level has one child policer associated, the fair-portion will be equal to the maximum of the min-thresh-separation value and the priority level's mbs-contribution value.
- When a parent policer instance's priority level has two or more child policers associated, the fair-portion will be equal to the maximum of the following:

- min-thresh-separation** value

- The priority level's **mbs-contribution** value less **min-thresh-separation** value

When the **mbs-contribution** command's optional fixed keyword is defined for a priority level within the policy, the system will treat the defined **mbs-contribution** value as an explicit definition of the priority level's MBS. While the system will continue to track child policer associations with the parent policer priority levels, the association counters will have no effect. Instead the following rules will be used to determine a fixed priority level's shared-portion and fair-portion:

- If a fixed priority level's **mbs-contribution** value is set to zero, both the shared-portion and fair-portion will be set to zero

- If the **mbs-contribution** value is not set to zero:

- The shared-portion will be set to the current **min-thresh-separation** value

- The fair-portion will be set to the maximum of the following:

- min-thresh-separation** value

- mbs-contribution** value less **min-thresh-separation** value

Each time the **min-thresh-separation** value is modified, the thresholds for all instances of the parent policer created through association with this **policer-control-policy** are reevaluated.

Determining the Correct Value for the Minimum Threshold Separation Value

The minimum value for **min-thresh-separation** should be set equal to the maximum size packet that will be handled by the parent policer. This ensures that when a lower priority packet is incrementing the bucket, the size of the increment will not cause the bucket's depth to equal or exceed a higher priority threshold. It also ensures that an unfair packet within a priority level cannot cause the PIR bucket to increment to the discard-all threshold within the priority.

When evaluating maximum packet size, each child policer's per-packet-offset setting should be taken into consideration. If the maximum size packet is 1518 bytes and a per-packet-offset parameter is configured to add 20 bytes per packet, min-thresh-separation should be set to 1538 due to the fact that the parent policer will increment its PIR bucket using the extra 20 bytes.

In most circumstances, a value larger than the maximum packet size is not necessary. Management of priority level aggregate burst tolerance is intended to be implemented using the priority level **mbs-contribution** command. Setting a value larger than the maximum packet size will not adversely affect the policer performance, but it may increase the aggregate burst tolerance for each priority level.

**NOTE:** One thing to note is that a priority level's shared-portion of the parent policer's PIR bucket depth is only necessary to provide some separation between a lower priority's discard-all threshold and this priority's discard-unfair threshold. It is expected that the burst tolerance for the unfair packets is relatively minimal since the child policers feeding the parent policer priority level all have some amount of fair burst before entering into an FIR exceed or unfair state. The fair burst amount for a priority level is defined using the mbs-contribution command.



The **no** form of this command returns the policy's **min-thresh-separation** value to the default value.

<b>Default</b>	<b>no min-thresh-separation</b>
<b>Parameters</b>	<p><i>size</i> [<b>bytes</b>   <b>kilobytes</b>] — The <i>size</i> parameter is required when executing the <b>min-thresh-separation</b> command. It is expressed as an integer and specifies the shared portion in bytes or kilobytes that is selected by the trailing <b>bytes</b> or <b>kilobytes</b> keywords. If both <b>bytes</b> and <b>kilobytes</b> are missing, <b>kilobytes</b> is the assumed value. Setting this value has no effect on parent policer instances where the <b>min-thresh-separation</b> value has been overridden.</p> <p><b>Values</b>      0 – 16777216</p> <p><b>Default</b>      none</p> <p>[<b>bytes</b>   <b>kilobytes</b>] — The <b>bytes</b> keyword is optional and is mutually exclusive with the <b>kilobytes</b> keyword. When specified, <i>size</i> is interpreted as specifying the size of <b>min-thresh-separation</b> in bytes.</p> <p>The <b>kilobytes</b> keyword is optional and is mutually exclusive with the <b>bytes</b> keyword. When specified, <i>size</i> is interpreted as specifying the size of <b>min-thresh-separation</b> in kilobytes.</p> <p><b>Values</b>      <b>bytes</b> or <b>kilobytes</b></p> <p><b>Default</b>      <b>kilobytes</b></p>

## priority

<b>Syntax</b>	<b>priority</b> <i>level</i>
<b>Context</b>	config>subscr-mgmt>sub-prof>ingress>policer-control-policy>priority-mbs-thresholds config>subscr-mgmt>sub-prof>egress>policer-control-policy>priority-mbs-thresholds
<b>Description</b>	<p>The <b>priority</b> level command contains the <b>mbs-contribution</b> configuration command for a given strict priority level. Eight levels are supported numbered 1 through 8 with 8 being the highest strict priority.</p> <p>Each of the eight priority CLI nodes always exists and do not need to be created. While parameters exist for each priority level, the parameters are only applied when the priority level within a parent policer instance is currently supporting child policers.</p>
<b>Default</b>	None.

## mbs-contribution

<b>Syntax</b>	<b>mbs-contribution</b> <i>size</i> [ <b>bytes</b>   <b>kilobytes</b> ] [ <b>fixed</b> ] <b>no mbs-contribution</b>
<b>Context</b>	config>subscr-mgmt>sub-prof>ingress>policer-control-policy>priority-mbs-thresholds>priority config>subscr-mgmt>sub-prof>egress>policer-control-policy>priority-mbs-thresholds>priority

**Description** The **mbs-contribution** command is used to configure the policy-based burst tolerance for a parent policer instance created when the policy is applied to a SAP or subscriber context. The system uses the parent policer's **min-thresh-separation** value, the priority level's **mbs-contribution** value and the number of child policers currently attached to the priority level to derive the priority level's shared-portion and fair-portion of burst tolerance within the local priority level. The shared-portion and fair-portions for each priority level are then used by the system to calculate each priority level's discard-unfair threshold and discard-all threshold.

The value for a priority level's **mbs-contribution** within the policer-control-policy may be overridden on the SAP or subscriber sub-profile where the policy is applied in order to allow fine tuning of the discard-unfair and discard-all thresholds relevant to the needs of the local child policers on the object.

### Accumulative Nature of Burst Tolerance for a Parent Policer Priority Level

When defining **mbs-contribution**, the specified size may only be a portion of the burst tolerance associated with the priority level. The packets associated with the priority level share the burst tolerance of lower within the parent policer. As the parent policer PIR bucket depth increases during congestion, the lower priority packets eventually experience discard based on each priority's discard-unfair and discard-all thresholds. Assuming congestion continues once all the lower priority packets have been prevented from consuming bucket depth, the burst tolerance for the priority level will be consumed by its own packets and any packets associated with higher priorities.

### The Effect of Fair and Unfair Child Policer Traffic at a Parent Policer Priority Level

The system continually monitors the offered rate of each child policer on each parent policer priority level and detects when the policer is in a congested state (the aggregate offered load is greater than the decrement rate defined on the parent policer). As previously stated, the result of congestion is that the parent policer's bucket depth will increase until it eventually hovers around either a discard-unfair or discard-all threshold belonging to one of the priority levels. This threshold is the point where enough packets are being discarded that the increment rate and decrement rate begin to even out. If only a single child policer is associated to the priority level, the discard-unfair threshold is not used since fairness is only applicable when multiple child policers are competing at the same priority level.

When multiple child policers are sharing the congested priority level, the system uses the offered rates and the parenting parameters of each child to determine the fair rate per child when the parent policer is unable to meet the bandwidth needs of each child. The fair rate represents the amount of bandwidth that each child at the priority level should receive relative to the other children at the same level according to the policer control policy instance managing the child policers. This fair rate is applied as the decrement rate for each child's FIR bucket. Changing a child's FIR rate does not modify the amount of packets forwarded by the parent policer for the child's priority level. It simply modifies the forwarded ratio between the children on that priority level. Since each child FIR bucket has some level of burst tolerance before marking its packets as unfair, the current parent policer bucket depth may at times rise above the discard-unfair threshold. The mbs-contribution value provides a means to define how much separation is provided between the priority level's discard-unfair and discard-all threshold to allow the parent policer to absorb some amount of FIR burst before reaching the priority's discard-all threshold.

This level of fair aggregate burst tolerance is based on the decrement rate of the parent policer's PIR bucket while the individual fair bursts making up the aggregate are based on each child's FIR decrement rate. The aggregate fair rate of the priority level is managed by the system with consideration of the current rate of traffic in higher priority levels. In essence, the system ensures that for each iteration of the child FIR rate calculation, the sum of the child FIR decrement rates plus the sum of the higher priority traffic increment rates equals the parent policers decrement rate. This means that dynamic

amounts of higher priority traffic can be ignored when sizing a lower priority's fair aggregate burst tolerance. Consider the following:

- The parent policer decrement rate is set to 20 Mbps (max-rate 20,000).
- A priority level's fair burst size is set to 30 Kbytes (mbs-contribution 30 kilobytes).
- Higher priority traffic is currently taking 12 Mbps.
- The priority level has three child policers attached.
- Each child's PIR MBS is set to 10 Kbytes, which makes each child's FIR MBS 10 Kbytes.
- The children want 10 Mbps, but only 8 Mbps is available,
- Based on weights, the children's FIR rates are set as follows:

	FIR Rate	FIR MBS
Child 1	4 Mbps	10 Kbytes
Child 2	3 Mbps	10 Kbytes
Child 3	1 Mbps	10 Kbytes

The 12 Mbps of the higher priority traffic and the 8 Mbps of fair traffic equal the 20 Mbps decrement rate of the parent policer.

It is clear that the higher priority traffic is consuming 12 Mbps of the parent policer's decrement rate, leaving 8 Mbps of decrement rate for the lower priority's fair traffic.

- The burst tolerance of child 1 is based on 10 Kbytes above 4 Mbps,
- The burst tolerance of child 2 is based on 10 Kbytes above 3 Mbps,
- The burst tolerance of child 3 is based on 10 Kbytes above 1 Mbps.

If all three children burst simultaneously (unlikely), they will consume 30 Kbytes above 8 Mbps. This is the same as the remaining decrement rate after the higher priority traffic.

#### Parent Policer Total Burst Tolerance and Downstream Buffering

The highest in-use priority level's discard-all threshold is the total burst tolerance of the parent policer. In some cases the parent policer represents downstream bandwidth capacity and the max-rate of the parent policer is set to prevent overrunning the downstream bandwidth. The burst tolerance of the parent policer defines how much more traffic may be sent beyond the downstream scheduling capacity. In the worst case scenario, when the downstream buffering is insufficient to handle the total possible burst from the parent policer, downstream discards based on lack of buffering may occur. However, in all likelihood, this is not the case.

In most cases, lower priority traffic in the policer will be responsible for the greater part of congestion above the parent policer rate. Since this traffic is discarded with a lower threshold, this lowers the effective burst tolerance even while the highest priority traffic is present.

#### Configuring a Priority Level's MBS Contribution Value

In the most conservative case, a priority level's **mbs-contribution** value may be set to be greater than the sum of child policer's mbs and one max-size-frame per child policer. This ensures that even in the absolute worst case where all the lower priority levels are simultaneously bursting to the maximum capacity of each child, enough burst tolerance for the priority's children will exist if they also burst to their maximum capacity.

Since simply adding up all the child policer's PIR MBS values may result in large overall burst tolerances that are not ever likely to be needed, you should consider some level of burst oversubscription when configuring the **mbs-contribution** value for each priority level. The amount of oversubscription should be determined based on the needs of each priority level.

### Using the Fixed Keyword to Create Deterministic Parent Policer Discard Thresholds

In the default behavior, the system ignores the **mbs-contribution** values for a priority level on a subscriber or SAP parent policer when a child policer is not currently associated with the level. This prevents additional burst tolerance from being added to higher priority traffic within the parent policer.

This does cause fluctuations in the defined threshold values when child policers are added or removed from a parent policer instance. If this behavior is undesirable, the fixed keyword may be used which causes the **mbs-contribution** value to always be included in the calculation of parent policer's discard thresholds. The defined **mbs-contribution** value may be overridden on a subscriber sla-profile or on a SAP instance, but the fixed nature of the contribution cannot be overridden.

If the defined **mbs-contribution** value for the priority level is zero, the priority level will have no effect on the parent policer's defined discard thresholds. A packet associated with the priority level will use the next lower priority level's discard-unfair and discard-all thresholds.

## Parameters

*size* [**bytes** | **kilobytes**] — The size parameter is required when executing the **mbs-contribution** command. It is expressed as an integer and specifies the priority's specific portion amount of accumulative MBS for the priority level in bytes or kilobytes which is selected by the trailing **bytes** or **kilobytes** keywords. If both **bytes** and **kilobytes** are missing, **kilobytes** is assumed. Setting this value has no effect on parent policer instances where the priority level's **mbs-contribution** value has been overridden.

**Values** 0 — 16777216

**Default** none

**bytes | kilobytes:** — The **bytes** keyword is optional and is mutually exclusive with the **kilobytes** keyword. When specified, size is interpreted as specifying the size of **min-thresh-separation** in bytes.

The **kilobytes** keyword is optional and is mutually exclusive with the **bytes** keyword. When specified, size is interpreted as specifying the size of min-thresh-separation in kilobytes.

**Default**      **kilobytes**

**fixed** — The optional **fixed** keyword is used to force the inclusion of the defined **mbs-contribution** value in the parent policer's discard threshold calculations. If the **mbs-contribution** command is executed without the **fixed** keyword, the fixed calculation behavior for the priority level is removed.

**Default**      **no mbs-contribution**

The **no mbs-contribution** command returns the policy's priority level's MBS contribution to the default value. When changed, the thresholds for the priority level and all higher priority levels for all instances of the parent policer will be recalculated.

## radius-accounting-policy

**Syntax** `radius-accounting-policy acct-policy-name [duplicate acct-policy-name]`  
`no radius-accounting-policy`

**Context** config>subscr-mgmt>sub-prof

<b>Description</b>	This command specifies an existing RADIUS accounting policy to use to collect accounting statistics on this subscriber profile by RADIUS. This command is used independently of the <b>collect-stats</b> command.
<b>Parameters</b>	<p><i>acct-policy-name</i> — Specifies an existing RADIUS based accounting policy.</p> <p><b>duplicate</b> <i>acct-policy-name</i> — Specifies the RADIUS accounting policy to be used to generate duplicate accounting information.</p>

## sla-profile-map

<b>Syntax</b>	<b>sla-profile-map</b>
<b>Context</b>	config>subscr-mgmt>sub-prof
<b>Description</b>	This command enables the context to configure SLA profile mapping.

## entry

<b>Syntax</b>	<b>entry key</b> <i>sub-profile-string</i> <b>sub-profile</b> <i>sub-profile-name</i> <b>no entry key</b> <i>sub-profile-string</i>
<b>Context</b>	config>subscr-mgmt>sub-prof>sla-prof-map
<b>Description</b>	This command configures SLA profile string mappings.
<b>Parameters</b>	<p><i>sub-profile-string</i> — Specifies the subscriber profile string.</p> <p><b>Values</b> 16 characters maximum</p> <p><i>sub-profile-name</i> — Specifies the subscriber profile name.</p> <p><b>Values</b> 32 characters maximum</p>

## use-direct-map-as-default

<b>Syntax</b>	<b>[no] use-direct-map-as-default</b>
<b>Context</b>	config>subscr-mgmt>sub-prof>sla-prof-map
<b>Description</b>	<p>This command enables direct mapping of the SLA profile as default.</p> <p>The <b>no</b> form of the command disables direct mapping,</p>

## sub-mcac-policy

<b>Syntax</b>	<b>sub-mcac-policy</b> <i>policy-name</i> <b>no sub-mcac-policy</b>
<b>Context</b>	config>subscr-mgmt>sub-prof

<b>Description</b>	<p>This command references the policy template in which the mcac bandwidth limits are defined. Mcac for the subscriber is effectively enabled with this command when the sub-profile is applied to the subscriber. The bandwidth of the channels is defined in a different policy (under the <b>configure&gt;router&gt;mcac</b> context) and this policy is applied on the interface level as follows:</p> <ul style="list-style-type: none"> <li>• For group-interfaces under the <b>configure&gt;service&gt;vrf&gt;igmp&gt;group-interface&gt;mcac</b> context</li> <li>• For regular interfaces under the <b>configure&gt;service/router&gt;igmp&gt;interface&gt;mcac</b> context</li> </ul> <p>In case of HQoS Adjustment, it is mandatory that the sub-mcac-policy be created and applied to the subscriber. The sub-mac-policy does not have to contain any bandwidth constraints, but it has to be in a no shutdown state in order for HQoS Adjustment to work.</p>
<b>Default</b>	none
<b>Parameters</b>	<i>policy-name</i> — Specifies the policy name configured in the config>subscr-mgmt>sub-mcac-policy context.

## volume-stats-type

<b>Syntax</b>	<b>volume-stats-type {ip default}</b> <b>no volume-stats-type</b>
<b>Context</b>	config>subscr-mgmt>sub-prof
<b>Description</b>	<p>This command enables the reporting of layer 3 (IP) based subscriber host volume accounting data.</p> <p>By default, subscriber host volume accounting data includes Layer 2 header octets and can be configured to include a fixed packet byte offset or last-mile encapsulation overhead.</p>
<b>Default</b>	<b>volume-stats-type default</b>
<b>Parameters</b>	<p><b>default</b> — subscriber host volume accounting data is reported including the Layer 2 header octets and optional delta's introduced by configuration (for example: packet byte offset, last mile aware shaping, etc.)</p> <p><b>ip</b> — subscriber host volume accounting data reporting is based on Layer 3 (IP) packet sizes. This includes subscriber host ingress/egress queue and policer stats in snmp, CLI show commands, RADIUS and XML accounting, and Diameter Gx usage monitoring. RADIUS and Diameter (DCCA) based credit control volume quota are interpreted as Layer 3 (IP).</p>

## hsmdda

<b>Syntax</b>	<b>hsmdda</b>
<b>Context</b>	config>subscr-mgmt>sub-prof
<b>Description</b>	This command enables the context to configure egress and ingress HSMDDA queue parameters.

## egress-qos

<b>Syntax</b>	<b>egress-queues</b>
<b>Context</b>	config>subscr-mgmt>sub-prof>hsmda
<b>Description</b>	This command enables the context to configure SAP egress QOS policy for the HSMDA egress queue.

## ingress-qos

<b>Syntax</b>	<b>ingress-queues</b>
<b>Context</b>	config>subscr-mgmt>sub-prof>hsmda>egress-queues
<b>Description</b>	This command enables the context to configure SAP egress QOS policy for the HSMDA ingress queue

## agg-rate

<b>Syntax</b>	<b>agg-rate</b> <i>rate</i> <b>no agg-rate</b>
<b>Context</b>	config>port>sonet-sdh>path>access>egress>vport config>port>ethernet>access>egress>vport
<b>Description</b>	This command configures an aggregate rate for the vport. The <b>agg-rate</b> <i>rate</i> , <b>port-scheduler-policy</b> and <b>scheduler-policy</b> commands are mutually exclusive. Changing between the use of a scheduler policy and the use of an agg-rate/port-scheduler-policy involves removing the existing command and applying the new command. Applying a <b>scheduler-policy</b> to a VPORT is only applicable to Ethernet interfaces.
<b>Parameters</b>	<i>rate</i> — Specifies the rate limit for the vport.
<b>Values</b>	1 — 800000000, max

## limit-unused-bandwidth

<b>Syntax</b>	<b>limit-unused-bandwidth</b>
<b>Context</b>	config>port>sonet-sdh>path>access>egress>vport config>port>ethernet>access>egress>vport
<b>Description</b>	Optional command used to enable (or disable) aggregate rate overrun protection on the agg-rate context.

## agg-rate-limit

<b>Syntax</b>	<b>agg-rate-limit</b> <i>agg-rate</i> <b>no agg-rate-limit</b>
---------------	---

<b>Context</b>	config>subscr-mgmt>sub-prof>hsmda>egress-qos
<b>Description</b>	<p>This command defines a maximum total rate for all subscriber egress queues for each subscriber associated with the sub-profile. The egress-agg-rate-limit command is mutually exclusive with the egress-scheduler-policy. When an egress-scheduler-policy is defined on the sub-profile, the egress-agg-rate-limit command will fail. If the egress-agg-rate-limit command is specified, at attempt to bind an egress-scheduler-policy to the sub-profile will fail.</p> <p>A port scheduler policy must be applied on the egress port or channel the subscriber instance is bound to in order for the defined egress-agg-rate-limit to take effect. The egress port scheduler enforces the aggregate queue rate as it distributes its bandwidth at the various port priority levels. The port scheduler stops offering bandwidth to member queues once it has detected that the aggregate rate limit has been reached.</p> <p>If a port scheduler is not defined on the egress port, the queues are allowed to operate based on their own bandwidth parameters.</p> <p>The <b>no</b> form of the command removes the aggregate rate limit from the sub-profile.</p>
<b>Default</b>	no agg-rate-limit
<b>Parameters</b>	<p><i>agg-rate</i> — Defines the maximum aggregate rate the egress queues associated with the subscriber profile may operate. The value is specified in kilobits per second in a base 10 context. A value of 1 indicates a rate of 1000 bits per second.</p> <p><b>Values</b> 1 — 40000000, max Kbps</p>

### qos

<b>Syntax</b>	<p><b>qos</b> <i>policy-id</i></p> <p><b>no qos</b></p>
<b>Context</b>	config>subscr-mgmt>sub-prof>hsmda>egress-qos
<b>Description</b>	This command assigns a SAP egress QoS policy to the HSMDA egress queue.
<b>Parameters</b>	<p><i>policy-id</i> — Specifies the policy ID of an existing QoS SAP egress policy.</p> <p><b>Values</b> 1 — 65535</p>

### qos

<b>Syntax</b>	<p><b>qos</b> <i>policy-id</i></p> <p><b>no qos</b></p>
<b>Context</b>	config>subscr-mgmt>sub-prof>hsmda>ingress-qos
<b>Description</b>	This command assigns a SAP ingress QoS policy to the HSMDA ingress queue.
<b>Parameters</b>	<p><i>policy-id</i> — Specifies the policy ID of an existing QoS SAP egress policy.</p> <p><b>Values</b> 1 — 65535</p>



## packet-byte-offset

<b>Syntax</b>	<b>packet-byte-offset</b> { <b>add</b> <i>add-bytes</i>   <b>subtract</b> <i>sub-bytes</i> } <b>no packet-byte-offset</b>
<b>Context</b>	config>subscr-mgmt>sub-prof>hsmda>egress-qos
<b>Description</b>	<p>This command adds or subtracts the specified number of bytes to the accounting function for each packet handled by the HSMDA queue. Normally, the accounting and leaky bucket functions are based on the Ethernet DLC header, payload and the 4 byte CRC (everything except the preamble and inter-frame gap). As an example, the packet-byte-offset command can be used to add the frame encapsulation overhead (20 bytes) to the queues accounting functions.</p> <p>The accounting functions affected include:</p> <ul style="list-style-type: none"> <li>• Offered High Priority / In-Profile Octet Counter</li> <li>• Offered Low Priority / Out-of-Profile Octet Counter</li> <li>• Discarded High Priority / In-Profile Octet Counter</li> <li>• Discarded Low Priority / Out-of-Profile Octet Counter</li> <li>• Forwarded In-Profile Octet Counter</li> <li>• Forwarded Out-of-Profile Octet Counter</li> <li>• Peak Information Rate (PIR) Leaky Bucket Updates</li> <li>• Committed Information Rate (CIR) Leaky Bucket Updates</li> <li>• Queue Group Aggregate Rate Limit Leaky Bucket Updates</li> </ul> <p>The secondary shaper leaky bucket, scheduler priority level leaky bucket and the port maximum rate updates are not affected by the configured packet-byte-offset. Each of these accounting functions are frame based and always include the preamble, DLC header, payload and the CRC regardless of the configured byte offset.</p> <p>The packet-byte-offset command accepts either add or subtract as valid keywords which define whether bytes are being added or removed from each packet traversing the queue. An example use case for subtracting bytes from each packet is an IP based accounting function. Given a Dot1Q encapsulation, the command packet-byte-offset subtract 14 would remove the DLC header and the Dot1Q header from the size of each packet for accounting functions only. The 14 bytes are not actually removed from the packet, only the accounting size of the packet is affected.</p> <p>As inferred above, the variable accounting size offered by the packet-byte-offset command is targeted at the queue and queue group level. The packet-byte-offset, when set, applies to all queues in the queue group. The accounting size of the packet is ignored by the secondary shapers, the scheduling priority level shapers and the scheduler maximum rate. The actual on-the-wire frame size is used for these functions to allow an accurate representation of the behavior of the subscriber's packets on an Ethernet aggregation network.</p> <p>The packet-byte-offset value may be overridden at the queue-group level.</p>
<b>Parameters</b>	<p><b>add</b> <i>add-bytes</i> — Indicates that the byte value should be added to the packet for queue and queue group level accounting functions. Either the <b>add</b> or <b>subtract</b> keyword must be specified. The corresponding byte value must be specified when executing the packet-byte-offset command. The <b>add</b> keyword is mutually exclusive with the <b>subtract</b> keyword.</p> <p><b>Values</b>      0 — 31</p>

**subtract** *sub-bytes* — Indicates that the byte value should be subtracted from the packet for queue and queue group level accounting functions. The **subtract** keyword is mutually exclusive with the **add** keyword. Either the **add** or **subtract** keyword must be specified. The corresponding byte value must be specified when executing the packet-byte-offset command. Note that the minimum resulting packet size used by the system is 1 byte.

**Values** 1 — 64

## queue

<b>Syntax</b>	<b>queue</b> <i>queue-id</i> [ <b>create</b> ] <b>no queue</b> <i>queue-id</i>
<b>Context</b>	config>subscr-mgmt>sub-prof>hsmda>ingress-qos>qos
<b>Description</b>	<p>This command specifies the HSM DA queue mapping for all packets in point-to-point services and unicast destined packets in multipoint services. Point-to-point services include epipe and other VLL type services. Multipoint services include IES, VPLS and VPRN services. The queue command does not apply to multicast, broadcast or unknown unicast packets within multipoint services (the multi-cast, broadcast and unknown commands must be used to define the queue mapping for non-unicast packets within a forwarding class). For Epipe services, the <b>queue</b> <i>queue-id</i> mapping applies to all packets, regardless of the packets destination MAC address.</p> <p>Each forwarding class has a default queue ID based on the intrinsic hierarchy between the forwarding classes. Executing the queue command within the HSM DA context of a forwarding class with a different queue ID than the default overrides the default mapping. Multiple forwarding classes may be mapped to the same HSM DA queue ID.</p> <p>The <b>no</b> form of the command returns the HSM DA queue mapping for queue to the default mapping for the forwarding class.</p>
<b>Parameters</b>	<p><i>queue-id</i> — Specifies the queue ID to override.</p> <p><b>Values</b> 1 — 8</p> <p><b>create</b> — This keyword is mandatory while creating a new queue override.</p>

## rate

<b>Syntax</b>	<b>rate</b> <i>pir-rate</i> [ <b>cir</b> <i>cir-rate</i> ] <b>no rate</b>
<b>Context</b>	config>subscr-mgmt>sub-prof>hsmda>egress-qos>qos>queue config>subscr-mgmt>sub-prof>hsmda>ingress-qos>queue config>subscr-mgmt>sub-prof>hsmda>ingress-qos>policer
<b>Description</b>	<p>This command can be used to override specific attributes of the specified queue's Peak Information Rate (PIR) and the Committed Information Rate (CIR) parameters. The PIR defines the maximum rate that the queue can transmit packets out an egress interface (for SAP egress queues). Defining a PIR does not necessarily guarantee that the queue can transmit at the intended rate. The actual rate sustained by the queue can be limited by oversubscription factors or available egress bandwidth.</p>

The CIR defines the rate at which the system prioritizes the queue over other queues competing for the same bandwidth. In-profile packets are preferentially queued by the system at egress and at subsequent next hop nodes where the packet can traverse. To be properly handled as in- or out-of-profile throughout the network, the packets must be marked accordingly for profiling at each hop.

The CIR can be used by the queue's parent commands *cir-level* and *cir-weight* parameters to define the amount of bandwidth considered to be committed for the child queue during bandwidth allocation by the parent scheduler.

The **rate** command can be executed at any time, altering the PIR and CIR rates for all queues created through the association of the SAP egress QoS policy with the *queue-id*.

The **no** form of the command returns all queues created with the *queue-id* by association with the QoS policy to the default PIR and CIR parameters (**max**, 0).

<b>Default</b>	<b>rate max cir 0</b> — The <b>max</b> default specifies the amount of bandwidth in kilobits per second (thousand bits per second). The <b>max</b> value is mutually exclusive to the <b>pir-rate</b> value.
<b>Parameters</b>	<p><i>pir-rate</i> — Defines the administrative PIR rate, in kilobits, for the queue. When the <b>rate</b> command is executed, a valid PIR setting must be explicitly defined. When the <b>rate</b> command has not been executed, the default PIR of <b>max</b> is assumed.</p> <p>Fractional values are not allowed and must be given as a positive integer.</p> <p>The actual PIR rate is dependent on the queue's <b>adaptation-rule</b> parameters and the actual hardware where the queue is provisioned.</p> <p><b>Values</b>      1 — 100000000</p> <p><b>Default</b>      max</p> <p><i>cir-rate</i> — The <b>cir</b> parameter overrides the default administrative CIR used by the queue. When the <b>rate</b> command is executed, a CIR setting is optional. When the <b>rate</b> command has not been executed or the <b>cir</b> parameter is not explicitly specified, the default CIR (0) is assumed.</p> <p>Fractional values are not allowed and must be given as a positive integer. The <b>sum</b> keyword specifies that the CIR be used as the summed CIR values of the children schedulers or queues.</p> <p><b>Values</b>      0 — 100000000, <b>max</b>, <b>sum</b></p> <p><b>Default</b>      0</p>

## slope-policy

<b>Syntax</b>	<b>slope-policy</b> <i>hsmda-slope-policy-name</i> <b>no slope-policy</b>
<b>Context</b>	config>subscr-mgmt>sub-prof>hsmda>egress-qos>qos>queue
<b>Description</b>	<p>This command specifies an existing slope policy name. The policy contains the Maximum Buffer Size (MBS) that will be applied to the queue and the high and low priority RED slope definitions. The function of the MBS and RED slopes is to provide congestion control for an HSMDBA queue. The MBS parameter defines the maximum depth a queue may reach when accepting packets. The low and high priority RED slopes provides for random early detection of congestion and slope based discards based on queue depth.</p> <p>An hsmda-slope-policy can be applied to queues defined in the sap-ingress and sap-egress QoS policy hsmda-queues context. Once an HSMDBA slope policy is applied to a SAP QoS policy queue, it can-</p>

not be deleted. Any edits to the policy are updated to all HSMDA queues indirectly associated with the policy.

### Default HSMDA Slope Policy

An hsmda-slope-policy named **default** always exists on the system and does not need to be created. The default policy is automatically applied to all HSMDA queues unless another HSMDA slope policy is specified for the queue. The default policy cannot be modified or deleted. Attempting to execute no hsmda-slope-policy default will result in an error.

The **no** form of the command removes the slope policy from the subscriber profile HSMDA configuration.

## stat-mode

<b>Syntax</b>	<b>stat-mode {v4-v6}</b> <b>no stat-mode</b>
<b>Context</b>	config>subscr-mgmt>sub-prof>hsmda>ingress-qos>qos>policer config>subscr-mgmt>sub-prof>hsmda>ingress-qos>qos>queue config>subscr-mgmt>sub-prof>hsmda>egress-qos>qos>queue
<b>Description</b>	<p>This command configures the forwarding plane octet and packet counters of a policer or queue to count packets of a specific type or state. For example separate counters for IPv4/IPv6.</p> <p>For HSMDA ingress policers, this command overrides the policer stat-mode configuration as defined in the sap-ingress qos policy. For details on sap-ingress and sap-egress policer stat-mode, refer to the 7750 SR OS Quality of Service Guide. For use in Enhanced Subscriber Management (ESM) context only, an additional stat-mode enables separate counters for IPv4 and IPv6 packets. <b>tat-mode v4-v6</b> is the only mode that can be configured as an HSMDA ingress policer override.</p> <p>An HSMDA policer's stat-mode cannot be changed while the sub profile is in use.</p> <p>For queues, this command sets the stat-mode. Queue stat-mode is only available for use in ESM context to enable separate IPv4/IPv6 counters.</p> <p>An HSMDA queue's stat-mode cannot be changed while the sub profile is in use.</p>
<b>Default</b>	<p>no stat-mode</p> <p>For policers, the default is no stat-mode override. The <b>sap-ingress stat-mode</b> is used instead.</p> <p>For queues, the default is to <b>count in-/out-of-profile</b> octets and packets.</p>
<b>Parameters</b>	<b>v4-v6</b> — Count IPv4 and IPv6 forwarded/dropped octets and packets separately

## wrr-weight

<b>Syntax</b>	<b>wrr-weight value</b> <b>no wrr-weight</b>
<b>Context</b>	config>subscr-mgmt>sub-prof>hsmda>egress-qos>qos>queue
<b>Description</b>	This command assigns the weight value to the HSMDA queue.

The **no** form of the command returns the weight value for the queue to the default value.

**Parameters** *percentage* — Specifies the weight for the HSMDA queue.  
**Values** 1— 32

## wrr-policy

**Syntax** **wrr-policy** *hsmda-wrr-policy-name*  
**no wrr-policy**

**Context** config>subscr-mgmt>sub-prof>hsmda>egress-qos>qos

**Description** This command associates an existing HSMDA weighted-round-robin (WRR) scheduling loop policy to the HSMDA queue.

**Parameters** *hsmda-wrr-policy-name* — Specifies the existing HSMDA WRR policy name to associate to the queue.



---

## Explicit Subscriber Mapping Commands

### explicit-sub-map

<b>Syntax</b>	<b>explicit-sub-map</b>
<b>Context</b>	config>subscr-mgmt
<b>Description</b>	This command configures an explicit subscriber mapping

### entry

<b>Syntax</b>	<b>entry key</b> <i>sub-ident-string</i> [ <b>sub-profile</b> <i>sub-profile-name</i> ] [ <b>alias</b> <i>sub-alias-string</i> ] [ <b>sla-profile</b> <i>sla-profile-name</i> ] <b>no entry key</b> <i>sub-profile-string</i>
<b>Context</b>	config>subscr-mgmt>explicit-sub-map
<b>Description</b>	This command configures a subscriber identification string.
<b>Parameters</b>	<i>sub-ident-string</i> — Specifies the profile string. <b>Values</b> 16 characters maximum <i>sub-profile-name</i> — Specifies an existing subscriber profile name. <b>Values</b> 32 characters maximum <b>alias</b> <i>sub-alias-string</i> — Specifies an alias for the subscriber identification string. <b>sla-profile</b> <i>sla-profile-name</i> — Specifies an existing SLA profile.

---

# Subscriber Management Service Commands

---

## SAP Subscriber Management Commands

---

### sub-sla-mgmt

<b>Syntax</b>	<b>[no] sub-sla-mgmt</b>
<b>Context</b>	config>service>vpls>sap
<b>Description</b>	This command enables the context to configure subscriber management parameters for this SAP.
<b>Default</b>	no sub-sla-mgmt

### def-sla-profile

<b>Syntax</b>	<b>def-sla-profile</b> <i>default-sla-profile-name</i> <b>no def-sla-profile</b>
<b>Context</b>	config>service>vpls>sap>sub-sla-mgmt
<b>Description</b>	<p>This command specifies a default SLA profile for this SAP.</p> <p>An SLA profile is a named group of QoS parameters used to define per service QoS for all subscriber hosts common to the same subscriber within a provider service offering. A single SLA profile may define the QoS parameters for multiple subscriber hosts.</p> <p>The <b>no</b> form of the command removes the default SLA profile from the SAP configuration.</p>
<b>Default</b>	no def-sla-profile
<b>Parameters</b>	<i>default-sla-profile-name</i> — Specifies a default SLA profile for this SAP.

### def-sub-profile

<b>Syntax</b>	<b>def-sub-profile</b> <i>default-subscriber-profile-name</i>
<b>Context</b>	config>service>vpls>sap>sub-sla-mgmt
<b>Description</b>	<p>This command specifies a default subscriber profile for this SAP.</p> <p>A subscriber profile defines the aggregate QoS for all hosts within a subscriber context. This is done through the definition of the egress and ingress scheduler policies that govern the aggregate SLA for subscriber using the subscriber profile.</p> <p>The <b>no</b> form of the command removes the default SLA profile from the SAP configuration.</p>
<b>Parameters</b>	<i>default-sub-profile</i> — Specifies a default subscriber profile for this SAP.



## sub-ident-policy

<b>Syntax</b>	<b>sub-ident-policy</b> <i>sub-ident-policy-name</i>
<b>Context</b>	config>service>vpls>sap>sub-sla-mgmt
<b>Description</b>	<p>This command associates a subscriber identification policy to this SAP.</p> <p>Subscribers are managed by the system through the use of subscriber identification strings. A subscriber identification string uniquely identifies a subscriber. For static hosts, the subscriber identification string is explicitly defined with each static subscriber host.</p> <p>For dynamic hosts, the subscriber identification string must be derived from the DHCP ACK message sent to the subscriber host. The default value for the string is the content of Option 82 CIRCUIT-ID and REMOTE-ID fields interpreted as an octet sting. As an option, the DHCP ACK message may be processed by a subscriber identification policy which has the capability to parse the message into an alternative ASCII or octet string value.</p> <p>When multiple hosts on the same port are associated with the same subscriber identification string they are considered to be host members of the same subscriber.</p> <p>The <b>no</b> form of the command removes the default subscriber identification policy from the SAP configuration.</p>
<b>Default</b>	no sub-ident-policy
<b>Parameters</b>	<i>sub-ident-policy-name</i> — Specifies a subscriber identification policy for this SAP.

## multi-sub-sap

<b>Syntax</b>	<b>multi-sub-sap</b> <i>number-of-sub</i> <b>no multi-sub-sap</b>
<b>Context</b>	config>service>vpls>sap>sub-sla-mgmt
<b>Description</b>	<p>This command defines the maximum number of subscribers (dynamic + static) that can be simultaneously active on this SAP.</p> <p>If the limit is reached, a new host will be denied access and the corresponding DHCP ACK will be dropped.</p>
<b>Default</b>	1
	The <b>no</b> form of the command reverts back to the default setting.
<b>Default</b>	no multi-sub-sap
<b>Parameters</b>	<i>multi-sub-sap</i> — Specifies the maximum allowed.

## single-sub-parameters

<b>Syntax</b>	<b>single-sub-parameters</b>
<b>Context</b>	config>service>vpls>sap>sub-sla-mgmt

**Description** This command configure single subscriber SAP parameters.

### non-sub-traffic

**Syntax** **non-sub-traffic sub-profile** *sub-profile-name* **sla-profile** *sla-profile-name* [**subscriber** *sub-ident-string*]  
**no non-sub-traffic**

**Context** config>service>vpls>sap>sub-sla-mgmt>single-sub

**Description** This command configures traffic profiles for non-IP traffic such as PPPoE. It is used in conjunction with the **profiled-traffic-only** on single subscriber SAPs and creates a subscriber host which is used to forward non-IP traffic through the single subscriber SAP without the need for SAP queues.  
The **no** form of the command removes any configured profile.

**Default** no non-sub-traffic

**Parameters** *sub-profile-name* — Identifies the subscriber profile name.  
**Values** 32 characters maximum  
*sla-profile-name* — Identifies the SLA profile name.  
**Values** 32 characters maximum

### profiled-traffic-only

**Syntax** [**no**] **profiled-traffic-only**

**Context** config>service>vpls>sap>sub-sla-mgmt>single-sub-parameters

**Description** This command specifies whether only profiled traffic is applicable for this SAP. The profiled traffic refers to single subscriber traffic on a dedicated SAP (in the VLAN-per-subscriber model). When enabled, subscriber queues are instantiated through the QOS policy defined in the *sla-profile* and the associated SAP queues are deleted. This can increase subscriber scaling by reducing the number of queues instantiated per subscriber (in the VLAN-per-subscriber model). In order for this to be achieved, any configured multi-sub-sap limit must be removed (leaving the default of 1).  
The **no** form of the command reverts to the default setting.

**Default** no profiled-traffic-only

**Default**

### mac-da-hashing

**Syntax** **mac-da-hashing**  
**no mac-da-hashing**

**Context** config>service>vpls>sap>sub-sla-mgmt

<b>Description</b>	This command specifies whether subscriber traffic egressing a LAG SAP has its egress LAG link selected by a function of the MAC destination address instead of the subscriber ID.  The <b>no</b> form of the command reverts to the default setting.
<b>Default</b>	no mac-da-hashing

## diameter-auth-policy

<b>Syntax</b>	<b>diameter-auth-policy</b> <i>name</i> <b>no diameter-auth-policy</b>
<b>Context</b>	config>service>vpls>sap
<b>Description</b>	This command is used to configure the Diameter NASREQ application policy to use for authentication.
<b>Parameters</b>	<i>name</i> — Specifies the name of the Diameter NASREQ application policy to use for authentication.

## host

<b>Syntax</b>	<b>host</b> {[ <i>ip ip-address</i> [ <i>mac mac-address</i> ]]} [ <b>subscriber-sap-id</b>   <b>subscriber</b> <i>sub-ident-string</i> [ <b>sub-profile</b> <i>sub-profile-name</i> [ <b>sla-profile</b> <i>sla-profile-name</i> [ <b>anccp-string</b> <i>anccp-string</i> ] [ <b>app-profile</b> <i>app-profile-name</i> ] [ <b>inter-dest-id</b> <i>intermediate-destination-id</i> ]]] <b>no host</b> {[ <i>ip ip-address</i> ] [ <i>mac ieee-address</i> ]]} <b>no host all</b>
<b>Context</b>	config>service>vpls>sap config>service>ies>if>sap
<b>Description</b>	<p>This command creates a static subscriber host for the SAP. Static subscriber hosts may be used by the system for various purposes. Applications within the system that make use of static host entries include anti-spoof, ARP reply agent and source MAC population into the VPLS forwarding database.</p> <p>Multiple static hosts may be defined on the SAP. Each host is identified by either a source IP address, a source MAC address or both a source IP and source MAC address. Every static host definition must have at least one address defined, IP or MAC.</p> <p>Static hosts can exist on the SAP even with anti-spoof and ARP reply agent features disabled. When enabled, each feature has different requirements for static hosts.</p> <p>Use the <b>no</b> form of the command to remove a static entry from the system. The specified <i>ip-address</i> and <i>mac-address</i> must match the host's exact IP and MAC addresses as defined when it was created. When a static host is removed from the SAP, the corresponding anti-spoof filter entry and/or FDB entry is also removed.</p>
<b>Default</b>	none
<b>Parameters</b>	<b>ip</b> <i>ip-address</i> — Specify this parameter to associate a subscriber with the static subscriber host. Only one static host can be configured on the SAP with a given IP address.  <b>mac</b> <i>mac-address</i> — Specify this optional parameter when defining a static host. The MAC address must be specified for <b>anti-spoof mac</b> , <b>anti-spoof ip-mac</b> and arp-reply-agent. Multiple static

hosts may be configured with the same MAC address given that each definition is distinguished by a unique IP address.

Every static host definition must have at least one address defined, IP or MAC.

**subscriber** *sub-ident-string* — Specify this parameter to configure an existing subscriber identification profile to be associated with the static subscriber host. The subscriber identification profile is configured in the **config>subscr-mgmt>sub-ident-policy** context. The subscriber information is used by the VPLS SAP arp-reply-agent to determine the proper handling of received ARP requests from subscribers.

- For VPLS SAPs with **arp-reply-agent** enabled with the optional *sub-ident* parameter, the static subscriber hosts *sub-ident-string* is used to determine whether an ARP request received on the SAP is sourced from a host belonging to the same subscriber as the destination host. When both the destination and source hosts from the ARP request are known on the SAP and the subscriber identifications do not match, the ARP request may be forwarded to the rest of the VPLS destinations.

If the static subscriber hosts *sub-ident* string is not defined, the host is not considered to belong to the same subscriber as another host on the SAP.

If source or destination host is unknown, the hosts are not considered to belong to the same subscriber. (ARP messages from unknown hosts are subject to anti-spoof filtering rules applied at the SAP.)

If *sub-ident* is not enabled on the SAP arp-reply-agent, subscriber identification matching is not performed on ARP requests received on the SAP.

ARP requests are never forwarded back to the same SAP or within the receiving SAP's Split Horizon Group.

**sub-profile** *sub-profile-name* — Specify this parameter to configure an existing subscriber profile name to be associated with the static subscriber host. The subscriber profile is configured in the **config>subscr-mgmt>sub-profile** context.

**sla-profile** *sla-profile-name* — Specify this parameter to configure an existing SLA profile name to be associated with the static subscriber host. The SLA profile is configured in the **config>subscr-mgmt>sla-profile** context.

Note that if Enhanced Subscriber Management is enabled on a SAP using the **sub-sla-mgmt** command, the **sub-ident**, **sub-profile**, and **sla-profile** must be configured for all static hosts defined on this SAP.

## Wireless Portal Protocol (WPP) Commands

### wpp

<b>Syntax</b>	<b>wpp</b>
<b>Context</b>	config>service>ies>sub-if>grp-if config>service>vprn>sub-if>grp-if
<b>Description</b>	This command enables the context to configure Wireless Portal Protocol (WPP) parameters.

### enable-triggered-hosts

<b>Syntax</b>	<b>[no] enable-triggered-hosts</b>
<b>Context</b>	config>service>vprn>sub-if>grp-if>wpp config>service>ies>sub-if>grp-if>wpp
<b>Description</b>	This command enables system to auto creates ESM hosts upon successful WPP authentication. Default host need to be configured under SAP on the subscriber SAP in order to redirection un-authentication client traffic to web portal.
<b>Default</b>	none

### initial-app-profile

<b>Syntax</b>	<b>initial-app-profile <i>app-profile-name</i></b> <b>no initial-app-profile</b>
<b>Context</b>	config>subscr-mgmt>loc-user-db>ipoe>host>wpp config>service>ies>sub-if>grp-if>wpp config>service>vprn>sub-if>grp-if>wpp
<b>Description</b>	This command specifies the initial app-profile for the hosts created on the group-interface. This initial app-profile will be replaced after hosts pass web portal authentication.
<b>Default</b>	none
<b>Parameters</b>	<i>app-profile-name</i> — Specifies the initial application profile, to be used during the WPP authentication phase of the IPoE hosts.

### initial-sla-profile

<b>Syntax</b>	<b>initial-sla-profile <i>sla-profile-name</i></b> <b>no initial-sla-profile</b>
---------------	---

## Subscriber Management Service Commands

<b>Context</b>	config>subscr-mgmt>loc-user-db>ipoe>host>wpp config>service>ies>sub-if>grp-if>wpp config>service>vprn>sub-if>grp-if>wpp
<b>Description</b>	This command specifies the initial sla-profile for the hosts created on the group-interface. This initial sla-profile will be replaced after hosts pass web portal authentication.
<b>Default</b>	none
<b>Parameters</b>	<i>sla-profile-name</i> — Specifies the initial SLA profile to be used during the WPP authentication phase of the IPOE host.

## initial-sub-profile

<b>Syntax</b>	<b>initial-sub-profile</b> <i>sub-profile-name</i> <b>no initial-sub-profile</b>
<b>Context</b>	config>subscr-mgmt>loc-user-db>ipoe>host>wpp config>service>ies>sub-if>grp-if>wpp config>service>vprn>sub-if>grp-if>wpp
<b>Description</b>	This command specifies the initial sub-profile for the hosts created on the group-interface. This initial sub-profile will be replaced after hosts pass web portal authentication.
<b>Default</b>	none
<b>Parameters</b>	<i>sub-profile-name</i> — specifies the initial subscriber profile, to be used during the WPP authentication phase of the IPoE host.

## portals

<b>Syntax</b>	<b>portals</b>
<b>Context</b>	config>router>wpp config>service>vprn>wpp
<b>Description</b>	This command enables the context to configure WPP portal server parameters.

## portal

<b>Syntax</b>	<b>portal router</b> <i>router-instance name</i> <i>wpp-portal-name</i> <b>no portal</b>
<b>Context</b>	config>subscr-mgmt>loc-user-db>ipoe>host>wpp config>service>ies>sub-if>grp-if>wpp config>service>vprn>sub-if>grp-if>wpp
<b>Description</b>	This command specifies the web portal server that system talks to for the hosts on the group-interface.
<b>Default</b>	none

**router** *router-instance* — Specifies the virtual router instance.

**Values**      router-name:      Base, management  
                  service-id:        1 — 2147483647  
                  service-name:    Specifies the service name up to 64 characters in length.

**Default**      Base

**name** *wpp-portal-name* — Specifies the name of the web portal server.

## lease-time

**Syntax**      **lease-time** [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]  
**no lease-time**

**Context**      config>service>vprn>sub-if>grp-if>wpp

**Description**    This command specifies the lease time of the trigger created by the ESM host by WPP authentication.

**Parameters**    **days** *days* — Specifies the lease time in days.

**Values**      0 — 3650

**hrs** *hours* — Specifies the lease time in hours.

**Values**      1 — 23

**min** *minutes* — Specifies the lease time in minutes.

**Values**      1 — 59

**sec** *seconds* — Specifies the lease time in seconds.

**Values**      0 — 50

## restore-disconnected

**Syntax**      **restore-disconnected** {**restore**|**no-restore**}  
**no restore-disconnected**

**Context**      config>subscr-mgmt>loc-user-db>ipoe>host>wpp  
                  config>service>ies>sub-if>grp-if>wpp  
                  config>service>vprn>sub-if>grp-if>wpp

**Description**    This command specifies the behavior that system will restore the initial-sla-profile/initial-sub-profile/initial-aa-prfprofile when hosts disconnects instead of removing them.

**Default**      none

**Parameters**    **restore** — Specifies that the initial profiles must be restored after a DHCP host has disconnected.

**no-restore** — Specifies that the initial profiles will not be restored after a DHCP host has disconnected.

## user-db

<b>Syntax</b>	<b>user-db</b> <i>local-user-db-name</i> <b>no user-db</b>
<b>Context</b>	config>subscr-mgmt>loc-user-db>ipoe>host>wpp config>service>ies>sub-if>grp-if>wpp config>service>vprn>sub-if>grp-if>wpp
<b>Description</b>	<p>This command configures the user database. Note that if configured, the values configured under grp-if will only be used if there is no corresponding value returned from LUDB lookup.</p> <p>This command specifies the LUDB system use to lookup while creating initial host before WPP authentication. LUDB could return WPP attributes such as portal name, initial-sla-profile, initial-sub-profile, etc. LUDB is configured in <b>config&gt;subscr-mgmt&gt;local-user-db</b> context.</p>
<b>Default</b>	none
<b>Parameters</b>	<i>local-user-db-name</i> — Specifies the Local User Database name.

## vpi

<b>Syntax</b>	<b>vpi</b> <i>vpi egress-traffic-desc atm-td-profile-id</i> <b>no vpi</b> <i>vpi</i>
<b>Context</b>	configure>port>sonet-sdh>path>atm
<b>Description</b>	<p>This command enables the ATM VP shaper under the ATM port. The type of ATM shaper are CBR or rt/nrt-VBR as defined by the traffic descriptor. It cannot be a UBR service-type.</p> <p>All VCs within the shaper will degrade into a UBR type service class. For example, when a CBR type VC is associated with the shaper, it will degrade into a UBR type VC. Scheduling traffic amongst VCs within the shaper is based on WRR using the weight parameter.</p> <p>If the VP shaper is deleted, the VCs that were under it is restored to their original service category.</p> <p>The VP shaper is statically configured and instantiated upon configuration.</p> <p>A VP shaper can be seamlessly added to or removed from the active VCs in the system.</p>
<b>Default</b>	none
<b>Parameters</b>	<i>atm-td-profile-id</i> — Specifies ATM traffic description id. <div style="margin-left: 40px;"><b>Values</b> [1..1000]</div> <i>vpi</i> — <div style="margin-left: 40px;"><b>Values</b> [0..4095]</div> <b>egress-traffic-desc</b> — References an atm traffic descriptor profile.



## Layer 3 Subscriber Interfaces SAP Commands

### anti-spoof

<b>Syntax</b>	<b>anti-spoof {ip   ip-mac   nh-mac}</b> <b>no anti-spoof</b>
<b>Context</b>	config>service>ies>sub-if>grp-if>sap config>subscr-mgmt>msap-policy
<b>Description</b>	<p>This command configures the anti-spoof type of the MSAP.</p> <p>The type of anti-spoof filtering defines what information in the incoming packet is used to generate the criteria to lookup an entry in the anti-spoof filter table. The type parameter (<b>ip</b>, <b>ip-mac</b>) defines the anti-spoof filter type enforced by the SAP when anti-spoof filtering is enabled.</p> <p>The <b>no</b> form of the command reverts back to the default.</p> <p>Note that for IES subscriber group interfaces, setting no anti-spoof will set the default anti-spoofing type which is <b>ip-mac</b>.</p>
<b>Default</b>	no anti-spoof
<b>Parameters</b>	<p><b>ip</b> — Configures SAP anti-spoof filtering to use only the source IP address in its lookup. If a static host exists on the SAP without an IP address specified, the anti-spoof type <b>ip</b> command will fail. Note that this parameter is not applicable in the <b>config&gt;subscr-mgmt&gt;msap-policy</b> context.</p> <p><b>ip-mac</b> — Configures SAP anti-spoof filtering to use both the source IP address and the source MAC address in its lookup. If a static host exists on the SAP without both the IP address and MAC address specified, the anti-spoof type <b>ip-mac</b> command will fail. This is also true if the default anti-spoof filter type of the SAP is <b>ip-mac</b> and the default is not overridden. The anti-spoof type <b>ip-mac</b> command will also fail if the SAP does not support Ethernet encapsulation.</p> <p><b>nh-mac</b> — Indicates that the ingress anti-spoof is based on the source MAC and egress anti-spoof is based on the nh-ip-address.</p>

### cpu-protection

<b>Syntax</b>	<b>cpu-protection <i>policy-id</i> [mac-monitoring]   [eth-cfm-monitoring [aggregate] [car]]</b> <b>no cpu-protection</b>
<b>Context</b>	config>service>ies>sub-if>grp-if>sap
<b>Description</b>	<p>This command assigns an existing CPU protection policy to the associated group interface. The CPU protection policies are configured in the <b>config&gt;sys&gt;security&gt;cpu-protection&gt;policy <i>cpu-protection-policy-id</i></b> context.</p> <p>If no CPU-Protection policy is assigned to a group interface SAP, then the default policy is used to limit the overall-rate. The default policy is policy number 254 for access interfaces and 255 for network interfaces.</p>

The **no** form of the command removes the association of the CPU protection policy from the associated interface and reverts to the default policy values.

**Default**     `cpu-protection 254` (for access interfaces)  
`cpu-protection 255` (for network interfaces)

The configuration of `no cpu-protection` returns the interface/SAP to the default policies as shown above.

**Parameters**     *policy-id* — Specifies an existing CPU protection policy.

**Values**         1 — 255

**mac-monitoring** — This keyword enables MAC monitoring.

**eth-cfm-monitoring** — This keyword enables Ethernet Connectivity Fault Management monitoring.

**aggregate** — This keyword applies the rate limit to the sum of the per peer packet rates.

**car** — (Committed Access Rate) This keyword causes Eth-CFM packets to be ignored when enforcing the overall-rate.

## SDP Binding Commands

### binding

<b>Syntax</b>	<b>binding</b>
<b>Context</b>	config>service>sdp
<b>Description</b>	The command enables the context to configure SDP bindings.

### port

<b>Syntax</b>	<b>port</b> [ <i>port-id</i>   <i>lag-id</i> ] <b>no port</b>
<b>Context</b>	config>service>sdp>binding
<b>Description</b>	This command specifies the port or lag identifier, to which the PW ports associated with the underlying SDP are bound. If the underlying SDP is re-routed to a port or lag other than the specified one, the PW ports on the SDP are operationally brought down.  The <b>no</b> form of the command removes the value from the configuration.
<b>Default</b>	none
<b>Parameters</b>	<i>port-id</i> — The identifier of the port in the slot/mda/port format. <i>lag-id</i> — Specifies the LAG identifier.

### pw-port

<b>Syntax</b>	<b>pw-port</b> <i>pw-port-id</i> [ <i>vc-id vc-id</i> ] [ <b>create</b> ] <b>no pw-port</b>
<b>Context</b>	config>service>sdp>binding
<b>Description</b>	This command creates a PW-port.  The <b>no</b> form of the command removes the PW-port ID from the configuration.
<b>Default</b>	none
<b>Parameters</b>	<i>pw-port-id</i> — Specifies a unique identifier of the PW port.  <b>Values</b> 1 — 10239 <i>vc-id vc-id</i> — Specifies a virtual circuit identifier signaled to the peer.  <b>Values</b> 1 — 4294967295

## description

<b>Syntax</b>	<b>description</b> <i>description-string</i> <b>no description</b>
<b>Context</b>	config>service>sdp>binding>pw-port
<b>Description</b>	This command creates a text description stored in the configuration file for a configuration context. The description command associates a text string with a configuration context to help identify the content in the configuration file. The <b>no</b> form of the command removes the string from the configuration.
<b>Default</b>	no description
<b>Parameters</b>	<i>description-string</i> — Specifies the description character string of the configuration context. <b>Values</b> Any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

## egress

<b>Syntax</b>	<b>egress</b>
<b>Context</b>	config>service>sdp>binding>pw-port
<b>Description</b>	This command enables the context to configure PW-port egress side parameters.

## encap-type

<b>Syntax</b>	<b>encap-type</b> {dot1q qinq} <b>no encap-type</b>
<b>Context</b>	config>service>sdp>binding>pw-port
<b>Description</b>	This command sets the encapsulation type for the PW-port as dot1q or qinq.
<b>Default</b>	dot1q
<b>Parameters</b>	<b>dot1q</b> — Specifies <b>dot1q</b> encapsulation type. <b>qinq</b> — Specifies <b>qinq</b> encapsulation type.

## shaper

<b>Syntax</b>	<b>shaper</b> <b>no shaper</b>
<b>Context</b>	config>service>sdp>binding>pw-port>egress

<b>Description</b>	This command configures an egress shaping option for use by a PW port..
<b>Default</b>	no shaper.

## int-dest-id

<b>Syntax</b>	<b>[no] int-dest-id</b> <i>int-dest-id</i>
<b>Context</b>	config>service>sdp>binding>pw-port>egress>shaper
<b>Description</b>	This command specifies the intermediate destination string configured for dynamic vport selection. The <b>no</b> form of the command removes the configured intermediate destination string. This command is only valid for PW ports used for enhanced subscriber management (ESM on PW).
<b>Default</b>	no .int-dest-id
<b>Parameters</b>	<i>int-dest-id</i> — A text string that describes the intermediate destination ID.

## pw-sap-secondary-shaper

<b>Syntax</b>	<b>[no] pw-sap-secondary-shaper</b> <i>secondary-shaper-name</i>
<b>Context</b>	config>service>sdp>binding>pw-port>egress>shaper
<b>Description</b>	This command configures the use of secondary shaper name as a reference to a shaper to use for a PW port on the HSMDA. The <b>no</b> form of the command removes the configured shaper. This command is valid for PW ports used PW SAPs on the HSMDA.
<b>Default</b>	no pw-sap-secondary-shaper
<b>Parameters</b>	<i>secondary-shaper-name</i> — specifies a text string representing the name of the secondary shaper.

## vport

<b>Syntax</b>	<b>[no] vport</b> <i>vport-name</i>
<b>Context</b>	config>service>sdp>binding>pw-port>egress>shaper
<b>Description</b>	This command configures the name of the vport to be used for the PW port. The <b>no</b> form of the command removes the configured vport name. This command is valid for PW ports used for enhanced subscriber management (ESM on pseudowire) and pseudowire SAPs on Ethernet ports. It is not valid for pseudowire ports on the HSMDA.
<b>Default</b>	no vport
<b>Parameters</b>	<i>vport-name</i> — Specifies a text string representing the name of the vport.

## vc-type

<b>Syntax</b>	<b>vc-type {ether vlan}</b> <b>no vc-type</b>
<b>Context</b>	config>service>sdp>binding>pw-port
<b>Description</b>	<p>This command sets the forwarding mode for PW-port. The vc-type is signaled to the peer, and must be configured consistently on both ends of the PW. vc-type VLAN is only configurable with dot1q encapsulation on the PW-port. The tag with vc-type vlan only has significance for transport, and is not used for service delineation or ESM. The top (provider tag) is stripped while forwarding out of the PW, and a configured vlan-tag (for vc-type vlan) is inserted when forwarding into the PW. With vc-type ether, the tags if present (max 2), are transparently preserved when forwarding in our out of the PW.</p> <p>The <b>no</b> form of the command reverts to the default value.</p>
<b>Default</b>	ether
<b>Parameters</b>	<p><b>ether</b> — Specifies <b>ether</b> as the virtual circuit (VC) associated with the SDP binding.</p> <p><b>vlan</b> — Specifies <b>vlan</b> as the virtual circuit (VC) associated with the SDP binding.</p>

## vlan-vc-tag

<b>Syntax</b>	<b>vlan-vc-tag <i>vlan-id</i></b> <b>no vc-type</b>
<b>Context</b>	config>service>sdp>binding>pw-port
<b>Description</b>	<p>This command sets tag relevant for vc-type vlan mode. This tag is inserted in traffic forwarded into the PW.</p> <p>The <b>no</b> form of the command reverts to the default value.</p>
<b>Default</b>	0
<b>Parameters</b>	<p><i>vlan-id</i> — Specifies the VLAN ID value.</p> <p><b>Values</b>      0 — 4094</p>

## RIP Commands

### rip-policy

<b>Syntax</b>	<b>rip-policy</b> <i>policy-name</i> [create] <b>no rip-</b> <i>policy-name</i>
<b>Context</b>	config>subscr-mgmt
<b>Description</b>	This command creates a RIP policy. This policy is applied to a subscriber IPv4 host to enable the BNG to learn RIP routes from the host. RIP routes are never sent to the hosts.
<b>Default</b>	none
<b>Parameters</b>	<i>policy-name</i> — Specifies the RIP policy name up to 32 characters in length.

### neighbor

<b>Syntax</b>	[no] <b>neighbor</b> <i>ip-int-name</i>
<b>Context</b>	config>router>rip>group config>service>vprn>rip>group
<b>Description</b>	<p>This command creates a context for configuring a RIP neighbor interface. By default, group interfaces are not activated with RIP, unless explicitly configured. The BNG will only learn RIP routes from IPv4 host on the group interface. Hence, RIP neighbor group interface will default send to “none”. The send operation is unchangeable for group-interface.</p> <p>The no form of the command deletes the RIP interface configuration for this group interface. The shutdown command in the <b>config&gt;router&gt;rip&gt;group group-name&gt;neighbor</b> context can be used to disable an interface without removing the configuration for the interface.</p>
<b>Default</b>	no neighbor — No RIP interfaces are defined.
<b>Parameters</b>	<i>ip-int-name</i> — The group interface name. Interface names must be unique within the group of defined group interfaces within config service vprn/ies sub-interface grp-interface commands. An interface name cannot be in the form of an IP address. Interface names can be any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes. If the group interface name does not exist, an error message will be returned.

### authentication-key

<b>Syntax</b>	<b>authentication-key</b> [ <i>authentication-key</i>   <i>hash-key</i> ] [ <b>hash</b>   <b>hash2</b> ] <b>no authentication-key</b>
<b>Context</b>	config>subscr-mgmt>rip-policy

<b>Description</b>	<p>This command configures the BGP authentication key.</p> <p>Authentication is performed between neighboring routers before setting up the BGP session by verifying the password. Authentication is performed using the MD-5 message-based digest. The authentication key can be any combination of letters or numbers from 1 to 16.</p> <p>The no form of the command removes the authentication password from the configuration and effectively disables authentication.</p>
<b>Default</b>	Authentication is disabled and the authentication password is empty.
<b>Parameters</b>	<p><i>authentication-key</i> — The authentication key. The key can be any combination of ASCII characters up to 255 characters in length (unencrypted). If spaces are used in the string, enclose the entire string in quotation marks (“ ”).</p> <p><i>hash-key</i> — The hash key. The key can be any combination of ASCII characters up to 342 characters in length (encrypted). If spaces are used in the string, enclose the entire string in quotation marks (“ ”).</p> <p>This is useful when a user must configure the parameter, but, for security purposes, the actual unencrypted key value is not provided.</p> <p><b>hash</b> — Specifies the key is entered in an encrypted form. If the <b>hash</b> parameter is not used, the key is assumed to be in a non-encrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the <b>hash</b> parameter specified.</p> <p><b>hash2</b> — Specifies the key is entered in a more complex encrypted form. If the <b>hash2</b> parameter is not used, the less encrypted <b>hash</b> form is assumed.</p>

## authentication-type

<b>Syntax</b>	<b>authentication-type {none password message-digest message-digest-20}</b> <b>no authentication-type</b>
<b>Context</b>	config>sub-mgmt>rip-policy>
<b>Description</b>	<p>This command sets the type of authentication to be used between RIP neighbors. The type and password must match exactly for the RIP message to be considered authentic and processed.</p> <p>The <b>no</b> form of the command removes the authentication type from the configuration and effectively disables authentication.</p>
<b>Default</b>	no authentication-type — No authentication enabled.
<b>Parameters</b>	<p><b>none</b> — The none parameter explicitly disables authentication at a given level (global, group, neighbor). If the command does not exist in the configuration, the parameter is inherited.</p> <p><b>password</b> — Specify password to enable simple password (plain text) authentication. If authentication is enabled and no authentication type is specified in the command, simple password authentication is enabled.</p> <p><b>message-digest</b> — Configures 16 byte message digest for MD5 authentication. If this option is configured, then at least one message-digest-key must be configured.</p>



**message-digest-20** — Configures 20 byte message digest for MD5 authentication in accordance with RFC 2082, RIP-2 MD5 Authentication. If this option is configured, then at least one message-digest-key must be configured.

## retail-svc-id

<b>Syntax</b>	<b>retail-svc-id</b> <i>service-id</i> <b>retail-svc-id</b>
<b>Context</b>	config>service>ies vprn>sub-if>grp-if>sap>static-host
<b>Description</b>	This command specifies the service id of the retailer IES/VRN service to which the static IPv6 host belongs. A corresponding retailer subscriber interface must exist in the specified service.
<b>Default</b>	no retail-svc-id
<b>Parameters</b>	<i>service-id</i> — Specifies the retailer service ID. <b>Values</b> 1 — 2148007978

## rip

<b>Syntax</b>	<b>[no] rip</b>
<b>Context</b>	config>service>vprn config>service>ies
<b>Description</b>	This command enables the RIP protocol on the given VPRN IP interface. The <b>no</b> form of the command disables the RIP protocol from the given VPRN IP interface.
<b>Default</b>	no rip

## group

<b>Syntax</b>	<b>[no] group</b> <i>group-name</i>
<b>Context</b>	config>service>vprn>rip config>service>ies>rip
<b>Description</b>	This command creates a context for configuring a RIP group of neighbors. RIP groups are a way of logically associating RIP neighbor interfaces to facilitate a common configuration for RIP interfaces. The <b>no</b> form of the command deletes the RIP neighbor interface group. Deleting the group will also remove the RIP configuration of all the neighbor interfaces currently assigned to this group.
<b>Default</b>	<b>no group</b> — No group of RIP neighbor interfaces defined
<b>Parameters</b>	<i>group-name</i> — The RIP group name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

## Vport Commands

---

### ethernet

<b>Syntax</b>	<b>ethernet</b>
<b>Context</b>	config>port
<b>Description</b>	<p>This command enables access to the context to configure Ethernet port attributes.</p> <p>This context can only be used when configuring Fast Ethernet, gigabit or 10Gig Fast Ethernet or Ethernet LAN ports on an appropriate MDA.</p>

### egress-scheduler-override

<b>Syntax</b>	<b>[no] egress-scheduler-override</b>
<b>Context</b>	config>port>ethernet
<b>Description</b>	<p>This command applies egress scheduler overrides. When a port scheduler is associated with an egress port, it is possible to override the following parameters:</p> <ul style="list-style-type: none"> <li>• The <b>max-rate</b> allowed for the scheduler.</li> <li>• The maximum <b>rate</b> for each priority level 8 through 1.</li> <li>• The CIR associated with each priority level 8 through 1.</li> </ul> <p>See the SR OS Quality of Service Guide for command syntax and usage for the <b>port-scheduler-policy</b> command.</p> <p>The <b>no</b> form of this command removes all override parameters from the egress port or channel scheduler context. Once removed, the port scheduler reverts all rate parameters back to the parameters defined on the port-scheduler-policy associated with the port.</p>

### level

<b>Syntax</b>	<b>level</b> <i>priority-level</i> <b>rate</b> <i>pir-rate</i> [ <b>cir</b> <i>cir-rate</i> ] <b>no level</b> <i>priority-level</i>
<b>Context</b>	config>port>ethernet>egress-scheduler-override
<b>Description</b>	<p>This command overrides the maximum and CIR rate parameters for a specific priority level on the port or channel's port scheduler instance. When the <b>level</b> command is executed for a priority level, the corresponding priority level command in the port-scheduler-policy associated with the port is ignored.</p> <p>The override level command supports the keyword <b>max</b> for the <b>rate</b> and <b>cir</b> parameter.</p> <p>When executing the level override command, at least the <b>rate</b> or <b>cir</b> keywords and associated parameters must be specified for the command to succeed.</p>

The **no** form of this command removes the local port priority level rate overrides. Once removed, the port priority level will use the port scheduler policies level command for that priority level.

<b>Parameters</b>	<i>priority-level</i> — Identifies which of the eight port priority levels are being overridden.
<b>Values</b>	1 — 8
<b>rate</b> <i>pir-rate</i> — Overrides the port scheduler policy's maximum level rate and requires either the <b>max</b> keyword or a rate defined in kilobits-per-second to follow.	
<b>Values</b>	1 — 40000000, max
<b>cir</b> <i>cir-rate</i> — Overrides the port scheduler policy's within-cir level rate and requires either the <b>max</b> keyword or a rate defined in kilobits-per-second to follow.	
<b>Values</b>	0 — 40000000, max
<b>max</b> — removes any existing rate limit imposed by the port scheduler policy for the priority level allowing it to use as much total bandwidth as possible.	

## access

<b>Syntax</b>	<b>access</b>
<b>Context</b>	config>port>ethernet
<b>Description</b>	This command configures Ethernet access port parameters.

## egress

<b>Syntax</b>	<b>egress</b>
<b>Context</b>	config>port>ethernet>access
<b>Description</b>	This command configures Ethernet access egress port parameters.

## vport

<b>Syntax</b>	<b>vport</b> <i>name</i> [ <b>create</b> ] <b>no vport</b> <i>name</i>
<b>Context</b>	config>port>ethernet>access>egress
<b>Description</b>	<p>This command configures a scheduling node, referred to as virtual port, within the context of an egress Ethernet port. The vport scheduler operates either like a port scheduler with the difference that multiple vport objects can be configured on the egress context of an Ethernet port, or it can be an aggregate rate when an egress port-scheduler policy is applied to the port.</p> <p>The vport is always configured at the port level even when a port is a member of a LAG.</p> <p>When a port scheduler policy is applied to a vport the following command is used:</p> <p><b>configure&gt;port&gt;ethernet&gt;access&gt;egress&gt;vport&gt;port-scheduler-policy</b> <i>port-scheduler-policy-name</i></p>

The CLI will not allow the user to apply a port scheduler policy to a vport if one has been applied to the port. Conversely, the CLI will not allow the user to apply a port scheduler policy to the egress of an Ethernet port if one has been applied to any vport defined on the access egress context of this port. The `agg-rate-limit`, along with an egress port-scheduler, can be used to ensure that a given vport does not oversubscribe the port's rate.

SAP and subscriber host queues can be port-parented to a vport scheduler in a similar way they port-parent to a port scheduler or can be port-parented directly to the egress port-scheduler if the `agg-rate-limit` is used.

When the vport uses an aggregate rate, the following command is used:

**configure>port>ethernet>access>egress>vport>agg-rate-limit**

**Parameters**     *name* — Specifies the name of the vport scheduling node and can be up to 32 ASCII characters in length. This does not need to be unique within the system but is unique within the port or a LAG.

### agg-rate-limit

**Syntax**     **agg-rate-limit *agg-rate***  
**no agg-rate-limit**

**Context**     configure>port>ethernet>access>egress>vport

**Description**     This command configures an aggregate rate for the vport. This command is mutually exclusive with the `port-scheduler-policy` command.

**Parameters**     *agg-rate* — Specifies the rate limit for the vport.

**Values**     **max**, 1 — 10000000

### egress-rate-modify

**Syntax**     **[no] egress-rate-modify**

**Context**     configure>port>ethernet>access>egress>vport

**Description**     This command is used to apply HQoS Adjustment to a vport. HQoS Adjustment refers to the dynamic adjustment of the rate limit at an QoS enforcement point within 7x50 when the multicast traffic stream is disjointed from the unicast traffic stream. This QoS enforcement point within 7x50 represents the physical point further down in the access part of the network where the two streams join each other and potentially can cause congestion.

An example would be a PON port which is shared amongst subscriber's multicast traffic (single copy of each channel) and subscriber's unicast traffic. The bandwidth control point for this PON port resides in the upstream 7x50 BNG node in the form of a vport. In case that the multicast delivery method in the 7x50 BNG utilizes redirection, the multicast traffic in the 7x50 BNG will flow outside of the subscriber or the vport context and thus will bypass any bandwidth enforcement in 7x50. To correct this, a vport bandwidth adjustment is necessary in 7x50 that will account for the multicast bandwidth consumption that is bypassing vport in 7x50 but is present in the PON port whose bandwidth is controlled by vport.

An estimate of the multicast bandwidth consumption on the PON port can be made at the vport level based on the IGMP messages sourced from the subscribers behind the PON port. This process is called HQoS Adjustment.

A multicast channel bandwidth is subtracted from or added to the vport rate limit according to the received IGMP Join/Leave messages and the channel bandwidth definition policy associated with the vport (indirectly through a group-interface). Since the multicast traffic on the PON port is shared amongst subscribers behind this PON port, only the first IGMP Join or the last IGMP Leave per multicast channel is tracked for the purpose of the vport bandwidth modification.

The vport rate that will be affected by this functionality depends on the configuration:

- In case the `agg-rate-limit` within the vport is configured, its value will be modified based on the IGMP activity associated with the subscriber under this vport.
- In case the `port-scheduler-policy` within the vport is referenced, the `max-rate` defined in the corresponding `port-scheduler-policy` will be modified based on the IGMP activity associated with the subscriber under this vport.

The channel bandwidth definition policy is defined in the `mcac` policy in the **config-ure>router>mcac>policy** context. The policy is applied under the group-interface or in case of redirection under the redirected-interface.

The rates in effect can be displayed with the following two commands:

```
show port 1/1/5 vport name
```

```
qos scheduler-hierarchy port port-id vport vport-name
```

The configuration of a scheduler policy under a Vport, which is only applicable to Ethernet interfaces, is mutually exclusive with the configuration of the `egress-rate-modify` parameter.

The configuration of a scheduler policy under a Vport, which is only applicable to Ethernet interfaces, is mutually exclusive with the configuration of the **egress-rate-modify** parameter.

**Context** HQoS Adjustment for vport is disabled.

## host-match

<b>Syntax</b>	<b>host-match dest</b> <i>destination-string</i> [ <b>create</b> ] <b>no host-match dest</b> <i>destination-string</i>
<b>Context</b>	config>port>ethernet>access>egr>qgrp
<b>Description</b>	This command configures host matching for the Ethernet port egress queue-group. The no form of the command removes
<b>Parameters</b>	<b>dest</b> <i>destination-string</i> — Specify a host match destination string up to 32 characters in length. <b>create</b> — Keyword used to create the host match. The <b>create</b> keyword requirement can be enabled/disabled in the <b>environment&gt;create</b> context.

## port-scheduler-policy

**Syntax** **port-scheduler-policy** *port-scheduler-policy-name*

**no port-scheduler-policy**

<b>Context</b>	config>port>ethernet>access>egress>vport
<b>Description</b>	<p>This command specifies the destination and organization strings to be used for matching subscriber hosts with this vport.</p> <p>The parent vport of a subscriber host queue, which has the port-parent option enabled, is determined by matching the destination string dest string associated with the subscriber and the organization string org string associated with the subscriber host with the strings defined under a vport on the port associated with the subscriber.</p> <p>If a given subscriber host queue does not have the port-parent option enabled, it will be foster-parented to the vport used by this subscriber and which is based on matching the dest string and org string. If the subscriber could not be matched with a vport on the egress port, the host queue will not be bandwidth controlled and will compete for bandwidth directly based on its own PIR and CIR parameters.</p> <p>By default, a subscriber host queue with the port-parent option enabled is scheduled within the context of the port's port scheduler policy.</p> <p>The <b>agg-rate rate</b>, <b>port-scheduler-policy</b> and <b>scheduler-policy</b> commands are mutually exclusive. Changing between the use of a scheduler policy and the use of an agg-rate/port-scheduler-policy involves removing the existing command and applying the new command. Applying a scheduler-policy to a VPORT is only applicable to Ethernet interfaces.</p> <p>The <b>no</b> form of the command removes the port-scheduler-policy-name from the configuration.</p> <p>The <b>agg-rate rate</b>, <b>port-scheduler-policy</b> and <b>scheduler-policy</b> commands are mutually exclusive. Changing between the use of a scheduler policy and the use of an <b>agg-rate/port-scheduler-policy</b> involves removing the existing command and applying the new command.</p>
<b>Parameters</b>	<i>port-scheduler-policy-name</i> — Specifies an existing port-scheduler-policy configured in the <b>config&gt;qos</b> context.

## scheduler-policy

<b>Syntax</b>	<b>scheduler-policy</b> <i>scheduler-policy-name</i> <b>no scheduler-policy</b>
<b>Context</b>	config>port>ethernet>access>egress>vport
<b>Description</b>	<p>This command specifies a scheduler policy to associate to the Vport. Scheduler policies are configured in the <b>configure&gt;qos&gt;scheduler&gt;policy</b> context. Each scheduler policy is divided up into groups of schedulers based on the tier each scheduler is created under. A tier is used to give structure to the schedulers within a policy and define rules for parent scheduler associations. The policy defines the hierarchy and operating parameters for virtual schedulers.</p> <p>The <b>no</b> form of this command removes the configured egress scheduler policy from the VPORT.</p> <p>The <b>agg-rate rate</b>, <b>port-scheduler-policy</b> and <b>scheduler-policy</b> commands are mutually exclusive. Changing between the use of a scheduler policy and the use of an agg-rate/port-scheduler-policy involves removing the existing command and applying the new command.</p> <p>The configuration of a scheduler policy under a Vport is mutually exclusive with the configuration of the egress-rate-modify parameter.</p>

**Parameters** *scheduler-policy-name* — The *scheduler-policy-name* parameter applies an existing scheduler policy that was created in the **config>qos>scheduler-policy** *scheduler-policy-name* context to create the hierarchy of egress virtual schedulers.

## parent-location

**Syntax** **parent-location {default | sla}**  
**no parent-location**

**Context** config>qos>sap-egress

**Description** This command determines the expected location of the parent schedulers for queues configured with a parent command within the SAP egress policy. All parent schedulers must be configured within a scheduler policy applied at the location corresponding to the parent-location parameter.

If a parent scheduler name does not exist at the specified location, the queue will not be parented and will be orphaned.

**Default** parent-location default

**Parameters** **default** — When the SAP egress policy is applied to an SLA profile for a subscriber, the parent schedulers of the queues need to be configured in the scheduler policy applied to the subscriber's SUB profile.

When the SAP egress policy is applied to a SAP, the parent schedulers of the queues need to be configured in the scheduler policy applied to the SAP or the multi-service site.

**sla** — When the SAP egress policy is applied to an SLA profile for a subscriber, the parent schedulers of the queues need to be configured in the scheduler policy applied to the same SLA profile.

If this parameter is configured within a SAP egress policy that is applied to any object except of the egress of an SLA profile, the configured parent schedulers will not be found and so the queues will not be parented and will be orphaned.

## parent-location

**Syntax** **parent-location {none | sub | vport}**  
**no parent-location**

**Context** config>qos>scheduler-policy

**Description** This command determines the expected location of the parent schedulers for the tier 1 schedulers configured with a parent command within the scheduler policy. The parent schedulers must be configured within a scheduler policy applied at the location corresponding to the parent location parameter.

If a parent scheduler name does not exist at the specified location, the schedulers will not be parented and will be orphaned.

The configuration of parent-location and frame-based-accounting in a scheduler policy is mutually exclusive in to ensure consistency between the different scheduling levels.

**Default** parent-location none

- Parameters**
- none** — This parameter indicates that the tier 1 schedulers do not have a parent scheduler and the configuration of the parent under a tier 1 scheduler is blocked. Conversely, this parameter is blocked when any tier 1 scheduler has a parent configured.
  - sub** — When the scheduler policy is applied to an SLA profile for a subscriber, the parent schedulers of the tier 1 schedulers need to be configured in the scheduler policy applied to the subscriber's SUB profile.  
If this parameter is configured within a scheduler policy that is applied to any object except for the egress of an SLA profile, the configured parent schedulers will not be found and so the tier 1 schedulers will not be parented and will be orphaned.
  - vport** — When the scheduler policy is applied to an SLA profile, a SUB profile for a subscriber or to the egress of a PW SAP, the parent schedulers of the tier 1 schedulers need to be configured in the scheduler policy applied to the VPORT to which the subscriber will be assigned.  
If this parameter is configured within a scheduler policy that is applied to to any object except for the egress of an SLA profile or SUB profile, or to the egress of a PW SAP, the configured parent schedulers will not be found and so the tier 1 schedulers will not be parented and will be orphaned.



---

## MLD Policy Commands

### mld-policy

<b>Syntax</b>	<b>mld-policy</b> <i>mld-policy-name</i> [ <b>create</b> ] <b>no mld-policy</b> <i>mld-policy-name</i>
<b>Context</b>	config>subscr-mgmt
<b>Description</b>	This command enables the context to create an MLD policy.

### egress-rate-modify

<b>Syntax</b>	<b>egress-rate-modify</b> <b>agg-rate-limit</b> <b>egress-rate-modify</b> <b>scheduler</b> <i>scheduler-name</i> <b>no egress-rate-modify</b>
<b>Context</b>	config>subscr-mgmt>mld-policy
<b>Description</b>	This command configures the egress rate modification. The <b>no</b> form of the command removes the values from the configuration.
<b>Parameters</b>	<b>agg-rate-limit</b> — specifies that the maximum total rate for all subscriber egress queues for each subscriber associated with the policy. <b>scheduler</b> <i>scheduler-name</i> — specifies the scheduler to be applied for egress rate modification.

### fast-leave

<b>Syntax</b>	<b>[no] fast-leave</b>
<b>Context</b>	config>subscr-mgmt>mld-policy
<b>Description</b>	This command enables fast leave. When fast leave processing is enabled, the router will immediately remove a SAP or SDP from the IP multicast group when it detects an MLD 'leave' on that SAP or SDP. Fast leave processing allows the switch to remove a SAP or SDP that sends a 'leave' from the forwarding table without first sending out group-specific queries to the SAP or SDP, and thus speeds up the process of changing channels ('zapping'). Fast leave should only be enabled when there is a single receiver present on the SAP or SDP. When fast leave is enabled, the configured last-member-query-interval value is ignored.
<b>Default</b>	no fast-leave

## import

<b>Syntax</b>	<b>import</b> <i>policy-name</i> <b>no import</b>
<b>Context</b>	config>subscr-mgmt>mld-policy
<b>Description</b>	This command specifies the import routing policy to be used. Only a single policy can be imported at a time.  The <b>no</b> form of the command removes the policy association.
<b>Default</b>	<b>no import</b> — No import policy is specified.
<b>Parameters</b>	<i>policy-name</i> — The import policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes. Routing policies are configured in the config>router>policy-options context. The router policy must be defined before it can be imported.

## max-num-groups

<b>Syntax</b>	<b>max-num-groups</b> <i>count</i> <b>no max-num-groups</b>
<b>Context</b>	config>subscr-mgmt>mld-policy
<b>Description</b>	This command defines the maximum number of multicast groups that can be joined. If the router receives a join message that would exceed the configured number of groups, the request is ignored.
<b>Default</b>	no max-num-groups
<b>Parameters</b>	<i>count</i> — Specifies the maximum number of groups that can be joined.
<b>Values</b>	1 — 1000

## max-num-grp-sources

<b>Syntax</b>	<b>max-num-grp-sources</b> [1..32000] <b>no max-num-grp-sources</b>
<b>Context</b>	config>subscr-mgmt>mld-policy
<b>Description</b>	This command configures the maximum number of group sources for which MLD can have local receiver information based on received MLD reports on this interface. When this configuration is changed dynamically to a value lower than currently accepted number of group sources, the group sources that are already accepted are not deleted. Only new group sources will not be allowed. When this object has a value of 0, there is no limit to the number of group sources.  The <b>no</b> form of the command removes the value from the configuration.
<b>Default</b>	no max-num-grp-sources

**Parameters**    **1..32000** — Specifies the maximum number of multicast sources allowed to be tracked per group

## max-num-sources

**Syntax**    **max-num-sources** *max-num-sources*  
**no max-num-sources**

**Context**    config>subscr-mgmt>mld-policy

**Description**    This command configures the maximum number of multicast sources allowed per group.  
 The **no** form of the command removes the value from the configuration.

**Parameters**    *max-num-sources* — Specifies the maximum number of multicast sources allowed per group.

**Values**    1 — 1000

## per-host-replication

**Syntax**    **[no] per-host-replication**

**Context**    config>subscr-mgmt>mld-policy

**Description**    This command enables per-host-replication. In the per-host-replication mode, multicast traffic is replicated per each host within the subscriber irrespective of the fact that some hosts may be subscribed to the same multicast stream. As a result, in case that multiple hosts within the subscriber are registered for the same multicast group, the multicast streams of that group will be generated. The destination MAC address of multicast streams will be changed to unicast so that each host receives its own copy of the stream. Multicast traffic in the per-host-replication mode can be classified via the existing QoS CLI structure. As such the multicast traffic will flow through the subscriber queues. HQoS Adjustment is not needed in this case.

The alternative behavior for multicast replication in IPoE environment is per-SAP- replication. In this model, only a single copy of the multicast stream is sent per SAP, irrespective of the number of hosts that are subscribed to the same multicast group. This behavior applies to 1:1 connectivity model as well as on 1:N connectivity model (SAP centric behavior as opposed to subscriber centric behavior).

In the per-SAP-replication model the destination MAC address is multicast (as opposed to unicast in the per-host-replication model). Multicast traffic is flowing via the SAP queue which is outside of the subscriber context. The consequence is that multicast traffic is not accounted in the subscriber HQoS. In addition, HQoS Adaptation is not supported in the per SAP replication model.

**Default**    disabled

## redirection-policy

**Syntax**    **redirection-policy** *policy-name*  
**no redirection-policy**

**Context**    config>subscr-mgmt>mld-policy

<b>Description</b>	<p>This command will apply multicast redirection action to the subscriber. The redirection action along with the redirected interface (and possibly service id) is defined in the referenced policy-name. MLD messages will be redirected to an alternate interface if that alternate interface has MLD enabled. The alternate interface does not have to have any multicast groups registered via MLD. Currently all MLD messages are redirected and there is no ability to selectively redirect MLD messages based on match conditions (multicast-group address, source IP address, etc.). Multicast redirection is supported between VPRN services and also between interfaces within the Global Routing Context. Multicast Redirection is not supported between the VPRN services and the Global Routing Table (GRT).</p> <p>MLD state is maintained per subscriber host and per redirected interface. Traffic is however forwarded only on the redirected interface.</p>
<b>Default</b>	none
<b>Parameters</b>	<i>policy-name</i> — This is a regular policy defined under the <b>configure&gt;router&gt;policy-option&gt;policy-statement</b> context.

## static

<b>Syntax</b>	<b>static</b>
<b>Context</b>	config>subscr-mgmt>mld-policy
<b>Description</b>	This command adds an MLD static group membership.

## group

<b>Syntax</b>	<b>[no] group grp-ipv6-address</b>
<b>Context</b>	config>subscr-mgmt>mld-policy>static
<b>Description</b>	This command configures a static multicast group.
<b>Parameters</b>	<i>grp-ipv6-address</i> — Specifies the IPv6 address.
<b>Values</b>	<p>&lt;grp-ipv6-address&gt; : ipv6-address - x:x:x:x:x:x:x (eight 16-bit pieces)</p> <p>x:x:x:x:x:d.d.d</p> <p>x - [0..FFFF]H</p> <p>d - [0..255]D</p> <p>- multicast group IPv6 address</p>

## source

<b>Syntax</b>	<b>[no] source ipv6-address</b>
<b>Context</b>	config>subscr-mgmt>mld-policy>static>group
<b>Description</b>	This command adds or removes a static multicast source.

<b>Parameters</b>	<i>grp-ipv6-address</i> — Specifies the IPv6 address.
<b>Values</b>	<grp-ipv6-address> : ipv6-address - 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 - multicast group IPv6 address

## starg

<b>Syntax</b>	<b>[no] starg</b>
<b>Context</b>	config>subscr-mgmt>mld-policy>static>group
<b>Description</b>	This command adds a static (*,G) entry. This command can only be enabled if no existing source addresses for this group are specified.  Use the <b>no</b> form of the command to remove the starg entry from the configuration.
<b>Default</b>	none

## version

<b>Syntax</b>	<b>version</b> <i>version</i> <b>no version</b>
<b>Context</b>	config>subscr-mgmt>mld-policy#
<b>Description</b>	This command configures the MLD version.
<b>Parameters</b>	<i>version</i> — <b>Values</b> 1, 2

---

## IPoE Session Commands

### ipoe-session-policy

<b>Syntax</b>	<b>ipoe-session-policy</b> <i>policy-name</i> [create] <b>no ipoe-session-policy</b> <i>policy-name</i>
<b>Context</b>	config>subscr-mgmt
<b>Description</b>	This command configures an IPoE session policy. The policies are referenced from subscriber interfaces, group interfaces and capture SAPs. Multiple IPoE session policies can be configured.
<b>Default</b>	none
<b>Parameters</b>	<i>policy-name</i> — Specifies the IPoE policy name up to 32 characters in length.

### description

<b>Syntax</b>	<b>description</b> <i>description-string</i> <b>no description</b>
<b>Context</b>	config>subscr-mgmt>ipoe-policy
<b>Description</b>	This command creates a text description stored in the configuration file for a configuration context.  The <b>description</b> command associates a text string with a configuration context to help identify the context in the configuration file.  The <b>no</b> form of this command removes any description string from the context.
<b>Default</b>	no description
<b>Parameters</b>	<i>description-string</i> — A text string describing the entity. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters excluding double quotes. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

### session-key

<b>Syntax</b>	<b>session-key sap mac [cid] [rid]</b> <b>no session-key</b>
<b>Context</b>	config>subscr-mgmt>ipoe-policy
<b>Description</b>	This command configures the key to logically group subscriber hosts that belong to the same dual stack end device in an IPoE session.  The SAP and MAC address are always part of the IPoE session key. Optionally the Circuit-Id/Interface-Id or Remote-Id can be added to the session key.

<b>Default</b>	session-key sap mac
<b>Parameters</b>	<p><b>sap</b> — Includes the SAP as part of the IPoE session key. The <b>sap</b> parameter is mandatory and cannot be removed from the key.</p> <p><b>mac</b> — Includes the MAC address as part of the IPoE session key. The <b>mac</b> parameter is mandatory and cannot be removed from the key.</p> <p><b>cid</b> — Optionally adds the DHCPv4 Relay Agent Circuit-Id (option 82, sub option 1) and DHCPv6 Interface-Id (option 18) field to the IPoE session key.</p> <p><b>rid</b> — Optionally adds the DHCPv4 Relay Agent Remote-Id (option 82, sub option 2) and DHCPv6 Remote-Id (option 37) field to the IPoE session key. For DHCPv6, the enterprise number is excluded from the key.</p> <p><b>NOTE:</b> <b>sap</b> and <b>mac</b> are mandatory parameters while <b>cid</b> and <b>rid</b> are optional and mutually exclusive. Valid IPoE session key parameters are: <b>sap mac</b>, <b>sap mac cid</b> and <b>sap mac rid</b>.</p>

## session-timeout

<b>Syntax</b>	<b>session-timeout</b> <i>timeout</i> <b>no session-timeout</b>
<b>Context</b>	config>subscr-mgmt>ipoe-policy
<b>Description</b>	This command defines the time in seconds between 1 second and 360 days before the IPoE session will be disconnected. The default value is unlimited session timeout.
<b>Default</b>	no session-timeout
<b>Parameters</b>	<i>timeout</i> — Specifies the session timeout in seconds.
<b>Values</b>	1 — 31104000

## ipoe-session

<b>Syntax</b>	<b>[no] ipoe-session</b>
<b>Context</b>	config>service>vpls>sap config>service>ies>sub-if>grp-if config>service>vprn>sub-if>grp-if config>service>ies>sub-if config>service>vprn>sub-if
<b>Description</b>	This command configures IPoE session parameters.
<b>Default</b>	none

## force-auth

<b>Syntax</b>	<b>force-auth [cid-change] [rid-change]</b> <b>force-auth disabled</b> <b>no force-auth</b>
<b>Context</b>	config>service>ies>sub-if>grp-if>ipoe-session config>service>vprn>sub-if>grp-if>ipoe-session
<b>Description</b>	By default, if the circuit-id/interface-id or remote-id in the IPoE session re-authentication trigger packet (such as a DHCP renewal) is not empty and different from the circuit-id/interface-id or remote-id stored in the IPoE session data, a forced re-authentication is performed, ignoring the configured min-auth-interval. This default behavior can be changed with the force-auth command.  The <b>no</b> form of the command, resets the default behavior.
<b>Default</b>	force-auth cid-change rid-change
<b>Parameters</b>	<b>cid-change</b> — Perform a forced re-authentication upon a circuit-id/interface-id change. An empty circuit-id/interface-id is not considered a change.  <b>rid-change</b> — Perform a forced re-authentication upon a remote-id change. an empty remote-id is not considered a change. For DHCPv6, the enterprise number is excluded from the comparison.  <b>disabled</b> — Does not perform a forced re-authentication upon a circuit-id/interface-id or remote-id change.

## ipoe-session-policy

<b>Syntax</b>	<b>ipoe-session-policy <i>policy-name</i></b> <b>no ipoe-session-policy</b>
<b>Default</b>	config>service>vpls>sap> ipoe-session config>service>ies>sub-if>grp-if>ipoe-session config>service>vprn>sub-if>grp-if>ipoe-session
<b>Description</b>	This command specifies the IPoE session policy applicable for this group interface or capture SAP.
<b>Default</b>	no ipoe-session-policy
<b>Parameters</b>	<i>policy-name</i> — Specifies the IPoE session policy name up to 32 characters in length

## min-auth-interval

<b>Syntax</b>	<b>min-auth-interval [days <i>days</i>] [hrs <i>hours</i>] [min <i>minutes</i>] [sec <i>seconds</i>]</b> <b>min-auth-interval infinite</b> <b>no min-auth-interval</b>
<b>Context</b>	config>service>ies>sub-if>grp-if>ipoe-session config>service>vprn>sub-if>grp-if>ipoe-session
<b>Description</b>	Re-authentication for IPoE sessions enable dynamic policy changes.



This command configures the maximum frequency of re-authentications by specifying a minimum interval between two non-forced authentications for the same IPoE session.

A forced authentication is by default triggered by a Circuit-Id/Interface-Id or Remote-Id change (see the [force-auth](#) command).

Re-authentications are by default disabled and can be enabled by configuring a min-auth-interval.

Setting the min-auth-interval to zero seconds will always re-authenticate on each trigger packet.

<b>Default</b>	min-auth-interval infinite
<b>Parameters</b>	<p><b>days</b> — Specifies the min number of days between two non-forced authentications for IPoE sessions</p> <p><b>Values</b> 0 — 365</p> <p><b>hrs</b> — Specifies the min number of hours between two non-forced authentications for IPoE sessions</p> <p><b>Values</b> 0 — 23</p> <p><b>min</b> — Specifies the min number of minutes between two non-forced authentications for IPoE sessions</p> <p><b>Values</b> 0 — 59</p> <p><b>sec</b> — Specifies the min number of seconds between two non-forced authentications for IPoE sessions</p> <p><b>Values</b> 0 — 59</p> <p><b>infinite</b> — Does not perform non-forced re-authentications for IPoE sessions (default).</p>

## sap-session-limit

<b>Syntax</b>	<b>sap-session-limit</b> <i>sap-session-limit</i> <b>no sap-session-limit</b>
<b>Context</b>	config>service>ies>sub-if>grp-if>ipoe-session config>service>vprn>sub-if>grp-if>ipoe-session
<b>Description</b>	This command specifies the number of IPoE sessions per SAP allowed for this group-interface
<b>Default</b>	sap-session-limit 1
<b>Parameters</b>	<p><i>sap-session-limit</i> — Specifies the number of allowed IPoE sessions.</p> <p><b>Values</b> 1 — 65535</p> <p>Note that the operational maximum value may be smaller due to equipped hardware dependencies.</p>

## session-limit

<b>Syntax</b>	<b>session-limit</b> <i>session-limit</i> <b>no session-limit</b>
---------------	--

<b>Context</b>	config>service>ies>sub-if>grp-if>ipoe-session config>service>vprn>sub-if>grp-if>ipoe-session config>service>ies>sub-if>ipoe-session config>service>vprn>sub-if>ipoe-session
<b>Description</b>	This command specifies the number of IPoE sessions allowed for this group interface or retail subscriber interface.
<b>Default</b>	session-limit 1
<b>Parameters</b>	<i>session-limit</i> — Specifies the number of allowed IPoE sessions.
<b>Values</b>	1 – 65535 1 – 262143 (retail subscriber interface) The operational maximum value may be smaller due to equipped hardware dependencies.

## user-db

<b>Syntax</b>	<b>user-db</b> <i>local-user-db-name</i> <b>no user-db</b>
<b>Context</b>	config>service>vpls>sap> ipoe-session config>service>ies>sub-if>grp-if>ipoe-session config>service>vprn>sub-if>grp-if>ipoe-session
<b>Description</b>	This command configures the local user database to use for IPoE session authentication. When configured on a capture SAP, the group interface must have the same local user database configured.
<b>Default</b>	no user-db
<b>Parameters</b>	<i>local-user-db-name</i> — Specifies the local user database name up to 32 characters in length.

## shutdown

<b>Syntax</b>	<b>[no] shutdown</b>
<b>Context</b>	config>service>vpls>sap> ipoe-session config>service>ies>sub-if>grp-if>ipoe-session config>service>vprn>sub-if>grp-if>ipoe-session
<b>Description</b>	The <b>shutdown</b> command enables or disables IPoE session management on a group-interface or capture SAP. A shutdown of the IPoE session CLI hierarchy on a group-interface will clear all active IPoE sessions on that interface, resulting in a deletion of all corresponding subscriber hosts.
<b>Default</b>	shutdown

## Show Commands

### radius-configuration

<b>Syntax</b>	<b>radius-configuration</b>
<b>Context</b>	show>aaa
<b>Description</b>	This command displays RADIUS configuration information.

#### Sample Output

```
# show aaa radius-configuration
=====
RADIUS configuration
=====
CoA Port                : 3799
=====
```

### ancp-policy

<b>Syntax</b>	<b>ancp-policy [policy-name]</b> <b>ancp-policy policy-name association</b>
<b>Context</b>	show>subscr-mgmt
<b>Description</b>	This command displays subscriber Access Node Control Protocol (ANCP) policy information.
<b>Parameters</b>	<i>policy-name</i> — Displays information for the specified ANCP policy. <b>association</b> — Displays the information configured with the ANCP policy.

#### Sample Output

```
A:cses-E11>config>subscr-mgmt>ancp# show subscriber-mgmt ancp-policy "test"
=====
ANCP Policy "test"
=====
I. Rate Reduction      : 0 kbps
I. Rate Adjustment    : 100 percent
I. Rate Monitor        : 63360 kbps
I. Rate Monitor Alarm  : Yes
I. Rate Modify         : N/A

E. Rate Reduction      : 0 kbps
E. Rate Adjustment    : 100 percent
E. Rate Monitor        : 0 kbps
E. Rate Monitor Alarm  : no
E. Rate Modify         : N/A
```

Port Down : N/A

Last Mgmt Change: 02/13/2013 19:15:28

=====

\*A:cses-E11>config>subscr-mgmt>ancp#

## ancp-string

Syntax	ancp-string ancp-string <i>ancp-string</i> ancp-string customer <i>customer-id</i> site <i>customer-site-name</i> ancp-string sap <i>sap-id</i>																																																																													
Context	show>subscr-mgmt																																																																													
Description	This command displays subscriber Access Node Control Protocol (ANCP) string information.																																																																													
Parameters	<i>ancp-string</i> — Specifies an Access Node Control Protocol (ANCP) string up to 63 characters in length.  <b>customer</b> <i>customer-id</i> — Specifies an existing customer ID.  <b>Values</b> 1..2147483647  <b>site</b> <i>customer-site-name</i> — Specifies an existing customer site name up to 32 characters in length.  <b>sap</b> <i>sap-id</i> — Displays ANCP string information for the specified SAP ID.  <b>Values</b> <table><tr><td>&lt;sap-id&gt;</td><td>null</td><td>&lt;port-id bundle-id bpgrp-id lag-id aps-id&gt;</td></tr><tr><td></td><td>dot1q</td><td>&lt;port-id bundle-id bpgrp-id lag-id aps-id pw-id&gt;:qtag1</td></tr><tr><td></td><td>qinq</td><td>&lt;port-id bundle-id bpgrp-id lag-id pw-id&gt;:qtag1.qtag2</td></tr><tr><td></td><td>atm</td><td>&lt;port-id aps-id&gt;[:vpi/vci vpi vpi1.vpi2 cp.conn-prof-id]</td></tr><tr><td></td><td></td><td>cp keyword</td></tr><tr><td></td><td></td><td>conn-prof-id - [1..8000]</td></tr><tr><td></td><td>frame</td><td>&lt;port-id aps-id&gt;:dlci</td></tr><tr><td></td><td>cisco-hdlc</td><td>slot/mda/port.channel</td></tr><tr><td></td><td>cem</td><td>slot/mda/port.channel</td></tr><tr><td></td><td>ima-grp</td><td>&lt;bundle-id&gt;[:vpi/vci vpi vpi1.vpi2 cp.conn-prof-id]</td></tr><tr><td></td><td></td><td>cp keyword</td></tr><tr><td></td><td></td><td>conn-prof-id [1..8000]</td></tr><tr><td></td><td>port-id</td><td>slot/mda/port[.channel]</td></tr><tr><td></td><td>bundle-id</td><td>bundle-&lt;type&gt;-slot/mda.&lt;bundle-num&gt;</td></tr><tr><td></td><td></td><td>bundle keyword</td></tr><tr><td></td><td></td><td>type ima fr ppp</td></tr><tr><td></td><td></td><td>bundle-num [1..336]</td></tr><tr><td></td><td>bpgrp-id</td><td>bpgrp-&lt;type&gt;-&lt;bpgrp-num&gt;</td></tr><tr><td></td><td></td><td>bpgrp keyword</td></tr><tr><td></td><td></td><td>type ima ppp</td></tr><tr><td></td><td></td><td>bpgrp-num [1..2000]</td></tr><tr><td></td><td>aps-id</td><td>aps-&lt;group-id&gt;[.channel]</td></tr><tr><td></td><td></td><td>aps keyword</td></tr><tr><td></td><td></td><td>group-id [1..64]</td></tr><tr><td></td><td>ccag-id</td><td>ccag-&lt;id&gt;.&lt;path-id&gt;[cc-type]:&lt;cc-id&gt;</td></tr></table>			<sap-id>	null	<port-id bundle-id bpgrp-id lag-id aps-id>		dot1q	<port-id bundle-id bpgrp-id lag-id aps-id pw-id>:qtag1		qinq	<port-id bundle-id bpgrp-id lag-id pw-id>:qtag1.qtag2		atm	<port-id aps-id>[:vpi/vci vpi vpi1.vpi2 cp.conn-prof-id]			cp keyword			conn-prof-id - [1..8000]		frame	<port-id aps-id>:dlci		cisco-hdlc	slot/mda/port.channel		cem	slot/mda/port.channel		ima-grp	<bundle-id>[:vpi/vci vpi vpi1.vpi2 cp.conn-prof-id]			cp keyword			conn-prof-id [1..8000]		port-id	slot/mda/port[.channel]		bundle-id	bundle-<type>-slot/mda.<bundle-num>			bundle keyword			type ima fr ppp			bundle-num [1..336]		bpgrp-id	bpgrp-<type>-<bpgrp-num>			bpgrp keyword			type ima ppp			bpgrp-num [1..2000]		aps-id	aps-<group-id>[.channel]			aps keyword			group-id [1..64]		ccag-id	ccag-<id>.<path-id>[cc-type]:<cc-id>
<sap-id>	null	<port-id bundle-id bpgrp-id lag-id aps-id>																																																																												
	dot1q	<port-id bundle-id bpgrp-id lag-id aps-id pw-id>:qtag1																																																																												
	qinq	<port-id bundle-id bpgrp-id lag-id pw-id>:qtag1.qtag2																																																																												
	atm	<port-id aps-id>[:vpi/vci vpi vpi1.vpi2 cp.conn-prof-id]																																																																												
		cp keyword																																																																												
		conn-prof-id - [1..8000]																																																																												
	frame	<port-id aps-id>:dlci																																																																												
	cisco-hdlc	slot/mda/port.channel																																																																												
	cem	slot/mda/port.channel																																																																												
	ima-grp	<bundle-id>[:vpi/vci vpi vpi1.vpi2 cp.conn-prof-id]																																																																												
		cp keyword																																																																												
		conn-prof-id [1..8000]																																																																												
	port-id	slot/mda/port[.channel]																																																																												
	bundle-id	bundle-<type>-slot/mda.<bundle-num>																																																																												
		bundle keyword																																																																												
		type ima fr ppp																																																																												
		bundle-num [1..336]																																																																												
	bpgrp-id	bpgrp-<type>-<bpgrp-num>																																																																												
		bpgrp keyword																																																																												
		type ima ppp																																																																												
		bpgrp-num [1..2000]																																																																												
	aps-id	aps-<group-id>[.channel]																																																																												
		aps keyword																																																																												
		group-id [1..64]																																																																												
	ccag-id	ccag-<id>.<path-id>[cc-type]:<cc-id>																																																																												

```

ccag      keyword
id        [1..8]
path-id   [a|b]
cc-type   [.sap-net|.net-sap]
cc-id     [0..4094]
eth-tunnel eth-tunnel-<id>[:<eth-tun-sap-id>]
id        [1..1024]
eth-tun-sap-id [0..4094]
lag-id    lag-<id>
lag       keyword
id        [1..800]
pw-id     pw-<id>
pw        keyword
id        [1..10239]
qtag1     [0..4094]
qtag2     [*|0..4094]
vpi       [0..4095] (NNI)
          [0..255] (UNI)
vci       [1|2|5..65535]
dlci      [16..1022]
tunnel-id tunnel-<id>.<private|public>:<tag>
          tunnel keyword
id        [1..16]
tag       [0..4094]

```

### Sample Output

```

show subscriber-mgmt ancp-string "ANCP-0000003-0000001"
=====
ANCP-String "ANCP-0000003-0000001"
=====
Type       : SUB - "4AACAHCU74"
State      : Up                               Ancp Policy: N/A
I. Rate    : 129 kbps                         E. Rate    : 130 kbps
Adj I. Rate: N/A                             Adj E. Rate: N/A
Act I. Rate: N/A                             Act E. Rate: N/A
Service Id : 50 (VPRN)
Group      : linux
Neighbor   : 10.0.0.2:34885
Persist Key: N/A
-----
Actual-Net-Data-Rate-Upstream      : 129 kbits/s
Actual-Net-Data-Rate-Downstream    : 130 kbits/s
Minimum-Net-Data-Rate-Upstream     : 131 kbits/s
Minimum-Net-Data-Rate-Downstream   : 132 kbits/s
Attainable-Net-Data-Rate-Upstream  : 133 kbits/s
Attainable-Net-Data-Rate-Downstream: 134 kbits/s
Maximum-Net-Data-Rate-Upstream     : 135 kbits/s
Maximum-Net-Data-Rate-Downstream   : 136 kbits/s
Minimum-Net-Low-Power-Data-Rate-Upstream : 137 kbits/s
Minimum-Net-Low-Power-Data-Rate-Downstream : 138 kbits/s
Maximum-Interleaving-Delay-Upstream : 139 ms
Actual-Interleaving-Delay-Upstream  : 140 ms
Maximum-Interleaving-Delay-Downstream : 141 ms
Actual-Interleaving-Delay-Downstream : 142 ms

```

```
DSL-Line-State : 2 (IDLE)
Access-Loop-Encapsulation : 16909056 (0x01020300)
=====
```

## authentication

- Syntax** **authentication** *policy-name* **association**  
**authentication** [*policy-name*]  
**authentication** [*policy-name*] **statistics**  
**authentication** **coa-statistics**
- Context** show>subscr-mgmt
- Description** This command displays subscriber management RADIUS authentication policy information and statistics.
- Parameters** *policy-name* — Specifies the subscriber management RADIUS authentication policy name, up to 32 characters, for which information is requested.
- association** — Displays SAP, interface, local user database host, AA and L2TP associations of this policy.
- coa-statistics** — Displays the overall statistics for incoming RADIUS Change of Authorization (CoA) messages and Disconnect Messages. For dropped requests, a counter for different drop reasons is available.
- statistics** — Displays a list of policies with basic statistics (without specifying a policy name) or detailed statistics, including per-server statistics for the specified policy-name. These statistics apply only to the legacy RADIUS server configuration method where the servers are directly configured in the authentication policy.

### Sample Output

```
# show subscriber-mgmt authentication
=====
Authentication Policies
=====
Name                               Description
-----
auth-policy-1                      Radius auth policy - servers
auth-policy-2                      Radius auth policy - radius-server-policy
-----
Number of Authentication Policies : 2
=====

# show subscriber-mgmt authentication "auth-policy-2"
=====
Authentication Policy auth-policy-2
=====
Description      : Radius auth policy - radius-server-policy
Re-authentication : Yes                               Username Format   : MAC Address
PPPoE Access Method : PAP/CHAP                       Username Mac-Format : "aa:"
PPP-Username Oper : None
```

```

PPP-Domain-Name      : N/A
Username Oper        : None
Domain-Name          : N/A
Acct-Stop-On-Fail    :
RADIUS Server Policy : "aaa-server-policy-1"
Fallback Action       : deny
Last Mgmt Change     : 06/24/2013 21:16:50

```

-----  
Include Radius Attributes

```

Remote Id             : Yes          Circuit Id             : Yes
NAS Port Id           : Yes          NAS Identifier         : Yes
PPPoE Service Name    : Yes          DHCP Vendor Class Id  : Yes
Access Loop Options   : Yes          MAC Address           : Yes
NAS Port Prefix       : None         NAS Port Suffix        : None
NAS-Port-Type         : Yes (standard) Acct Session Id       : Host
Calling Station Id    : Yes (sap-string) Called Station Id     : Yes
Tunnel Server Attr    : Yes          DHCP Options           : Yes
NAS Port              : Yes
NAS Port Bits Spec    : *3s*1m*4p*12o*12i

```

-----  
Radius Servers

```

Router                : management + Base Source Address       : N/A
Access Algorithm      : Direct          Retry                 : 3
Timeout (s)           : 5              Hold down time (s)    : 30

```

```

Index IP Address      Port  Pend-Req-Limit Out/Overload time (s) Oper State

```

-----  
No Radius Servers configured.-----  
Accept Radius Attributes-----  
No Matching Entries-----  
Radius Script Policies

```

Access-Request        : "N/A"
Access-Accept         : "N/A"
Change-of-Authorization : "N/A"

```

=====

```

# show subscriber-mgmt authentication "auth-policy-2" association

```

-----  
Authentication Policy auth-policy-2-----  
SAP Associations-----  
No associations found.-----  
Interface Associations

```

Service-Id : 3000 (VPRN)
- If Name : group-int-ws-1-1

```

-----  
Local-User-Db PPPoE Host Associations

```

Local-User-Db : ludb-1
- Host : host-1

```

```
-----
Local-User-Db DHCP Host Associations
-----
```

```
Local-User-Db : ludb-1
- Host : default
-----
```

```
Application Assurance Associations
-----
```

```
No associations found.
=====
```

```
No associated L2TP groups found.
No associated L2TP tunnels found.
```

```
# show subscriber-mgmt authentication statistics
=====
```

```
Authentication Policy Statistics
=====
```

Policy Name	Subscr. Pkts Authenticated	Subscr. Pkts Rejected	Subscr. Pkts Rejected Send Failed
-------------	-------------------------------	--------------------------	---

auth-policy-1	0	0	0
auth-policy-2	0	0	0

```
-----
Number of Authentication Policies : 2
=====
```

```
# show subscriber-mgmt authentication "auth-policy-1" statistics
=====
```

```
Authentication Policy Statistics
=====
```

```
-----
Policy name : auth-policy-1
subscriber packets authenticated : 0
subscriber packets rejected : 0
subscriber packets rejected send failed : 0
-----
```

radius server idx IP-address	requests accepted	requests rejected	requests no reply	requests md5 failed	requests pending	requests send failed
---------------------------------	----------------------	----------------------	----------------------	------------------------	---------------------	-------------------------

1 172.16.1.1	0	0	0	0	0	0
--------------	---	---	---	---	---	---

```
# show subscriber-mgmt authentication coa-statistics
=====
```

```
Radius Notify Statistics Change-Of-Authorization Disconnect-Messages
=====
```

Requests Received	7	10
Requests Accepted	5	6
Requests Rejected	2	4
Requests Dropped	0	0
No Auth Policy found	0	0
Invalid message	0	0
Out of resources	0	0
Authentication failure	0	0



## diameter-application-policy

<b>Syntax</b>	<b>diameter-application-policy</b> [ <i>name</i> ]
<b>Context</b>	show>subscr-mgmt
<b>Description</b>	This command displays Diameter application policy information.

**Sample Output**

```
# show subscriber-mgmt diameter-application-policy

=====
DIAMETER application policies
=====
Name                                Description
-----
diameter-gx-policy-1                Diameter Gx policy
diameter-gy-policy-1                Diameter Gy policy
diameter-nasreq-policy-1            Diameter NASREQ policy
-----
No. of policies: 3
=====

# show subscriber-mgmt diameter-application-policy "diameter-nasreq-policy-1"

=====
DIAMETER application policy "diameter-nasreq-policy-1"
=====
Description                        : Diameter NASREQ policy
Session failover                    : enabled
Failover handling                   : continue
Peer policy                         : diameter-peer-policy-1
Application                         : nasreq
Tx timer (s)                       : 10
Last management change              : 02/28/2015 14:53:49
-----
NASREQ
-----
Include AVP                        : nas-port-id
                                   : nas-port-type
NAS-Port-Id prefix type             : none
NAS-Port-Id suffix type             : user-string
NAS-Port-Id suffix                  : @bng1
NAS-Port-Type type                  : standard

User name format                    : mac
User name operation                  : no-operation
MAC address format                  : aa:
Last management change              : 02/28/2015 14:53:49
=====
Interfaces using diameter-auth-policy "diameter-nasreq-policy-1"
-----
Interface-name                      Service-id Type
-----
group-int-1-1                      1000      IES
-----
No. of interfaces: 1
-----
```

```
VPLS SAP's with diameter-auth-policy "diameter-nasreq-policy-1"
```

```
-----
Service      SAP
-----
```

```
10           1/1/4:*. *
-----
```

```
No. of SAP's: 1
-----
```

## explicit-subscriber-map

<b>Syntax</b>	<b>explicit-subscriber-map</b>
<b>Context</b>	show>subscriber-mgmt
<b>Description</b>	This command displays explicit subscriber mappings.

### Sample Output

```
B:Dut-A>show>subscr-mgmt# explicit-subscriber-map
```

```
=====
Explicit Subscriber Map
=====
```

```
Key                               Sub profile
                                SLA profile
-----
```

```
sub_ident_A_1                    sub_prof80
                                sla_prof80
-----
```

```
Number of Explicit Subscriber Mappings : 1
=====
```

```
B:Dut-A>show>subscr-mgmt#
```

## host-lockout-policy

<b>Syntax</b>	<b>host-lockout-policy</b> <b>host-lockout-policy</b> <i>policy-name</i> <b>association</b> <b>host-lockout-policy</b> <i>policy-name</i> <b>host-lockout-policy</b> <i>policy-name</i> <b>all</b> <b>host-lockout-policy</b> <i>policy-name</i> <b>sap</b> <i>sap-id</i> [ <b>circuit-id</b>   <b>mac</b>   <b>remote-id</b> ]
<b>Context</b>	show>subscriber-mgmt
<b>Description</b>	This command displays host lockout policy information.
<b>Parameters</b>	<i>policy-name</i> — Specifies a specific subscriber Host Lockout policy name up to 32 characters. <b>association</b> — Specifies <b>all</b> — Specifies to display all information for the specified policy ID. <b>sap</b> <i>sap-id</i> — Specifies to display SAP ID information.

**circuit-id** — Specifies to display circuit ID information.

**mac** — Specifies to display MAC address information.

**remote-id** — Specifies to display remote ID information.

### Sample Output

```
*A:cses-E11# show subscriber-mgmt host-lockout-policy
=====
Host Lockout Policies
=====
Lockout Policy                               Last Mgmt Change
  Lockout Time Min                          Lockout Time Max
Description
  Lockout Reset Time                        Max Lockout Hosts
-----
test                                         04/20/2012 19:51:02
  10                                         3600
test
  60                                         100
=====
*A:cses-E11#

*A:cses-E11# show subscriber-mgmt host-lockout-policy "test"
=====
Host Lockout Policy "test"
=====
Description                                test
Last Mgmt Change                          04/20/2012 19:51:02
Lockout time min                          10
Lockout time max                         3600
Lockout reset time                       60
Max lockout hosts                        100
Host key                                 all
=====
*A:cses-E11#
```

## ipoe-session-policy

<b>Syntax</b>	<b>ipoe-session-policy</b> <i>ipoe-session-policy-name</i> [ <b>association</b> ] <b>ipoe-session-policy</b>
<b>Context</b>	show>subscr-mgmt
<b>Description</b>	This command displays IPoE session policy information.
<b>Parameters</b>	<i>ipoe-session-policy-name</i> — Specifies the IPoE session policy name up to 32 characters in length. <b>association</b> — Displays the interface and captures SAPs that reference the IPoE session policy.

**Sample Output**

```
# show subscriber-mgmt ipoe-session-policy "ipoe-policy-1"

=====
IPoE Session Policy "ipoe-policy-1"
=====
Description          : IPoE policy
Last Mgmt Change     : 02/28/2015 11:51:25
Session Key          : sap-mac
Session Timeout      : unlimited
=====

# show subscriber-mgmt ipoe-session-policy "ipoe-policy-1" association

=====
IPoE Session Policy "ipoe-policy-1"
=====
-----
IPoE Interface Associations
-----
Service-Id : 1000 (IES)
- group-int-1-1
Service-Id : 2000 (VPRN)
- group-int-1-1

-----
Capture SAP Associations
-----
Service-Id : 10 (VPLS)
- 1/1/4:*. *
=====
```

**local-user-db**

**Syntax** **local-user-db** *local-user-db-name* **association** [dhcp] [ppp] [12tp] [radius] [pppoe] [dhcp6] [capture-sap] [rtr-solicit] [wpp] [ipoe]  
**local-user-db** *local-user-db-name* **ipoe-all-hosts**  
**local-user-db** *local-user-db-name* **ipoe-host** *ipoe-host-name*  
**local-user-db** *local-user-db-name* **ipoe-unmatched-hosts**  
**local-user-db** [*local-user-db-name*]  
**local-user-db** *local-user-db-name* **pppoe-all-hosts**  
**local-user-db** *local-user-db-name* **pppoe-host** *pppoe-host-name*  
**local-user-db** *local-user-db-name* **pppoe-unmatched-hosts**

**Context** show>subscriber-mgmt

**Description** This command displays local user database information.

**Sample Output**

```
*A:ALA-48>show>subscr-mgmt# local-user-db
=====
Local User Databases
```

```

=====
Name                               Admin Host  Description
                                State Count
-----
database01                         Down    1
database02 Provider001/Class0002 Down    0      This is a long testdescription wi*
test                               Down    2
-----
Number of Local User Databases : 3      Number of Hosts : 3
=====
* indicates that the corresponding row element may have been truncated.

*A:ALA-48>show>subscr-mgmt# local-user-db database01
=====
Local User Database "database01"
=====
Admin State      : Down
Last Mgmt Change : 11/08/2007 12:27:36
Host Count       : 1
DHCP Match Types : circ-id
DHCP CircId Mask Pfx : test
DHCP CircId Mask Sfx : N/A
PPPoE Match Types : N/A
PPPoE CircId Mask Pfx : N/A
PPPoE CircId Mask Sfx : N/A
=====
*A:ALA-48>show>subscr-mgmt#

*A:ALA-48>show>subscr-mgmt# local-user-db database01 dhcp-all-hosts
=====
Local User Database "database01" DHCP hosts
=====
Name                               Admin      Matched objects
                                State
-----
host001                           Down       -
-----
Number of DHCP Hosts : 1
=====

*A:ALA-48>show>subscr-mgmt# local-user-db "database01" dhcp-host host001
=====
DHCP Host "host001"
=====
Admin State      : Down
Last Mgmt Change : 11/08/2007 12:13:42

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

Matched Objects   : N/A

```

## Vport Commands

Address : N/A

### Identification Strings

Subscriber Id : N/A

SLA Profile String : N/A

Sub Profile String : N/A

App Profile String : N/A

ANCP String : N/A

Inter Destination Id: N/A

=====

```

*A:ALA-48>show>subscr-mgmt# local-user-db "database01" dhcp-unmatched-hosts
=====
Local User Database "database01" DHCP unmatched hosts
=====
Name                               Reason          Duplicate Host
-----
host002                           No match        N/A
host003                           Duplicate       host001
host004                           No match        N/A
host005                           Duplicate       host001
-----
Number of DHCP Unmatched Hosts : 4
=====
*A:ALA-48>show>subscr-mgmt#

*A:ALA-48>show>subscr-mgmt# local-user-db "database01" association
=====
DHCP Servers where database01 is used
=====
Server-Name                        Router-Name
-----
dhcpS1                             1000
-----
No. of Server(s): 1
=====
Interfaces where database01 is used for authentication
=====
Interface-Name                     Service-Id Type
-----
No. of Interface(s): 0
=====
*A:ALA-48>show>subscr-mgmt#

*A:ALA-48>show>subscr-mgmt# local-user-db "database01" association dhcp
=====
DHCP Servers where database01 is used
=====
Server-Name                        Router-Name
-----
dhcpS1                             1000
-----
No. of Server(s): 1
=====
*A:ALA-48>show>subscr-mgmt#
=====

# show subscriber-mgmt local-user-db "ludb-1" association ipoe
=====
IPoE client interface associations for ludb-1
=====
Interface-Name                     Svc-Id          Type
-----
group-int-1-1                     1000            IES
group-int-1-1                     2000            VPRN
-----
No. of Interface(s): 2
=====

```

```

=====
Capture SAP associations for luidb-1
=====
SAP                               Svc-Id    Type      PPpOE  PPP  IPoE  DHCP  DHCP6  RS
-----
1/1/4:1202.*                     10       VPLS      Y      Y      Y      Y      Y
1/1/4:*.*                       10       VPLS      Y      Y      Y      Y      Y
-----
No. of SAP(s): 2
=====

```

## msap-policy

**Syntax** **msap-policy** [*msap-policy-name* [**association**]]

**Context** show>subscr-mgmt

**Description** This command displays Managed SAP policy information.

### Sample Output

```

*A:ALA-48>show>subscr-mgmt# msap-policy
=====
Managed SAP Policies
=====
Name                               Num      Description
                                MSAPs
-----
test                               0        (Not Specified)
test 1                             0        (Not Specified)
-----
Number of MSAP Policies : 2
Number of MSAPs         : 0
=====
*A:ALA-48>show>subscr-mgmt#

```

## sla-profile

**Syntax** **sla-profile** [*sla-profile-name* [**association**]]

**Context** show>subscriber-mgmt

**Description** This command displays SLA profile information.

**Parameters** *sla-profile-name* — Specifies an existing SLA profile name.

**association** — Displays the information configured with the specified *sla-profile-name*.

### Sample Output

```

A:Dut-A# show subscriber-mgmt sla-profile
=====

```



```

SLA Profiles
=====
Name                               Description
-----
sla_default
sla_prof100_VOIP
sla_prof110_VOIP
sla_prof120_VOIP
sla_prof130_VOIP
sla_prof140_VOIP
sla_prof230_VOIP
sla_prof80
sla_prof80_VOIP
sla_prof81_VOIP
sla_prof90_VOIP
sla_profPC1
sla_profPC2
sla_profPC3
-----
Number of SLA Profiles : 14
=====
A:Dut-A#

A:Dut-A# show subscriber-mgmt sla-profile sla_prof100_VOIP
=====
SLA Profile sla_prof100_VOIP
=====
Host Limit      : 3 (Remove Oldest)
Ingress Qos-Policy : 100                Egress Qos-Policy : 100
Ingress Queuing Type : Service-queuing
Ingress Filter-Id   : N/A                Egress Filter-Id   : N/A
Last Mgmt Change    : 07/10/2006 12:55:33
-----
Ingress Queue Overrides
-----
Queue Rate      CIR      HiPrio  CBS      MBS
-----
2      4000      -      -      -      -
3      2500      -      -      -      -
-----
Egress Queue Overrides
-----
Queue Rate      CIR      HiPrio  CBS      MBS
-----
2      4000      -      -      -      -
3      2500      -      -      -      -
=====
A:Dut-A#

A:Dut-A# show subscriber-mgmt sla-profile sla_prof100_VOIP association
=====
SLA Profile sla_prof100_VOIP
-----
SAP Default-Profile Associations
-----
No associations found.
-----
SAP Static Host Associations

```

```

-----
No associations found.
-----
SAP Non-Sub-Traffic-Profile Associations
-----
No associations found.
-----
Sub-Ident-Policy Profile Map Associations
-----
Policy-name : sub_ident_all
- Key : sla_prof100_VOIP
-----
Sub-Profile Map Associations
-----
No associations found.
-----
Explicit Subscriber Map Associations
-----
No associations found.
=====
A:Dut-A#

```

## sla-profile

<b>Syntax</b>	<b>subscriber</b> <i>sub-ident-string</i> <b>sla-profile</b> <i>sla-profile-name</i> <b>sap</b> <i>sap-id</i> [ <b>scheduler</b> <i>scheduler-name</i> ]
<b>Context</b>	show>qos>scheduler-stats
<b>Description</b>	This command displays the subscriber's SLA profile scheduler stats.
<b>Parameters</b>	<b>subscriber</b> <i>sub-ident-string</i> — Displays information for the specified subscriber profile name. <b>sla-profile</b> <i>sla-profile-name</i> — Displays information for the specified <i>sla-profile-name</i> . <b>sap</b> <i>sap-id</i> — Displays information for the specified SAP. <b>scheduler</b> <i>scheduler-name</i> — Displays information for the specified <i>scheduler-name</i> .

### Sample Output

```

*A:BNG# show qos scheduler-stats subscriber "sub1" sla-profile "sla-profile.1" sap 1/
1/1:1 scheduler

"session-sched"

=====
Scheduler Stats
=====
Scheduler                                Forwarded Packets      Forwarded Octets
-----
Egress Schedulers

session-sched                            0                      0
=====
*A:BNG#

```

port

<b>Syntax</b>	<b>port</b> <i>port-id</i> <b>vport</b> <i>name</i> [ <b>scheduler</b> <i>scheduler-name</i> ] [ <b>detail</b> ]
<b>Context</b>	show>qos>scheduler-hierarchy
<b>Description</b>	This command displays the subscriber's SLA profile scheduler stats.
<b>Parameters</b>	<b>port</b> <i>port-id</i> — Displays information for the specified port. <b>vport</b> <i>name</i> — Displays information for the specified vport. <b>scheduler</b> <i>scheduler-name</i> — Displays information for the specified <i>scheduler-name</i> . <b>detail</b> — Displays detailed information.

### Sample Output

```
*A:BN# show qos scheduler-hierarchy port 1/1/1 vport "dslam1" scheduler "dslam-sched"
```

```

Scheduler Hierarchy - Port 1/1/1
=====
Scheduler-policy dslam-sched-pol
| slot(1)
|--(S) : subscriber-sched (VPort dslam1 1/1/1)
|
|   |--(S) : session-sched
|   |
|   |   |--(Q) : Sub=sub2:sla-profile.2 200->1/1/1:2->3
|   |   |--(Q) : Sub=sub2:sla-profile.2 200->1/1/1:2->2
|   |   |--(Q) : Sub=sub2:sla-profile.2 200->1/1/1:2->1
|   |
|   |--(S) : session-sched
|   |
|   |   |--(Q) : Sub=sub2:sla-profile.1 200->1/1/1:2->3
|   |   |--(Q) : Sub=sub2:sla-profile.1 200->1/1/1:2->2
|   |   |--(Q) : Sub=sub2:sla-profile.1 200->1/1/1:2->1
|   |
|   |--(S) : subscriber-sched (VPort dslam1 1/1/1)
|   |
|   |   |--(S) : session-sched
|   |   |
|   |   |   |--(Q) : Sub=sub1:sla-profile.2 200->1/1/1:1->3
|   |   |   |--(Q) : Sub=sub1:sla-profile.2 200->1/1/1:1->2
|   |   |   |--(Q) : Sub=sub1:sla-profile.2 200->1/1/1:1->1
|   |   |
|   |   |--(S) : session-sched

```

```

|   |   |   | -- (Q) : Sub=sub1:sla-profile.1 200->1/1/1:1->3
|   |   |   | -- (Q) : Sub=sub1:sla-profile.1 200->1/1/1:1->2
|   |   |   | -- (Q) : Sub=sub1:sla-profile.1 200->1/1/1:1->1
|   |   |   |
=====
*A:BNG#

```

## vport

- Syntax** **port** *port-id* **vport** *vport name* [**scheduler** *scheduler-name*]
- Context** show>qos>scheduler-stats
- Description** This command displays the vport scheduler stats.
- Parameters** **port** *port-id* — Displays information for the specified port.  
**vport** *vport name* — Displays information for the specified vport.  
**scheduler** *scheduler-name* — Displays information for the specified *scheduler-name*.

### Sample Output

```

*A:BNG# show qos scheduler-stats port 1/1/1 vport "dslam1" scheduler "dslam-sched"
=====
Scheduler Stats
=====
Scheduler                               Forwarded Packets      Forwarded Octets
-----
Egress Schedulers

dslam-sched                             0                      0
=====
*A:BNG#

```

## statistics

- Syntax** **statistics iom** (*slot* | **all**) [**host**|**session**|**subscriber**|**summary**] [**non-zero-value-only**]  
**statistics mda** (*mda* | **all**) [**host**|**session**|**subscriber**|**summary**] [**non-zero-value-only**]  
**statistics port** (*port-id* | **all**) [**host**|**session**|**subscriber**|**summary**] [**non-zero-value-only**]  
**statistics pw-port** (*pw-port* | **all**) [**host**|**session**|**subscriber**|**summary**] [**non-zero-value-only**]  
**statistics system** [**host**|**session**|**subscriber**|**summary**] [**non-zero-value-only**]
- Context** show>subscr-mgmt
- Description** This command displays enhanced subscriber management statistics per port/pw-port/MDA/IOM/system.  
For each statistic, there is current value and peak value, peak value is the highest value since last reset via system boot or command **clear subscriber-mgmt peakvalue-stats**.

Note that the peak values can be reset via the **clear subscriber-mgmt peakvalue-stats** command.

- Parameters.**
- iom slot** — Displays specified IOM slot information.
  - mda mda** — Displays specified slot/mda information.
  - port port-id** — Specifies to display information for both the physical port ID and LAG.
  - pw-port pw-port** — Specifies to display information for a pseudowire port ID.
  - Values** 1 — 10239
  - all** — displays statistics of all IOM or MDA or port or pseudowire port in the system.
  - host** — Displays v4/v6 host statistics only.
  - session** — Displays PPPoX/LAC/LNS session statistics only.
  - subscriber** — Displays subscriber statistics only.
  - summary** — Displays summary statistics only.
  - non-zero-value-only** — Displays only non-zero value counters.

The following tables describe the counters available in the **show subscriber management statistics** command output.

The following terminology is used to indicate applicability of the stats:

- ESM — Enhanced Subscriber Management. Subscriber traffic forwarded via subscriber queues. Enabled with SAP sub-sla-mgmt in no shutdown state.
- BSM — Basic Subscriber Management. Subscriber traffic forwarded via SAP queues. SAP sub-sla-mgmt must be in shutdown state. For DHCP, dhcp lease-populate or dhcp6-relay lease-populate must be enabled to count the leases. For IPv4, if anti-spoof is enabled on the SAP, a subscriber host is instantiated.
- Routed CO — IES or VPRN service with subscriber-interface and group-interface constructs.
- Bridged CO — VPLS service with DHCPv4 lease management enabled (lease-populate)
- regular interface — IES or VPRN interface (none subscriber-interface or group-interface)
- Host (also subscriber host) — A resource in the system that is used for traffic forwarding and security related actions. The creation of a subscriber host entry is linked to anti-spoof being enabled on a SAP. For ESM, anti-spoof is mandatory and hence every connected {IP/MAC} consumes by default a subscriber host entry. A DHCP6 IA-PD can also be modeled as a managed route. In this case, no subscriber host is instantiated. For BSM, anti-spoof is optional on regular interfaces. An IPv4 static-host and DHCPv4 lease do not result in a subscriber host instantiation when anti-spoof is disabled on the SAP.

Host and Protocol Statistics			
Section	Counter	Counts	Applies to
IPv4	1. PPP Hosts - IPCP	IPv4 local terminated PPP hosts (PTA, LNS)	ESM, Routed CO
	2. IPOE Hosts - DHCP	DHCPv4 hosts (lease states)	ESM, Routed CO, Bridged CO
	3. IPOE Hosts - ARP	ARP hosts	ESM, Routed CO, Bridged CO

<b>Host and Protocol Statistics (Continued)</b>			
<b>Section</b>	<b>Counter</b>	<b>Counts</b>	<b>Applies to</b>
	4. IPOE Hosts – Static	IPv4 static hosts	ESM, Routed CO, Bridged CO
	5. IPOE Hosts BSM - DHCP	DHCPv4 hosts (lease states: anti-spoof and lease-populate enabled)	BSM, Routed CO, Bridged CO, regular interface
	6. IPOE Hosts BSM – Static	IPv4 static hosts (with anti-spoof enabled)	BSM, Routed CO, Bridged CO, regular interface
	7. IPOE BSM - DHCP	DHCPv4 lease states (with lease-populate enabled, no anti-spoof)	BSM, Routed CO, Bridged CO, regular interface
	8. IPOE BSM – Static	IPv4 static hosts (no anti-spoof)	BSM, Routed CO, Bridged CO, regular interface

Host and Protocol Statistics (Continued)			
Section	Counter	Counts	Applies to
IPv6	9. PPP Hosts – SLAAC	Local terminated IPv6 wan-host – SLAAC (PTA, LNS)	ESM, Routed CO
	10. PPP Hosts - DHCP6 (PD)	Local terminated IPv6 pd-host (PTA, LNS) – DHCP6 IA-PD leases over PPP (excluding PD as managed route)	ESM, Routed CO
	11. PPP Hosts - DHCP6 (NA)	Local terminated IPv6 wan-host (PTA, LNS) – DHCP6 IA-NA leases over PPP	ESM, Routed CO
	12. PPP Mngd Rt - DHCP6 (PD)	IPv6 (PTA, LNS) – DHCP6 IA-PD leases over PPP (PD as managed route only)	ESM, Routed CO
	13. IPOE Hosts – SLAAC	IPv6 wan-host – SLAAC	ESM, Routed CO
	14. IPOE Hosts - DHCP6 (PD)	IPv6 pd-host – DHCP6 IA-PD leases (excluding PD as managed route)	ESM, Routed CO
	15. IPOE Hosts - DHCP6 (NA)	IPv6 wan-host – DHCP6 IA-NA leases	ESM, Routed CO
	16. IPOE Mngd Rt - DHCP6 (PD)	IPv6 – DHCP6 IA-PD leases (PD as managed route only)	ESM, Routed CO
	17. IPOE Hosts – Static (PD)	IPv6 static hosts with prefix-length shorter than /128	ESM, Routed CO
	18. IPOE Hosts – Static (WAN)	IPv6 static hosts with prefix-length equal to /128	ESM, Routed CO
	19. IPOE BSM - DHCP6 (PD)	IPv6 – DHCP6 IA-PD leases (lease-populate)	BSM, regular interface
	20. IPOE BSM - DHCP6 (NA)	IPv6 – DHCP6 IA-NA leases (lease-populate)	BSM, regular interface
Total	21. PPP Hosts	Local terminated PPP hosts (PTA, LNS) Sum of counters 1, 9, 10 and 11	ESM
	22. IPOE Hosts	Total IPv4 and IPv6 IPOE hosts. Sum of counters 2, 3, 4, 5, 6, 13, 14, 15, 17 and 18	ESM
	23. IPv4 Hosts	Total IPv4 hosts. PPP (PTA, LNS) and IPOE. Sum of counters 1, 2, 3, 4, 5 and 6	ESM

Host and Protocol Statistics (Continued)			
Section	Counter	Counts	Applies to
Total (Cont)	24. IPv6 Hosts	Total IPv6 hosts. PPP (PTA, LNS) and IPOE. Sum of counters 9, 10, 11, 13, 14, 15, 17 and 18	ESM
	25. IPv6 PD Mngd Routes	Total DHCP6 IA-PD leases modeled as a managed route. PPP (PTA, LNS) and IPOE. Sum of counters 12 and 16	ESM
	26. L2TP LAC Hosts	L2TP LAC hosts – single host per single or dual stack PPP session. Counter also increases for outgoing LTS sessions.	ESM, Routed CO
	27. Internal Hosts	Subscriber hosts for internal use. For example: LNS redirect hosts (for LTS, an LNS redirect host is also instantiated).	ESM
	28. Non-Sub-Traffic L2-Hosts	Host on a single subscriber SAP in a VPLS service that enables non-IP traffic to be forwarded using the specified SLA profile instance queues. Host on a single subscriber SAP attached to an IES/VRPN group-interface that enables traffic normally forwarded via the SAP queues to flow via the specified SLA profile instance queues. configure service vpls <service-id> sap <sap-id> sub-sla-mgmt single-sub-parameters non-sub-traffic sub-profile <sub-profile-name> sla-profile <sla-profile-name> [subscriber <sub-ident-string>] [app-profile <app-profile-name>]	ESM, Routed CO, Bridged CO
	29. DHCP leases	Total number of DHCPv4 lease states. Sum of counters 2, 5 and 7	ESM, BSM
	30. DHCPv6 leases	Total number of DHCPv6 lease states. Sum of counters 10, 11, 12, 14, 15, 16, 19 and 20	ESM, BSM



Host and Protocol Statistics (Continued)			
Section	Counter	Counts	Applies to
Total (Cont)	31. Subscriber Hosts	Counter displayed in the output of “show subscriber-mgmt statistics iom   mda   port   pw-port” This counter matches the number of hosts accounted for in the per line card limit Sum of counters 1, 2, 3, 4, 5, 6, 9, 10, 11, 13, 14, 15, 17, 18 and 26	ESM
	32. System Hosts Scale	Counter displayed in the output of “show subscriber-mgmt statistics system” This counter matches the number of hosts accounted for in the system wide limit Sum of counters 1, 2, 3, 4, 5, 6, 9, 10, 11, 13, 14, 15, 17, 18, 26 and 27	ESM

PPP Session Statistics			
Section	Counter	Counts	Applies to
Local	33. PPP Sessions - PPPoE	Local terminated PPPoE sessions (PTA)	ESM, Routed CO
	34. PPP Sessions - PPPoEoA	Local terminated PPPoEoA sessions (PTA)	ESM, Routed CO
	35. .PPP Sessions - PPPoA	Local terminated PPPoA sessions (PTA)	ESM, Routed CO
	36. PPP Sessions - L2TP (LNS)	Local terminated PPP sessions (L2TP LNS)	ESM, Routed CO
LAC	37. PPP Sessions - PPPoE	Tunneled PPPoE session (L2TP LAC)	ESM, Routed CO
	38. PPP Sessions - PPPoEoA	Tunneled PPPoEoA session (L2TP LAC)	ESM, Routed CO
	39. PPP Sessions - PPPoA	Tunneled PPPoA session (L2TP LAC)	ESM, Routed CO
	40. PPP Sessions - L2TP (LTS)	Tunneled PPP session (L2TP LTS)	ESM, Routed CO
Total	41. PPP Sessions - established	PPP sessions that are established (at least one active host attached) – PTA/LAC/LTS/LNS	ESM, Routed CO

PPP Session Statistics (Continued)			
Section	Counter	Counts	Applies to
Total (Cont)	42. PPP Sessions - in setup	PPP sessions in setup (session created, host setup in progress) – PTA/LAC/LTS/LNS	ESM, Routed CO
	43. PPP Sessions - local	Local terminated PPPoX sessions (PTA, L2TP LNS) Sum of counters 33, 34, 35 and 36	ESM, Routed CO
	44. PPP Sessions - LAC	Tunneled PPPoX session (L2TP LAC, L2TP LTS) Sum of counters 37, 38, 39 and 40	ESM, Routed CO
L2TP	45. L2TP Tunnels - originator	Number of L2TP Tunnels originated on this node. (LAC/ LTS)	ESM, Routed CO
	46. L2TP Tunnels - receiver	Number of L2TP Tunnels terminated on this node. (LNS/LTS)	ESM, Routed CO
	47. Total L2TP Tunnels	Number of L2TP Tunnels originated or terminated on this node Sum of counters 45 and 46	ESM, Routed CO

IPoE Session Statistics			
Section	Counter	Counts	Applies to
Total	48. IPOE Sessions - established	IPoE sessions that are established (at least one active host attached).	ESM, Routed CO
	49. IPOE Sessions- in setup	IPoE sessions in setup (session created, host setup in progress).	ESM, Routed CO

Subscriber Statistics			
Section	Counter	Counts	Applies to
Total	50. Subscribers	Total number of active subscribers.	ESM, Routed CO, Bridged CO

SubMgmt Statistics Summary		
Section	Counter	Counts
Hosts	IPv4	Total IPv4 hosts (counter 23 in tables above)
	IPv6	Total IPv6 hosts (counter 24 in tables above)
Sessions	PPP	Total PPP sessions - established (counter 41 in tables above)
	IPOE	Total IPOE sessions – established (counter 48 in tables above)
Subscribers		Total number of active subscribers (counter 50 in tables above)

### Sample Output

```

A:PE-1# show subscriber-mgmt statistics system
=====
Subscriber Management Statistics for System
=====
Type                               Current      Peak      Peak Timestamp
-----
Host & Protocol Statistics
-----
IPv4  PPP Hosts      - IPCP                1          1 02/28/2015 16:25:43
      IPOE Hosts    - DHCP                0          2 02/28/2015 12:38:58
      IPOE Hosts    - ARP                 1          1 02/28/2015 13:46:10
      IPOE Hosts    - Static              0          0
      IPOE Hosts BSM - DHCP                0          0
      IPOE Hosts BSM - Static              0          0
      IPOE BSM      - DHCP                0          0
      IPOE BSM      - Static              0          0
-----
IPv6  PPP Hosts      - SLAAC              0          0
      PPP Hosts      - DHCP6 (PD)         0          0
      PPP Hosts      - DHCP6 (NA)         0          0
      PPP Mngd Rt    - DHCP6 (PD)         0          0
      IPOE Hosts      - SLAAC              0          0
      IPOE Hosts      - DHCP6 (PD)         0          0
      IPOE Hosts      - DHCP6 (NA)         0          0
      IPOE Mngd Rt    - DHCP6 (PD)         0          0
      IPOE Hosts      - Static (PD)         0          0
      IPOE Hosts      - Static (WAN)        0          0
      IPOE BSM        - DHCP6 (PD)         0          0
      IPOE BSM        - DHCP6 (NA)         0          0
-----
Total  PPP Hosts      1          1 02/28/2015 16:25:43
      IPOE Hosts      1          2 02/28/2015 12:38:58
      IPv4 Hosts      2          2 02/28/2015 16:25:43
      IPv6 Hosts      0          0
      IPv6 PD Mngd Routes 0          0
      L2TP LAC Hosts  0          0
      Internal Hosts   0          0
      Non-Sub-Traffic L2-Hosts 0          0

```

```

DHCP Leases                                0          2 02/28/2015 12:38:58
DHCPv6 Leases                             0          0
System Hosts Scale                         2          2 02/28/2015 16:25:43
-----
PPP Session Statistics
-----
Local  PPP Sessions - PPPoE                1          1 02/28/2015 16:25:43
      PPP Sessions - PPPoEoA              0          0
      PPP Sessions - PPPoA                0          0
      PPP Sessions - L2TP (LNS)           0          0
-----
LAC    PPP Sessions - PPPoE                0          0
      PPP Sessions - PPPoEoA              0          0
      PPP Sessions - PPPoA                0          0
      PPP Sessions - L2TP (LTS)           0          0
-----
Total  PPP Sessions - established          1          1 02/28/2015 16:25:43
      PPP Sessions - in setup              0          1 02/28/2015 16:25:43
      PPP Sessions - local                 1          1 02/28/2015 16:25:43
      PPP Sessions - LAC                   0          0
-----
L2TP   L2TP Tunnels - originator           0          0
      L2TP Tunnels - receiver             0          0
      Total L2TP Tunnels                   0          0
-----
IPOE Session Statistics
-----
Total  IPOE Sessions - established          0          0
      IPOE Sessions - in setup              0          0
-----
Subscriber Statistics
-----
Total  Subscribers                        2          2 02/28/2015 16:25:43
=====
Peak values last reset at : n/a

```

### Sample Output (summary view)

```
A:PE-1# show subscriber-mgmt statistics port 1/1/4 summary
```

```
=====
SubMgmt Statistics
```

```

=====
Port Id      |      Hosts      |      Sessions   |      Subscribers
              |      IPv4      |      IPv6      |      PPP      IPOE      |
-----
1/1/4        |      2          |      2          |      1          1      |      2      (Curr)
              |      3          |      3          |      1          2      |      3      (Peak)
=====

```

## sub-ident-policy

<b>Syntax</b>	<b>sub-ident-policy</b> [ <i>sub-ident-policy-name</i> [ <b>association</b> ]] <b>sub-ident-policy</b> <i>sub-ident-policy-name</i> <b>script</b> { <b>primary</b>   <b>secondary</b>   <b>tertiary</b> }
<b>Context</b>	show>subscriber-mgmt
<b>Description</b>	This command displays subscriber identification policy information.
<b>Parameters</b>	<i>sub-ident-policy-name</i> — Specifies an existing subscriber identification policy name. <b>association</b> — Displays information configured with the specified <i>sub-ident-policy-name</i> . <b>script</b> { <b>primary</b>   <b>secondary</b>   <b>tertiary</b> } — Displays information for the specified identification script.

## Sample Output

```

B:Dut-A>show>subscr-mgmt# sub-ident-policy
=====
Subscriber Identification Policies
=====
Name                               Description
-----
sub_ident_all
sub_ident_pc
-----
Number of Subscriber Identification Policies : 2
=====
B:Dut-A>show>subscr-mgmt#

B:Dut-A>show>subscr-mgmt# sub-ident-policy sub_ident_all
=====
Subscriber Identification Policy sub_ident_all
=====
Sub Profile Map
-----
Key                               Sub profile
-----
sub_prof100                       sub_prof100
sub_prof110                       sub_prof110
sub_prof120                       sub_prof120
sub_prof130                       sub_prof130
sub_prof140                       sub_prof140
sub_prof230                       sub_prof230
sub_prof80                        sub_prof80
sub_prof81                        sub_prof81
sub_prof90                        sub_prof90
-----
SLA Profile Map
-----
Key                               SLA profile
-----
sla_prof100_VOIP                  sla_prof100_VOIP
sla_prof110_VOIP                  sla_prof110_VOIP
sla_prof120_VOIP                  sla_prof120_VOIP
sla_prof130_VOIP                  sla_prof130_VOIP
sla_prof140_VOIP                  sla_prof140_VOIP

```

```

sla_prof230_VOIP          sla_prof230_VOIP
sla_prof80_VOIP           sla_prof80_VOIP
sla_prof81_VOIP           sla_prof81_VOIP
sla_prof90_VOIP           sla_prof90_VOIP
-----
Python Scripts
-----
#           Admin Oper  Script
           State State Name
-----
Primary    Down  Down  pyTom.py
Secondary  Up    Up    pyTomDebug.py
Tertiary   Up    Up    hardcoded.py
=====
B:Dut-A>show>subscr-mgmt#
B:Dut-A>show>subscr-mgmt# sub-ident-policy sub_ident_all association
=====
Subscriber Identification Policy sub_ident_all
=====
SAP Associations
-----
Service-Id : 80 (VPLS)
- SAP : 1/2/1:80
Service-Id : 90 (VPLS)
- SAP : 1/2/1:90
Service-Id : 100 (VPLS)
- SAP : 1/2/1:100
- SAP : 1/2/1:101
- SAP : 1/2/1:102
Service-Id : 110 (VPLS)
- SAP : 1/2/1:110
- SAP : 1/2/1:111
- SAP : 1/2/1:112
Service-Id : 120 (VPLS)
- SAP : 1/2/1:120
- SAP : 1/2/1:121
- SAP : 1/2/1:122
Service-Id : 130 (VPLS)
- SAP : 1/2/1:130
Service-Id : 140 (VPLS)
- SAP : 1/2/1:140
=====
B:Dut-A>show>subscr-mgmt#

B:Dut-A>show>subscr-mgmt# sub-ident-policy sub_ident_all script primary
=====
Subscriber Identification Policy sub_ident_all
=====
Primary Script
-----
URL           : ftp://xxx:yyy@a.b.c.d/pyTom.py
Admin State   : Down           Oper State : Down
-----
Source (dumped from memory)
-----
Script is not active.
-----
=====
B:Dut-A>show>subscr-mgmt#

```

```

B:Dut-A>show>subscr-mgmt# sub-ident-policy sub_ident_all script secondary
=====
Subscriber Identification Policy sub_ident_all
=====
Secondary Script
-----
URL           : ftp://xxx:yyy@a.b.c.d/pyTomDebug.py
Admin State  : Up                               Oper State : Up
-----
Source (dumped from memory)
-----
1 import alc
2 yiaddr = alc.dhcp.yiaddr
3 # Subscriber ID equals full client IP address.
4 # Note: IP address 10.10.10.10 yields 'sub-168430090'
5 # and not 'sub-10.10.10.10'
6 alc.dhcp.sub_ident = 'sub-' + str(yiaddr)
7 # DHCP server is configured such that the third byte (field) of the IP
8 # address indicates the session Profile ID.
9 alc.dhcp.sla_profile = 'sp-' + str((yiaddr & 0x0000FF00) >> 8)
=====
B:Dut-A>show>subscr-mgmt#

B:Dut-A>show>subscr-mgmt# sub-ident-policy sub_ident_all script tertiary
=====
Subscriber Identification Policy sub_ident_all
=====
Tertiary Script
-----
URL           : ftp://xxx:yyy@a.b.c.d/hardcoded.py
Admin State  : Up                               Oper State : Up
-----
Source (dumped from memory)
-----
1 from alc import dhcp
2
3 dhcp.sub_ident = 'sub_ident_A_1'
4 dhcp.sub_profile_string = 'sub_prof_B_2'
5 dhcp.sla_profile_string = 'sla_prof_C_3'
6
=====
B:Dut-A>show>subscr-mgmt#

```

## sub-profile

<b>Syntax</b>	<b>sub-profile</b> [ <i>sub-profile-name</i> [ <i>association</i> ]]
<b>Context</b>	show>subscriber-mgmt
<b>Description</b>	This command displays subscriber profile information.
<b>Parameters</b>	<i>sub-profile-name</i> — Specifies an existing subscriber profile name. <i>association</i> — Displays the information configured with the specified <i>sub-profile-name</i> .

**Sample Output**

```

A:Dut-A# show subscriber-mgmt sub-profile
=====
Subscriber Profiles
=====
Name                                     Description
-----
sub_default
sub_prof100
sub_prof110
sub_prof120
sub_prof130
sub_prof140
sub_prof230
sub_prof80
sub_prof81
sub_prof90
sub_profPC1
sub_profPC2
sub_profPC3
-----
Number of Subscriber Profiles : 13
=====
A:Dut-A#

A:Dut-A# show subscriber-mgmt sub-profile sub_prof100
=====
Subscriber Profile sub_prof100
=====
I. Sched. Policy : service100
E. Sched. Policy : service100
Acct. Policy      : 1                               Collect Stats : Enabled
Last Mgmt Change : 07/10/2006 12:55:33
-----
Ingress Scheduler Overrides
-----
Scheduler          Rate      CIR
-----
serv100            8000      sum
-----
Egress Scheduler Overrides
-----
Scheduler          Rate      CIR
-----
serv100            8000      sum
-----
SLA Profile Map
-----
Key                SLA Profile
-----
No mappings configured.
=====
A:Dut-A#

A:Dut-A# show subscriber-mgmt sub-profile sub_prof100 association
=====
Subscriber Profile sub_prof100
-----

```



```

SAP Default-Profile Associations
-----
No associations found.
-----
SAP Static Host Associations
-----
No associations found.
-----
SAP Non-Sub-Traffic-Profile Associations
-----
No associations found.
-----
Sub-Ident-Policy Profile Map Associations
-----
Policy-name : sub_ident_all
- Key : sub_prof100
-----
Explicit Subscriber Map Associations
-----
No associations found.
=====
A:Dut-A#

```

## pw-port

- Syntax** **pw-port** [*pw-port-id*] [**detail**]  
**pw-port sdp** *sdp-id*  
**pw-port sdp none**
- Context** show>pw-port
- Description** Displays pseudo-wire port information.  
 If no optional parameters are specified, the command displays a summary of all defined PW ports. The optional parameters restrict output to only ports matching the specified properties.
- Parameters** *pw-port-id* — Specifies the pseudo-wire port identifier.  
**Values** 1 — 10239  
**detail** — Displays detailed port information that includes all the **pw-port** output fields.  
**sdp** *sdp-id* — The SDP ID for which to display matching PW port information.  
**Values** 1 — 17407
- Output** **Show PW-Port** — The following table describes **show pw-port** output fields:

Label	Description
PW Port	The PW Port identifier.
Encap	The encapsulation type of the PW Port.
SDP	The SDP identifier.
IfIndex	The interface index used for the PW Port.

Label	Description (Continued)
VC-Id	The Virtual Circuit identifier.
Description	The description string for the PW Port.

### Sample Output

```
*A:ALA-48>config>service# show pw-port
```

```
=====
PW Port Information
=====
PW Port   Encap      SDP      IfIndex    VC-Id
-----
1         dot1q      1        1526726657 1
2         qinq      1        1526726658 2
3         dot1q      1        1526726659 3
4         qinq      1        1526726660 4
=====
```

```
*A:ALA-48>config>service# show pw-port 3
```

```
=====
PW Port Information
=====
PW Port   Encap      SDP      IfIndex    VC-Id
-----
3         dot1q      1        1526726659 3
=====
```

```
*A:ALA-48>config>service# show pw-port 3 detail
```

```
=====
PW Port Information
=====
PW Port      : 3
Encap        : dot1q
SDP          : 1
IfIndex      : 1526726659
VC-Id        : 3
Description  : 1-Gig Ethernet dual fiber
=====
```

```
*A:ALA-48>config>pw-port$ show pw-port sdp none
```

```
=====
PW Port Information
=====
PW Port   Encap      SDP      IfIndex    VC-Id
-----
5         dot1q      1        1526726661
=====
```

```
*A:ALA-48>config>pw-port$ show pw-port sdp 1
```

```
=====
PW Port Information
=====
PW Port   Encap      SDP      IfIndex    VC-Id
-----
```

1	dot1q	1	1526726657	1
2	qinq	1	1526726658	2
3	dot1q	1	1526726659	3
4	qinq	1	1526726660	4
=====				

## port-scheduler-policy

**Syntax** **port-scheduler-policy** [*port-scheduler-policy-name*] [*association*]  
**port-scheduler-policy** *port-scheduler-policy-name* **network-policy** *network-queue-policy-name*  
**port-scheduler-policy** *port-scheduler-policy-name* **sap-egress** *policy-id*  
**port-scheduler-policy** *port-scheduler-policy-name* **scheduler-policy** *scheduler-policy-name*  
**port-scheduler-policy** *port-scheduler-policy-name* **scheduler-policy** *scheduler-policy-name* **sap-egress** *policy-id*

**Context** show>qos

**Description** This command displays scheduler policy information.

### Sample Output

```
A:NS072860910>config>qos>port-sched-plcy# info
-----
max-rate 10000
group "group1" create
    rate 3000 cir 1000
exit
group "group2" create
    rate 2000 cir 500
exit
level 7 rate 7000 cir 700 group "group1" weight 3
level 6 rate 6000 cir 600 group "group1" weight 2
level 5 rate 5000 cir 500 group "group1" weight 1
level 2 rate 2000 cir 200 group "group2" weight 2
level 1 rate 1000 cir 100 group "group2" weight 1
-----

A:NS072860910# show qos scheduler-hierarchy port 5/1/2 vport "fred"
=====
Scheduler Hierarchy - Port 5/1/2, Vport "fred"
=====
Port-scheduler-policy pspl
    Port Bandwidth : 1000000    Max Rate : 10000
    Consumed : 0                Offered : 0

[Within CIR Level 8]
    Rate      : max
    Consumed : 0                Offered : 0

    (Q) : 1->5/1/2:1->1
    (Q) : 1->5/1/2:2->1

[Within CIR Group "group1"]
```

```
Rate      : 1000
Consumed  : 0           Offered : 0

[Within CIR Level 7]
  Weight   : 3
  Rate     : 700
  Consumed : 0           Offered : 0

  (Q) : 1->5/1/2:1->2
  (Q) : 1->5/1/2:2->2

[Within CIR Level 6]
  Weight   : 2
  Rate     : 600
  Consumed : 0           Offered : 0

  (Q) : 1->5/1/2:1->3
  (Q) : 1->5/1/2:2->3

[Within CIR Level 5]
  Weight   : 1
  Rate     : 500
  Consumed : 0           Offered : 0

  (Q) : 1->5/1/2:1->4
  (Q) : 1->5/1/2:2->4

[Within CIR Level 4]
  Rate     : max
  Consumed : 0           Offered : 0

[Within CIR Level 3]
  Rate     : max
  Consumed : 0           Offered : 0

  (Q) : 1->5/1/2:1->5
  (Q) : 1->5/1/2:2->5

[Within CIR Group "group2"]
  Rate : 500
  Consumed : 0           Offered : 0

[Within CIR Level 2]
  Weight   : 2
  Rate     : 200
  Consumed : 0           Offered : 0

  (Q) : 1->5/1/2:1->6
  (Q) : 1->5/1/2:2->6

[Within CIR Level 1]
  Weight   : 1
  Rate     : 200
  Consumed : 0           Offered : 0

  (Q) : 1->5/1/2:1->7
  (Q) : 1->5/1/2:2->7

[Within CIR Level 0]
  Rate     : 0
  Consumed : 0           Offered : 0
```

```

(Q) : 1->5/1/2:1->8
(Q) : 1->5/1/2:2->8

[Above CIR Level 8]
Rate      : max
Consumed  : 0           Offered : 0

[Above CIR Group "group1"]
Rate      : 3000
Consumed  : 0           Offered : 0

[Above CIR Level 7]
Weight    : 3
Rate      : 7000
Consumed  : 0           Offered : 0

[Above CIR Level 6]
Weight    : 2
Rate      : 6000
Consumed  : 0           Offered : 0

[Above CIR Level 5]
Weight    : 1
Rate      : 5000
Consumed  : 0           Offered : 0

[Above CIR Level 4]
Rate      : max
Consumed  : 0           Offered : 0

[Above CIR Level 3]
Rate      : max
Consumed  : 0           Offered : 0

[Above CIR Group "group2"]
Rate      : 2000
Consumed  : 0           Offered : 0

[Above CIR Level 2]
Weight    : 2
Rate      : 2000
Consumed  : 0           Offered : 0

[Above CIR Level 1]
Weight    : 1
Rate      : 1000
Consumed  : 0           Offered : 0

(Q) : 1->5/1/2:1->1
(Q) : 1->5/1/2:1->2
(Q) : 1->5/1/2:1->3
(Q) : 1->5/1/2:1->4
(Q) : 1->5/1/2:1->5
(Q) : 1->5/1/2:1->6
(Q) : 1->5/1/2:1->7
(Q) : 1->5/1/2:1->8
(Q) : 1->5/1/2:2->1
(Q) : 1->5/1/2:2->2
(Q) : 1->5/1/2:2->3
(Q) : 1->5/1/2:2->4

```

```

(Q) : 1->5/1/2:2->5
(Q) : 1->5/1/2:2->6
(Q) : 1->5/1/2:2->7
(Q) : 1->5/1/2:2->8
=====
A:NS072860910#

*A:B-Dut-A>config>qos>port-sched-plcy# show qos port-scheduler-policy "psp"
=====
QoS Port Scheduler Policy
=====
Policy-Name      : psp
Description      : (Not Specified)
Max Rate         : max                      Last changed      : 04/15/2010 00:37:02
Group           : 1
Group PIR        : 80000                    Group CIR         : max

Group           : 2
Group PIR        : 80000                    Group CIR         : max

Group           : 3
Group PIR        : 80000                    Group CIR         : max

Group           : 4
Group PIR        : 80000                    Group CIR         : max

Lvl1 PIR         : max                      Lvl1 CIR          : max
Lvl1 Group       : 1                      Lvl1 Grp Weight   : 10

Lvl2 PIR         : max                      Lvl2 CIR          : max
Lvl2 Group       : 1                      Lvl2 Grp Weight   : 20

Lvl3 PIR         : max                      Lvl3 CIR          : max
Lvl3 Group       : 2                      Lvl3 Grp Weight   : 30

Lvl4 PIR         : max                      Lvl4 CIR          : max
Lvl4 Group       : 2                      Lvl4 Grp Weight   : 40

Lvl5 PIR         : max                      Lvl5 CIR          : max
Lvl5 Group       : 3                      Lvl5 Grp Weight   : 50

Lvl6 PIR         : max                      Lvl6 CIR          : max
Lvl6 Group       : 3                      Lvl6 Grp Weight   : 60

Lvl7 PIR         : max                      Lvl7 CIR          : max
Lvl7 Group       : 4                      Lvl7 Grp Weight   : 70

Lvl8 PIR         : max                      Lvl8 CIR          : max
Lvl8 Group       : 4                      Lvl8 Grp Weight   : 80

Orphan Lvl       : default                  Orphan Weight     : default
Orphan CIR-Lvl   : default                  Orphan CIR-Weight : default
=====
*A:Bennet-Dut-A>config>qos>port-sched-plcy#

*A:B-Dut-A# show qos port-scheduler-policy "psp" association
=====
QoS Port Scheduler Policy
=====

```

```

Policy-Name      : psp
Description      : (Not Specified)
-----
Associations
-----
- Port : 1/1/2 VPort : vp1
=====
*A:B-Dut-A#

*A:B-Dut-A# show qos port-scheduler-policy "psp" sap-egress 1000
=====
Compatibility : Port-scheduler Policy psp & Sap Egress Queue 1000
=====
Orphan Queues :

None Found

Hierarchy      :

Root
|
|--- (Q) : 1
|
|--- (Q) : 2
|
|--- (Q) : 3
|
|--- (Q) : 4
|
|--- (Q) : 5
|
|--- (Q) : 6
|
|--- (Q) : 7
|
|--- (Q) : 8
=====
*A:B-Dut-A#

```

## sap-egress

<b>Syntax</b>	<b>sap-egress</b> [ <i>policy-id</i> ] [ <b>association</b>   <b>detail</b> ]
<b>Context</b>	show>qos
<b>Description</b>	This command displays SAP egress policy information.
<b>Parameters</b>	<p><i>policy-id</i> — Displays information for the specified SAP egress policy.</p> <p><b>association</b> — Displays the information configured with the specified <i>sap-egress</i> policy.</p> <p><b>detail</b> — Displays detailed information.</p>

**Sample Output**

```

*A:Dut-A# show qos sap-egress
=====
Sap Egress Policies
=====
Policy-Id  Scope      Name      Description
-----
1          Template  default   Default SAP egress QoS policy.
30         Template
31         Template   1 video channel, 1 BE
80         Template   1 video EF, 2xvideo AF, 1 BE
100        Template   Limit outgoing 80
110        Template   User100
120        Template   User110
130        Template   User120
140        Template   User130
140        Template   User140
901        Template   User90_1
902        Template   User90_2
903        Template   User90_3
904        Template   User90_4
905        Template   User90_5
1000       Template   Service all
-----
Number of Policies : 15
=====
*A:Dut-A#

A:Dut-A# show qos sap-egress 31 detail
=====
QoS Sap Egress
=====
Sap Scheduler Policy (31)
-----
Policy-id      : 31                      Scope      : Template
Description    : 1 video EF, 2xvideo AF, 1 BE

-----
Queue          CIR Admin PIR Admin CBS      HiPrio  PIR Lvl/Wt  Parent
              CIR Rule  PIR Rule  MBS
-----
1              0          max      def      def      1/1        limit_8000
              closest closest  def      0/1
2              0          max      def      def      2/1        limit_8000
              closest closest  def      0/1
3              0          max      def      def      3/1        limit_8000
              closest closest  def      0/1
-----
FC Name          Queue-id  Explicit/Default
-----
be               1          Explicit (0)
af               2          Explicit (0)
ef               3          Explicit (0)
-----
Associations
-----
Service-Id      : 23 (VPLS)              Customer-Id   : 1
- SAP : 1/2/2:4000
Service-Id      : 30 (VPLS)              Customer-Id   : 2

```



```

- SAP : lag-1
- SAP : lag-2:5
Service-Id      : 31 (VPLS)                Customer-Id   : 2
- SAP : 1/2/1:31
SLA Profiles :
- sla_profPC1                override
-----
Mirror SAPs
-----
No Mirror SAPs Found.
=====
A:Dut-A#

```

## sap-ingress

<b>Syntax</b>	<b>sap-ingress</b> [ <i>policy-id</i> ] [ <b>association</b>   <b>match-criteria</b>   <b>detail</b> ]
<b>Context</b>	show>qos
<b>Description</b>	This command displays SAP ingress policy information.
<b>Parameters</b>	<p><i>policy-id</i> — Displays information for the specified SAP ingress policy.</p> <p><b>association</b> — Displays the information configured with the specified <i>sap-ingress</i> policy.</p> <p><b>match-criteria</b> — Displays information about the matching criteria.</p> <p><b>detail</b> — Displays detailed information.</p>

### Sample Output

```

A:Dut-A# show qos sap-ingress
=====
Sap Ingress Policies
=====
Policy-Id  Scope      Name      Description
-----
1          Template  default   Default SAP ingress QoS policy.
80         Template
90         Template
100        Template
110        Template
120        Template
130        Template
140        Template
901        Template
902        Template
903        Template
904        Template
905        Template
1000       Template
          User90_1
          User90_2
          User90_3
          User90_4
          User90_5
          Dot1p mappings/service for all s*
-----
Number of Policies : 14
-----
* indicates that the corresponding row element may have been truncated.
A:Dut-A#

```

A:Dut-A# show qos sap-ingress 80 detail

=====

QoS Sap Ingress

=====

Sap Ingress Policy (80)

-----

Policy-id	: 80	Scope	: Template
Default FC	: be	Priority	: Low
Criteria-type	: IP		
Description	: Dot1p mappings/service for service 80		

-----

Queue	Mode	CIR Admin	PIR Admin	CBS	HiPrio	PIR Lvl/Wt	Parent
		CIR Rule	PIR Rule	MBS		CIR Lvl/Wt	
1	Prio	0	7000	def	def	1/1	serv80
		closest	closest	def		0/1	
2	Prio	0	3500	def	def	2/1	serv80
		closest	closest	def		0/1	
3	Prio	0	2000	def	def	3/1	serv80
		closest	closest	def		0/1	
11	Prio	0	max	def	def	1/1	None
		closest	closest	def		0/1	

-----

FC	UCastQ	MCastQ	BCastQ	UnknownQ
be	1	def	def	def
af	2	def	def	def
ef	3	def	def	def

-----

SubFC	Profile	In-Remark	Out-Remark
af	None	None	None
be	None	None	None
ef	None	None	None

-----

Dot1p	FC	Priority
0	be	Default
2	af	Default
5	ef	Default

-----

DSCP	FC	Priority
------	----	----------

-----

No DSCP-Map Entries Found.

-----

Prec Value	FC	Priority
------------	----	----------

-----

No Prec-Map Entries Found.

-----

Match Criteria

-----

IP Match Criteria

-----

Entry	: 1		
Source IP	: Undefined	Source Port	: None
Dest. IP	: Undefined	Dest. Port	: None
Protocol	: None	DSCP	: None
Fragment	: Off		
FC	: Default	Priority	: Default

-----

```

Associations
-----
Service-Id      : 80 (VPLS)                      Customer-Id   : 80
- SAP : 1/2/1:80

SLA Profiles :
- sla_prof80                      override
- sla_prof80_VOIP                 override
- sla_prof81_VOIP                 override
=====
A:Dut-A#

```

## scheduler-hierarchy

<b>Syntax</b>	<b>scheduler-hierarchy</b>
<b>Context</b>	show>qos
<b>Description</b>	This command enables the context to display information about policies that use this scheduler.

## customer

<b>Syntax</b>	<b>customer</b> <i>customer-id</i> <b>site</b> <i>customer-site-name</i> [ <b>scheduler</b> <i>scheduler-name</i> ] [ <b>ingress egress</b> ] [ <b>detail</b> ]
<b>Context</b>	show>qos>scheduler-hierarchy show>qos>scheduler-stats
<b>Description</b>	This command displays the scheduler hierarchy per customer multi-service-site.
<b>Parameters</b>	<b>customer</b> <i>customer-id</i> — Displays information for the specified customer ID. <b>site</b> <i>customer-site-name</i> — Displays information for the specified multi-service <i>customer-site-name</i> . <b>scheduler</b> <i>scheduler-name</i> — Displays information for the specified scheduler-name. <b>ingress</b> — Displays information for the ingress policy. <b>egress</b> — Displays information for the egress policy. <b>detail</b> — Displays detailed information.

## sap

<b>Syntax</b>	<b>sap</b> <i>sap-id</i> [ <b>scheduler</b> <i>scheduler-name</i> ] [ <b>ingress egress</b> ] [ <b>detail</b> ]
<b>Context</b>	show>qos>scheduler-hierarchy show>qos>scheduler-stats
<b>Description</b>	This command displays the scheduler stats per SAP.
<b>Parameters</b>	<i>sap-id</i> — Specifies the physical port identifier portion of the SAP definition. See <a href="#">Common Service Commands on page 2062</a> for <i>sap-id</i> command syntax.

**scheduler** *scheduler-name* — Displays information for the specified scheduler-name.

**ingress** — Displays information for the ingress policy.

**egress** — Displays information for the egress policy.

**detail** — Displays detailed information.

## subscriber

**Syntax** **subscriber** *sub-ident-string* [**scheduler** *scheduler-name*] [**ingress|egress**] [**detail**]  
**subscriber** *sub-ident-string* **sla-profile** *sla-profile-name* **sap** *sap-id* [**scheduler** *scheduler-name*] [**detail**]

**Context** show>qos>scheduler-hierarchy

**Description** This command displays the scheduler hierarchy rooted at the SLA profile scheduler.

**Parameters** **subscriber** *sub-ident-string* — Displays information for the specified subscriber profile name.  
**sla-profile** *sla-profile-name* — Displays information for the specified sla-profile-name.  
**sap** *sap-id* — Displays information for the specified SAP.  
**scheduler** *scheduler-name* — Displays information for the specified scheduler-name.  
**detail** — Displays detailed information.

Note that if the SLA profile scheduler is orphaned (that is when the scheduler has a parent which does not exist) then the hierarchy is only shown when the show command includes the sla-profile and sap parameters.

### Sample Output

```
*A:BNG# show qos scheduler-hierarchy subscriber "sub1" sla-profile "sla-profile.1"
sap 1/1/1:1 scheduler "session-sched"
=====
Scheduler Hierarchy - Subscriber sub1 SLA-Profile sla-profile.1 SAP 1/1/1:1
=====
Egress Scheduler Policy : session-sched-pol
-----
session-sched (Egr)
| slot(1)
|--(Q) : Sub=sub1:sla-profile.1 200->1/1/1:1->3
|
|--(Q) : Sub=sub1:sla-profile.1 200->1/1/1:1->2
|
|--(Q) : Sub=sub1:sla-profile.1 200->1/1/1:1->1
|

B:Dut-A# show qos scheduler-hierarchy subscriber alcatel_100 scheduler serv_all
=====
Scheduler Hierarchy - Subscriber alcatel_100
=====
serv_all (Ing)
| slot(1)
```

```

--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:101->11 MCast
--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:100->11 MCast
--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:102->11 MCast
--(S) : AccessIngress:Sub=6:1 100->1/2/1:100->2
|
|   --(Q) : Sub=alcatel_100:sla_default 100->1/2/1:100->2 1/1
|   --(Q) : Sub=alcatel_100:sla_default 100->1/2/1:100->2 3/2
|   --(Q) : Sub=alcatel_100:sla_default 100->1/2/1:100->2 1/2
|
--(S) : AccessIngress:Sub=6:1 100->1/2/1:100->1
|
|   --(Q) : Sub=alcatel_100:sla_default 100->1/2/1:100->1 1/1
|   --(Q) : Sub=alcatel_100:sla_default 100->1/2/1:100->1 3/2
|   --(Q) : Sub=alcatel_100:sla_default 100->1/2/1:100->1 1/2
|
--(S) : AccessIngress:Sub=6:1 100->1/2/1:100->3
|
|   --(Q) : Sub=alcatel_100:sla_default 100->1/2/1:100->3 1/1
|   --(Q) : Sub=alcatel_100:sla_default 100->1/2/1:100->3 3/2
|   --(Q) : Sub=alcatel_100:sla_default 100->1/2/1:100->3 1/2
|
--(S) : AccessIngress:Sub=6:1 100->1/2/1:102->1
|
|   --(Q) : Sub=alcatel_100:sla_default 100->1/2/1:102->1 1/1
|   --(Q) : Sub=alcatel_100:sla_default 100->1/2/1:102->1 3/2
|   --(Q) : Sub=alcatel_100:sla_default 100->1/2/1:102->1 1/2
|
--(S) : AccessIngress:Sub=6:1 100->1/2/1:102->2
|
|   --(Q) : Sub=alcatel_100:sla_default 100->1/2/1:102->2 1/1
|   --(Q) : Sub=alcatel_100:sla_default 100->1/2/1:102->2 3/2
|   --(Q) : Sub=alcatel_100:sla_default 100->1/2/1:102->2 1/2
|
...

```

```

=====
B:Dut-A#

```

```

B:Dut-A# show qos scheduler-hierarchy subscriber alcatel_100 scheduler serv_all
detail

```

```

=====
Scheduler Hierarchy - Subscriber alcatel_100
=====

```

```

Legend :

```

(U) - Unrestricted (P) - Provisioned  
 (A) - Administrative (O) - Operational  
 MIR - Measured Info Rate

```
-----
serv_all (Ing)
| slot(1)
|--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:101->11 MCast
|
|   PIR Lvl:4          PIR Wt :1
|   CIR Lvl:0          CIR Wt :1
|
|   MIR :0
|   PIR (P):0          PIR (U):7000
|   CIR (P):0          CIR (U):0
|
|   PIR (A):1000000    PIR (O):7000
|   CIR (A):0          CIR (O):0
|   CBS :0             MBS :1280
|   Depth :0           Hi Prio:256
|
|--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:102->11 MCast
|
|   PIR Lvl:4          PIR Wt :1
|   CIR Lvl:0          CIR Wt :1
|
|   MIR :0
|   PIR (P):0          PIR (U):7000
|   CIR (P):0          CIR (U):0
|
|   PIR (A):1000000    PIR (O):7000
|   CIR (A):0          CIR (O):0
|   CBS :0             MBS :1280
|   Depth :0           Hi Prio:256
|
|--(S) : AccessIngress:Sub=6:1 100->1/2/1:102->1
|
|   PIR Lvl:1          PIR Wt :1
|   CIR Lvl:0          CIR Wt :1
|
|   MIR :1687
|   PIR (P):1690       PIR (U):3510
|   CIR (P):0          CIR (U):0
|
|   PIR (A):7000
|   CIR (A):0
|
|--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:102->1 1/1
|
|   PIR Lvl:1          PIR Wt :1
|   CIR Lvl:1          CIR Wt :1
|
|   MIR :0
|   PIR (P):0          PIR (U):1830
|   CIR (P):0          CIR (U):0
|
|   PIR (A):7000       PIR (O):1850
|   CIR (A):0          CIR (O):0
|   CBS :0             MBS :64
|   Depth :0           Hi Prio:8
|
...

```

```

--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:102->3
|
|
|   PIR Lvl:3           PIR Wt :1
|   CIR Lvl:0           CIR Wt :1
|
|   MIR      :0
|   PIR (P):0           PIR (U):2000
|   CIR (P):0           CIR (U):0
|
|   PIR (A):2000        PIR (O):2000
|   CIR (A):0           CIR (O):0
|   CBS      :0         MBS      :64
|   Depth   :0         Hi Prio:8
|
=====
B:Dut-A#

```

## scheduler-name

<b>Syntax</b>	<b>scheduler-name</b> <i>scheduler-name</i>
<b>Context</b>	show>qos
<b>Description</b>	This command displays information about the specified scheduler name.
<b>Parameters</b>	<i>scheduler-name</i> — Displays information about the specified scheduler.

## scheduler-policy

<b>Syntax</b>	<b>scheduler-policy</b> [ <i>scheduler-policy-name</i> ] [ <b>association</b>   <b>sap-ingress</b> <i>policy-id</i>   <b>sap-egress</b> <i>policy-id</i> ]
<b>Context</b>	show>qos
<b>Description</b>	This command displays information about the specified scheduler policy.
<b>Parameters</b>	<i>scheduler-policy-name</i> — Displays information for the specified scheduler policy. <b>sap-ingress</b> <i>policy-id</i> — Displays information for the ingress policy. <b>sap-egress</b> <i>policy-id</i> — Displays information for the egress policy. <b>association</b> — Displays the information currently configured with the specified <i>scheduler-policy-name</i> .

### Sample Output

```

B:Dut-A# show qos scheduler-policy
=====
Sap Scheduler Policies
=====
Policy-Id           Description
-----
maximum_4000_1xEF_1xBE
maximum_8000_1xEF_2xAF_1xBE

```

```

multiservice-site
root
scheduler-7Mbps
service100
service110
service120
service130
service140
service80
service90
service_all
=====
B:Dut-A#

B:Dut-A# show qos scheduler-policy root association
=====
QoS Scheduler Policy
=====
Policy-Name      : root
-----
Associations
-----
No Association Found.
=====
B:Dut-A#

B:Dut-A# show qos scheduler-policy association
=====
QoS Scheduler Policy
=====
Policy-Name      : maximum_4000_1xEF_1xBE
-----
Associations
-----
No Association Found.

Policy-Name      : maximum_8000_1xEF_2xAF_1xBE
-----
Associations
-----
Service-Id       : 23 (VPLS)                Customer-Id      : 1
- SAP : 1/3/2:4000 (Egr)
Service-Id       : 30 (VPLS)                Customer-Id      : 2
- SAP : lag-1 (Egr)
- SAP : lag-2:5 (Egr)
Policy-Name      : multiservice-site
-----
Associations
-----
Service-Id       : 90 (VPLS)                Customer-Id      : 90
- SAP : 1/1/12:95 (Ing) (Egr) MSS : sitel
- SAP : 1/1/20:94 (Ing) (Egr) MSS : sitel

- Customer : 2          MSS : sitel (Ing) (Egr)
- Customer : 90         MSS : sitel (Ing) (Egr)

Policy-Name      : root
-----
Associations

```



-----  
 No Association Found.

Policy-Name : scheduler-7Mbps  
 -----

Associations  
 -----

No Association Found.

Policy-Name : service100  
 -----

Associations  
 -----

Service-Id : 100 (VPLS) Customer-Id : 100  
 - SAP : 1/2/1:100 (Ing) (Egr)  
 - SAP : 1/2/1:101 (Ing) (Egr)  
 - SAP : 1/2/1:102 (Ing) (Egr)  
  
 - Customer : 100 MSS : site100 (Ing) (Egr)

Sub Profiles :  
 - sub\_prof100 (Ing) (Egr)

Policy-Name : service110  
 -----

Associations  
 -----

Service-Id : 110 (VPLS) Customer-Id : 110  
 - SAP : 1/2/1:110 (Ing) (Egr)  
 - SAP : 1/2/1:111 (Ing) (Egr)  
 - SAP : 1/2/1:112 (Ing) (Egr)

Sub Profiles :  
 - sub\_prof110 (Ing) (Egr)

Policy-Name : service120  
 -----

Associations  
 -----

Service-Id : 120 (VPLS) Customer-Id : 120  
 - SAP : 1/2/1:120 (Ing) (Egr)  
 - SAP : 1/2/1:121 (Ing) (Egr)  
 - SAP : 1/2/1:122 (Ing) (Egr)

Sub Profiles :  
 - sub\_prof120 (Ing) (Egr)

Policy-Name : service130  
 -----

Associations  
 -----

Service-Id : 130 (VPLS) Customer-Id : 130  
 - SAP : 1/2/1:130 (Ing) (Egr)

Sub Profiles :  
 - sub\_prof130 (Ing) (Egr)

Policy-Name : service140  
 -----

Associations  
 -----

```

Service-Id      : 140 (VPLS)                      Customer-Id   : 140
- SAP : 1/2/1:140 (Ing) (Egr)

Sub Profiles :
- sub_prof140 (Ing) (Egr)

Policy-Name     : service80
-----
Associations
-----
Service-Id      : 80 (VPLS)                      Customer-Id   : 80
- SAP : 1/2/1:80 (Ing) (Egr)

- Customer : 80      MSS : site80 (Ing) (Egr)

Sub Profiles :
- sub_prof80 (Ing) (Egr)
- sub_prof81 (Ing) (Egr)

Policy-Name     : service90
-----
Associations
-----
Service-Id      : 90 (VPLS)                      Customer-Id   : 90
- SAP : 1/2/1:90 (Ing) (Egr)

Sub Profiles :
- sub_prof90 (Ing) (Egr)

Policy-Name     : service_all
-----
Associations
-----
Sub Profiles :
- sub_default (Ing) (Egr)
=====
B:Dut-A#

```

## scheduler-stats

<b>Syntax</b>	<b>scheduler-stats</b>
<b>Context</b>	show>qos
<b>Description</b>	This command enables the context to display scheduler statistics information.

### Sample Output

```

A:Dut-A# show qos scheduler-stats subscriber alcatel_100
=====
Scheduler Stats
=====
Scheduler                                Forwarded Packets    Forwarded Octets
-----
Ingress Schedulers
root                                     112777              25218126
serv_all                               112777              25218126

```

```

Egress Schedulers
root          113781          26008462
serv_all      113781          26008462
=====
A:Dut-A#

A:Dut-A# show qos scheduler-stats subscriber alcatel_100 scheduler root
=====
Scheduler Stats
=====
Scheduler          Forwarded Packets      Forwarded Octets
-----
Ingress Schedulers
root                0                      0
Egress Schedulers
root                0                      0
=====
A:Dut-A#

```

## shared-queue

**Syntax** `shared-queue [shared-queue-policy-name] [detail]`

**Context** `show>qos`

**Description** This command displays shared policy information.

### Sample Output

```

A:Dut-A# show qos shared-queue
=====
Shared Queue Policies
=====
Policy-Id          Description
-----
default            Default Shared Queue Policy
=====
A:Dut-A#

A:Dut-A# show qos shared-queue detail
=====
QoS Network Queue Policy
-----
Shared Queue Policy (default)
-----
Policy              : default
Description         : Default Shared Queue Policy
-----
Queue CIR          PIR          CBS          MBS          HiPrio  Multipoint
-----
1      0           100          1           50         10      FALSE
2      25          100          3           50         10      FALSE
3      25          100          10          50         10      FALSE
4      25          100          3           25         10      FALSE
5      100         100          10          50         10      FALSE

```

## Vport Commands

6	100	100	10	50	10	FALSE
7	10	100	3	25	10	FALSE
8	10	100	3	25	10	FALSE
9	0	100	1	50	10	TRUE
10	25	100	3	50	10	TRUE
11	25	100	10	50	10	TRUE
12	25	100	3	25	10	TRUE
13	100	100	10	50	10	TRUE
14	100	100	10	50	10	TRUE
15	10	100	3	25	10	TRUE
16	10	100	3	25	10	TRUE
17	0	100	1	50	10	TRUE
18	25	100	3	50	10	TRUE
19	25	100	10	50	10	TRUE
20	25	100	3	25	10	TRUE
21	100	100	10	50	10	TRUE
22	100	100	10	50	10	TRUE
23	10	100	3	25	10	TRUE
24	10	100	3	25	10	TRUE
25	0	100	1	50	10	TRUE
26	25	100	3	50	10	TRUE
27	25	100	10	50	10	TRUE
28	25	100	3	25	10	TRUE
29	100	100	10	50	10	TRUE
30	100	100	10	50	10	TRUE
31	10	100	3	25	10	TRUE
32	10	100	3	25	10	TRUE

---

FC	UCastQ	MCastQ	BCastQ	UnknownQ
----	--------	--------	--------	----------

---

be	1	9	17	25
l2	2	10	18	26
af	3	11	19	27
l1	4	12	20	28
h2	5	13	21	29
ef	6	14	22	30
h1	7	15	23	31
nc	8	16	24	32

---

### Associations

---

Service : 10	SAP : 1/1/4:101 (shared Q)
Service : 10	SAP : 1/1/4:102 (shared Q)
Service : 10	SAP : 1/1/4:103 (shared Q)
Service : 10	SAP : 1/1/4:104 (shared Q)
Service : 10	SAP : 1/1/4:105 (shared Q)
Service : 10	SAP : 1/1/4:106 (shared Q)
Service : 10	SAP : 1/1/4:107 (shared Q)
Service : 10	SAP : 1/1/4:108 (shared Q)
Service : 10	SAP : 1/1/4:109 (shared Q)
Service : 10	SAP : 1/1/4:110 (shared Q)
Service : 10	SAP : 1/1/4:111 (shared Q)
Service : 10	SAP : 1/1/4:112 (shared Q)
Service : 10	SAP : 1/1/4:113 (shared Q)
Service : 10	SAP : 1/1/4:114 (shared Q)
Service : 10	SAP : 1/1/4:115 (shared Q)
Service : 10	SAP : 1/1/4:116 (shared Q)
Service : 10	SAP : 1/1/4:117 (shared Q)
Service : 10	SAP : 1/1/4:118 (shared Q)
Service : 10	SAP : 1/1/4:119 (shared Q)
Service : 10	SAP : 1/1/4:120 (shared Q)

```

Service : 10          SAP : 1/1/4:121 (shared Q)
Service : 10          SAP : 1/1/4:122 (shared Q)
Service : 10          SAP : 1/1/4:123 (shared Q)
Service : 10          SAP : 1/1/5:279 (shared Q)

```

```

=====
A:Dut-A#

```

## ancp-policy

**Syntax** `ancp-policy [policy-name]`

**Context** `show>subscriber-management`

**Description** This command displays subscriber ANCP policy information.

### Sample Output

```

A:active# show subscriber-mgmt ancp-policy
=====
ANCP Policies
=====
adsl-operator1
vdsl-operator1
-----
Number of ANCP policies : 2
=====
A:active#

A:active# show subscriber-mgmt ancp-policy adsl-operator1
=====
ANCP Policy "adsl-operator1"
=====
I. Rate Reduction      : 0 kbps
I. Rate Adjustment     : 100 percent
I. Rate Monitor        : 0 kbps
I. Rate Monitor Alarm  : no
I. Rate Modify         : scheduler "root"

E. Rate Reduction      : 10 kbps
E. Rate Adjustment     : 100 percent
E. Rate Monitor        : 0 kbps
E. Rate Monitor Alarm  : no
E. Rate Modify         : scheduler "root"
Port Down : N/A
Last Mgmt Change: 01/26/2007 17:10:51
=====
A:active#

A:active# show subscriber-mgmt ancp-policy adsl-operator1 association
=====
ANCP Policy "adsl-operator1" associations
=====
SAP Static Map Associations
-----
- SAP      : 1/1/3                      Svc-id   : 333 (VPLS)
  String   : "ANCP-String-1"
  String   : "ANCP-String-2"

```

```

-----
MSS Static Map Associations
-----
- Cust-id : 1                      MSS-name: mss1
  String : "ANCP-String-3"
-----
Subscriber Associations
-----
No associations found.
Number of associations : 3
=====
A:active#

```

## ancp-string

**Syntax**     **ancp-string**  
**ancp-string** *ancp-string*  
**ancp-string** **customer** *customer-id* **site** *customer-site-name*  
**ancp-string** **sap** *sap-id*

**Context**     show>subscriber-management

**Description**     This command displays subscriber ANCP string information.

**Parameters**     *ancp-string* — Specify the ASCII representation of the DSLAM circuit-id name.  
**customer** *customer-id* — Specify the associated existing customer name.  
**site** *customer-site-name* — Specify the associated customer's configured MSS name.  
**sap** *sap-id* — Specifies the physical port identifier portion of the SAP definition. See [Common Service Commands on page 2062](#) for *sap-id* command syntax.

### Sample Output

```

A:active# show subscriber-mgmt ancp-string
=====
ANCP-Strings
=====
ANCP-String                      Assoc State
-----
"ANCP-String-1"                  SAP      Up
"ANCP-String-2"                  SAP      Down
"ANCP-String-3"                  MSS      Up
"ANCP-String-4"                  MSS      Unknown
"ANCP-String-5"                  ANCP     Up
"ANCP-String-6"                  MSS      Unknown
-----
Number of ANCP-Strings : 6
=====
A:active#

*A:Dut-C# show subscriber-mgmt ancp-string hpolSub43
=====
ANCP-String "hpolSub43"
=====

```

```

Type      : SUB - "hpolSub43"
State     : Up                      Ancp Policy: ancpPol
I. Rate   : 100 kbps                E. Rate    : 200 kbps
Adj I. Rate: N/A                    Adj E. Rate: 200 kbps
Act I. Rate: N/A                    Act E. Rate: 182 kbps
Service Id : 1 (VPRN)
Group      : Alu
Neighbor   : 100.100.100.1:49063

```

```
=====
*A:Dut-C#
```

### Other applicable show command output:

```
A:active# show service id 333 sap 1/1/3 detail
```

```
=====
Service Access Points(SAP)
```

```
=====
Service Id      : 333
SAP              : 1/1/3                      Encap          : null
...
-----
```

```
ANCP Override
```

```
-----
Ing Sched Name: root
- PIR      : 100 kbps
- String   : "ANCP-String-1"
Egr Sched Name: root
- PIR      : 100 kbps
- String   : "ANCP-String-1"
-----
```

```
...
Dro. InProf      : 0                      0
Dro. OutProf     : 0                      0
=====
```

```
A:active#
```

```
A:active# show service customer 1 site mss1
```

```
=====
Customer 1
```

```
=====
Customer-ID      : 1
Description      : Default customer
...
-----
```

```
ANCP Override
```

```
-----
Egr Sched Name: root
- PIR      : 90 kbps
- String   : "ANCP-String-3"
-----
```

```
-----
Service Association
```

```
-----
No Service Association Found.
=====
```

```
A:active#
```

## wpp

**Syntax**     **wpp**  
**wpp** [**portal** *wpp-portal-name*] [**host** *ip-address*] **hosts**  
**wpp portal** *wpp-portal-name*  
**wpp statistics**

**Context**     show>router

**Description**     This command displays WPP port-related information in the specified routing instance.

**Parameters**     **portal** *wpp-portal-name* — Specifies the name of this WPP portal.  
**host** *ip-address* — Specifies the host IP address.  
**hosts** — Displays the hosts enabled on the portal.

**Sample Output**

```
show router wpp
=====
WPP portals
=====
Portal                               Address           Controlled-Rtr     Num-Itf
-----
svr1                                 1.1.1.1           0                   0
svr2                                 2.2.2.2           0                   0
-----
No. of portals: 2
=====

show router wpp portal "svr1"
=====
WPP Portal "svr1"
=====
Address                               : 1.1.1.1
Controlled router                      : 0
Number of enabled interfaces          : 0
Triggered hosts                       : disabled
Last management change                : 01/27/2014 00:48:45
=====
```



## ipoe session

<b>Syntax</b>	<b>ipoe session</b> [sap <i>sap-id</i> ] [mac <i>ieee-address</i> ] [circuit-id <i>circuit-id</i> ] [remote-id <i>remote-id</i> ] [interface <i>ip-int-name</i>   <i>ip-address</i> ] [inter-dest-id <i>intermediate-destination-id</i> ] [no-inter-dest-id] [ip-address <i>ip-prefix</i> [/ <i>prefix-length</i> ]] [port <i>port-id</i> ] [subscriber <i>sub-ident-string</i> ] [sap-session-id <i>sap-session-index</i> ] [wholesaler <i>service-id</i> ] <b>session</b> [sap <i>sap-id</i> ] [mac <i>ieee-address</i> ] [circuit-id <i>circuit-id</i> ] [remote-id <i>remote-id</i> ] [interface <i>ip-int-name</i>   <i>ip-address</i> ] [inter-dest-id <i>intermediate-destination-id</i> ] [no-inter-dest-id] [ip-address <i>ip-prefix</i> [/ <i>prefix-length</i> ]] [port <i>port-id</i> ] [subscriber <i>sub-ident-string</i> ] [sap-session-id <i>sap-session-index</i> ] [wholesaler <i>service-id</i> ] <b>detail</b>
<b>Context</b>	show>service>id
<b>Description</b>	This command displays the identified IPoE session details active on the specified service instance.
<b>Parameters</b>	<b>detail</b> — Displays all IPoE session details.

## Sample Output

```
# show service id 4000 ipoe session
=====
IPoE sessions for svc-id 4000
=====
Sap Id                Mac Address          Up Time          MC-Stdbby
  Subscriber-Id
    [CircuitID] | [RemoteID]
-----
1/1/4:1201.27         00:51:00:00:00:0c   0d 00:00:18
    ipoe-session-001
-----
CID | RID displayed when included in session-key
Number of sessions : 1
=====

# show service id 4000 ipoe session detail
=====
IPoE sessions for service 4000
=====

SAP                   : 1/1/4:1201.27
Mac Address           : 00:51:00:00:00:0c
Circuit-Id            : circuit-id-1
Remote-Id             : remote-id-1
Session Key           : sap-mac

MC-Standby            : No

Subscriber-interface   : sub-int-1
Group-interface        : group-int-1

Up Time               : 0d 00:01:01
Session Time Left     : N/A
Last Auth Time        : 02/28/2015 01:01:09
Min Auth Intvl (left) : 0d 00:05:00 (0d 00:03:59)
Persistence Key       : N/A
```

## Vport Commands

```
Subscriber           : "ipoe-session-001"
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      : "144DFF0000001354D806D5"
Sap-Session-Index    : 1
```

```
IP Address           : 10.10.1.201/24
IP Origin             : Radius
Primary DNS           : N/A
Secondary DNS         : N/A
Primary NBNS          : N/A
Secondary NBNS        : N/A
Address-Pool          : N/A
```

```
IPv6 Prefix          : 2001:db8:a:111::/64
IPv6 Prefix Origin    : Radius
IPv6 Prefix Pool      : ""
IPv6 Del.Pfx.         : 2001:db8:a001:a100::/56
IPv6 Del.Pfx. Origin  : Radius
IPv6 Del.Pfx. Pool    : ""
IPv6 Address          : 2001:db8:a:101::aaa:1
IPv6 Address Origin   : Radius
IPv6 Address Pool     : ""
Primary IPv6 DNS      : N/A
Secondary IPv6 DNS    : N/A
```

```
Radius Session-TO     : N/A
Radius Class          : 
Radius User-Name       : 00:51:00:00:00:0c
```

```
-----
Number of sessions : 1
```

```
=====
```

## Clear Commands

### anccp-sub-string

<b>Syntax</b>	<b>anccp-sub-string</b> <i>string</i>
<b>Context</b>	clear>subscr-mgmt>anccp>anccp
<b>Description</b>	This command clears subscriber ANCP data.
<b>Parameters</b>	<i>string</i> — Clears the ANCP string corresponding to this subscriber ID.

### arp

<b>Syntax</b>	<b>arp</b> { <b>all</b>   <i>ip-address</i> } <b>arp interface</b> [ <i>ip-int-name</i>   <i>ip-address</i> ]
<b>Context</b>	clear>router
<b>Description</b>	This command clears all or specific ARP entries. The scope of ARP cache entries cleared depends on the command line option(s) specified.
<b>Parameters</b>	<b>all</b> — Clears all ARP cache entries. <i>ip-addr</i> — Clears the ARP cache entry for the specified IP address. <b>interface</b> <i>ip-int-name</i> — Clears all ARP cache entries for the interface with the specified name. <b>interface</b> <i>ip-addr</i> — Clears all ARP cache entries for the specified interface with the specified address.

### authentication

<b>Syntax</b>	<b>authentication</b> [ <i>policy-name</i> ] <b>authentication coa-statistics</b>
<b>Context</b>	clear
<b>Description</b>	This command clears subscriber authentication data.
<b>Parameters</b>	<i>policy-name</i> — Clears the authentication policy name. The policy must be already configured. <b>coa-statistics</b> — Clears statistics for incoming RADIUS Change of Authorization requests.

### msap-policy

<b>Syntax</b>	<b>msap-policy</b> <i>msap-policy-name</i>
---------------	--

<b>Context</b>	clear> subscriber-mgmt
<b>Description</b>	This command deletes managed SAPs created by the managed SAP policy.
<b>Parameters</b>	<i>msap-policy-name</i> — Specifies an existing managed SAP policy name. Any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

## peakvalue-stats

<b>Syntax</b>	<b>peakvalue-stats iom</b> ( <i>slot</i>   <b>all</b> ) [ <b>recursive</b> ] <b>peakvalue-stats mda</b> ( <i>mda</i>   <b>all</b> ) [ <b>recursive</b> ] <b>peakvalue-stats port</b> ( <i>port-id</i>   <b>all</b> ) <b>peakvalue-stats pw-port</b> ( <i>pw-port</i>   <b>all</b> ) <b>peakvalue-stats system</b> [ <b>recursive</b> ]
<b>Context</b>	clear> subscriber-mgmt
<b>Description</b>	<p>This command resets the most recent peak counter.</p> <p>Note that clearing one counter will not impact other counters. For example, clearing one IOM's most recent peak value will not impact chassis peak value.</p>
<b>Parameters</b>	<p><b>iom</b> <i>slot</i> — Clears IOM host peak value statistics for the specified IOM.</p> <p><b>mda</b> <i>mda</i> — Clears MDA host peak value statistics for the specified MDA.</p> <p><b>port</b> <i>port-id</i> — Clears port host peak value statistics for the specified port ID.</p> <p><b>pw-port</b> <i>pw-port</i> — Clears pseudowire port host peak value statistics for the specified port.</p> <p><b>Values</b> 1 — 10239</p> <p><b>system</b> — Clears system host peak value statistics.</p> <p><b>all</b> — Clears all host peak value statistics.</p> <p><b>recursive</b> — Resets the sub-level counters. For example, clearing IOM counters with the <b>recursive</b> keyword will also clear counters of all ports counters on that IOM.</p>

## radius-accounting

<b>Syntax</b>	<b>radius-accounting</b> [ <i>policy-name</i> ]
<b>Context</b>	clear> subscriber-mgmt
<b>Description</b>	This command clears RADIUS accounting data for the specified policy.
<b>Parameters</b>	<i>policy-name</i> — The name of the policy. The string is case sensitive and limited to 32 ASCII 7-bit printable characters with no spaces

## scheduler-stats

<b>Syntax</b>	<b>scheduler-stats</b>
<b>Context</b>	clear>qos
<b>Description</b>	This command clears scheduler statistics.

## subscriber

<b>Syntax</b>	<b>subscriber</b> <i>sub-ident-string</i> [ <b>scheduler</b> <i>scheduler-name</i> ] [ <b>ingress</b>   <b>egress</b> ]
<b>Context</b>	clear>qos>scheduler-stats
<b>Description</b>	This command clears scheduler stats per subscriber.
<b>Parameters</b>	<i>sub-ident-string</i> — Clears information for the subscriber profile name. <b>scheduler</b> <i>scheduler-name</i> — Clears information for the specified scheduler-name. <b>egress</b> — Clears egress information for the subscriber. <b>ingress</b> — Clears ingress information for the subscriber.

## sla-profile

<b>Syntax</b>	<b>subscriber</b> <i>sub-ident-string</i> <b>sla-profile</b> <i>sla-profile-name</i> <b>sap</b> <i>sap-id</i> [ <b>scheduler</b> <i>scheduler-name</i> ]
<b>Context</b>	clear>qos>scheduler-stats
<b>Description</b>	This command clears the subscriber's SLA profile scheduler stats.
<b>Parameters</b>	<b>subscriber</b> <i>sub-ident-string</i> — Clears information for the specified subscriber profile name. <b>sla-profile</b> <i>sla-profile-name</i> — Clears information for the specified <i>sla-profile-name</i> . <b>sap</b> <i>sap-id</i> — Clears information for the specified SAP. <b>scheduler</b> <i>scheduler-name</i> — Clears information for the specified <i>scheduler-name</i> .

## srrp

<b>Syntax</b>	<b>srrp</b>
<b>Context</b>	clear>router
<b>Description</b>	This command enables the context to clear and reset SRRP virtual router instances.

## interface

<b>Syntax</b>	<b>interface</b> <i>subscriber-interface</i> [ <b>id</b> <i>srrp-id</i> ]
---------------	---

<b>Context</b>	clear>router>srrp
<b>Description</b>	This command clears and resets SRRP interface instances.
<b>Parameters</b>	<i>subscriber-interface</i> — Specifies an existing subscriber interface name. <div> <b>Values</b> 32 chars max </div> <i>id srrp-id</i> — Specifies an existing SRRP ID. <div> <b>Values</b> 1 — 4294967295 </div>

## statistics

<b>Syntax</b>	<b>statistics interface</b> <i>subscriber-interface</i> [ <i>id srrp-id</i> ]
<b>Context</b>	clear>router>srrp
<b>Description</b>	This command clears statistics for SRRP instances.
<b>Parameters</b>	<i>subscriber-interface</i> — Specifies an existing subscriber interface name. <div> <b>Values</b> 32 chars max </div> <i>id srrp-id</i> — Specifies an existing SRRP ID. <div> <b>Values</b> 1 — 4294967295 </div>

## route-downloader

<b>Syntax</b>	<b>route-downloader</b> <i>name</i> [ <i>vprn vprn</i> ] [ <i>family family</i> ]
<b>Context</b>	clear>aaa
<b>Description</b>	<p>This command clears all the radius-downloaded routes from the internal downloader cache (or protocol RIB/db) (and thus eventually from the RTM itself). The parameters <b>vprn</b> and/or <b>family</b> allow to restrict the deletion of those routes learned in a particular address family (IPv4 or IPv6) and/or a particular VPRN.</p> <p>By default, all VPRNs and both IPv4 and IPv6 families are affected.</p> <p>Note that A clear of the internal protocol DB means the corresponding prefix that were deleted should be removed from the RTM (and from any other exports) as well.</p>
<b>Parameters</b>	<p><b>vprn</b> — Specifies to limit the removal of prefixes to only the specific VPRN. The parameter can be either the service-id or service-name that identifies a VPRN.</p> <p><b>family</b> <i>family</i> — Specifies to limit the removal of prefixes only belonging to the address family IPv4 or IPv6. Only these two values will be accepted.</p> <div> <b>Values</b> ipv4, ipv6 </div>

## vport

<b>Syntax</b>	<b>port</b> <i>port-id</i> <b>vport</b> <i>name</i> [ <b>scheduler</b> <i>scheduler-name</i> ]
<b>Context</b>	clear>qos>scheduler-stats
<b>Description</b>	This command clears the vport scheduler stats.
<b>Parameters</b>	<p><b>port</b> <i>port-id</i> — Clears information for the specified port.</p> <p><b>vport</b> <i>name</i> — Clears information for the specified vport.</p> <p><b>scheduler</b> <i>scheduler-name</i> — Clears information for the specified <i>scheduler-name</i>.</p>

## ipoe session

<b>Syntax</b>	<b>ipoe session</b> [ <b>sap</b> <i>sap-id</i> ] [ <b>interface</b> <i>ip-int-name</i>   <i>ip-address</i> ] [ <b>mac</b> <i>ieee-address</i> ] [ <b>circuit-id</b> <i>circuit-id</i> ] [ <b>remote-id</b> <i>remote-id</i> ] [ <b>inter-dest-id</b> <i>intermediate-destination-id</i> ] [ <b>no-inter-dest-id</b> ] [ <b>ip-address</b> <i>ip-prefix</i> [/ <i>prefix-length</i> ]] [ <b>port</b> <i>port-id</i> ] [ <b>subscriber</b> <i>sub-ident-string</i> ] [ <b>sap-session-id</b> <i>sap-session-index</i> ] <b>ipoe session all</b>
<b>Context</b>	clear>service>id
<b>Description</b>	This commands clears all identified IPoE sessions for the specified service instance. All associated subscriber hosts will be deleted from the system.
<b>Parameters</b>	<b>all</b> — clears all active IPoE sessions for the specified service instance.

---

## Tools Commands

### tools

<b>Syntax</b>	<b>tools</b>
<b>Context</b>	<root>
<b>Description</b>	The context to enable useful tools for debugging purposes.
<b>Default</b>	none
<b>Parameters</b>	<b>dump</b> — Enables dump tools for the various protocols. <b>perform</b> — Enables tools to perform specific tasks.

### perform

<b>Syntax</b>	<b>perform</b>
<b>Context</b>	tools
<b>Description</b>	This command enables the context to enable tools to perform specific tasks.
<b>Default</b>	none

### persistence

<b>Syntax</b>	<b>persistence</b>
<b>Context</b>	tools>perform
<b>Description</b>	This command enables the context to configure downgrade parameters.

### downgrade

<b>Syntax</b>	<b>downgrade target-version <i>target</i> [reboot]</b>
<b>Context</b>	tools>perform>persistence
<b>Description</b>	This command downgrades persistence files to a previous version.
<b>Parameters</b>	<b>target-version <i>target</i></b> — Specifies the downgrade version. <b>reboot</b> — Specifies to reboot the system after a successful conversion.



## subscriber-mgmt

<b>Syntax</b>	<b>subscriber-mgmt</b>
<b>Context</b>	tools>perform
<b>Description</b>	This command enables tools to control subscriber management.

## edit-lease-state

<b>Syntax</b>	<b>edit-lease-state sap</b> <i>sap-id</i> <b>ip</b> <i>ip-address</i> [ <b>subscriber</b> <i>sub-ident-string</i> ] [ <b>sub-profile-string</b> <i>sub-profile-string</i> ] [ <b>sla-profile-string</b> <i>sla-profile-string</i> ] <b>edit-lease-state svc-id</b> <i>service-id</i> <b>ip</b> <i>ip-address</i> [ <b>subscriber</b> <i>sub-ident-string</i> ] [ <b>sub-profile-string</b> <i>sub-profile-string</i> ] [ <b>sla-profile-string</b> <i>sla-profile-string</i> ]				
<b>Context</b>	tools>perform>subscr-mgmt				
<b>Parameters</b>	<p><b>sap</b> <i>sap-id</i> — Specifies the physical port identifier portion of the SAP definition. See <a href="#">Common Service Commands on page 2062</a> for <i>sap-id</i> command syntax.</p> <p><b>ip</b> <i>ip-address</i> — Modifies lease state information for the specified IP address.</p> <p><b>subscriber</b> <i>sub-ident-string</i> — Modifies lease state information for the specified subscriber identification.</p> <p><b>sub-profile-string</b> <i>sub-profile-string</i> — Modifies lease state information for the specified subscriber profile.</p> <p><b>sla-profile-string</b> <i>sla-profile-string</i> — Modifies lease state information for the specified SLA profile.</p> <p><b>svc-id</b> <i>service-id</i> — Modifies lease state information for the specified service ID.</p>				
<b>Values</b>	<table> <tr> <td><i>service-id:</i></td><td>1 — 2147483647</td></tr> <tr> <td><i>svc-name:</i></td><td>64 characters maximum</td></tr> </table>	<i>service-id:</i>	1 — 2147483647	<i>svc-name:</i>	64 characters maximum
<i>service-id:</i>	1 — 2147483647				
<i>svc-name:</i>	64 characters maximum				

## credit-reset

<b>Syntax</b>	<b>credit-reset sap</b> <i>sap-id</i> <b>subscriber</b> <i>sub-ident-string</i> <b>sla-profile</b> <i>sla-profile-name</i> { <b>category</b> <i>category-name</i>   <b>all-categories</b> } <b>credit-reset sap</b> <i>sap-id</i> <b>ip</b> <i>ip-address</i> { <b>category</b> <i>category-name</i>   <b>all-categories</b> } <b>credit-reset svc</b> <i>service-id</i> <b>ip</b> <i>ip-address</i> { <b>category</b> <i>category-name</i>   <b>all-categories</b> }
<b>Context</b>	tools>perform>subscr-mgmt
<b>Description</b>	This command resets the credit for an SLA-profile instance.
<b>Parameters</b>	<p><b>sap</b> <i>sap-id</i> — Specifies the physical port identifier portion of the SAP definition. See <a href="#">Common Service Commands on page 2062</a> for <i>sap-id</i> command syntax.</p> <p><b>ip</b> <i>ip-address</i> — Modifies lease state information for the specified IP address.</p>

**subscriber** *sub-ident-string* — Modifies lease state information for the specified subscriber identification.

**sub-profile-string** *sub-profile-string* — Modifies lease state information for the specified subscriber profile.

**sla-profile-string** *sla-profile-string* — Modifies lease state information for the specified SLA profile.

**svc-id** *service-id* — Modifies lease state information for the specified service ID.

<b>Values</b>	<i>service-id:</i>	1 — 2147483647
	<i>svc-name:</i>	64 characters maximum

## edit-ipoe-session

<b>Syntax</b>	<b>edit-ipoe-session</b> <b>sap</b> <i>sap-id</i> <b>mac</b> <i>mac-address</i> [ <b>subscriber</b> <i>sub-ident-string</i> ] [ <b>sub-profile-string</b> <i>sub-profile-string</i> ] [ <b>sla-profile-string</b> <i>sla-profile-string</i> ] [ <b>inter-dest-id</b> <i>intermediate-destination-id</i> ] [ <b>ancp-string</b> <i>ancp-string</i> ] [ <b>app-profile-string</b> <i>app-profile-string</i> ] [ <b>circuit-id</b> <i>circuit-id</i> ] [ <b>remote-id</b> <i>remote-id</i> ]
<b>Context</b>	tools>perform>subscr-mgmt
<b>Description</b>	This command updates the data of the IPoE session identified with the given MAC address and SAP identifier. Optionally the remote-id and circuit-id can be specified to identify the IPoE session to update.  Note that the changes take immediate effect.

## eval-ipoe-session

<b>Syntax</b>	<b>eval-ipoe-session</b> [ <b>svc-id</b> <i>service-id</i> ] [ <b>sap</b> <i>sap-id</i> ] [ <b>mac</b> <i>mac-address</i> ] [ <b>circuit-id</b> <i>circuit-id</i> ] [ <b>remote-id</b> <i>remote-id</i> ] [ <b>subscriber</b> <i>sub-ident-string</i> ]
<b>Context</b>	tools>perform>subscr-mgmt
<b>Description</b>	This command re-evaluates the mapping between authentication strings such as the SLA profile string and the actual profiles for the identified IPoE sessions.

## eval-lease-state

<b>Syntax</b>	<b>eval-lease-state</b> [ <b>svc-id</b> <i>service-id</i> ] [ <b>sap</b> <i>sap-id</i> ] [ <b>subscriber</b> <i>sub-ident-string</i> ] [ <b>ip</b> <i>ip-address</i> ]				
<b>Context</b>	tools>perform>subscr-mgmt				
<b>Description</b>	This command evaluates lease state information.				
<b>Parameters</b>	<b>svc-id</b> <i>service-id</i> — Evaluates lease state information for the specified service.				
<b>Values</b>	<table> <tr> <td><i>service-id:</i></td> <td>1 — 2147483647</td> </tr> <tr> <td><i>svc-name:</i></td> <td>64 characters maximum</td> </tr> </table>	<i>service-id:</i>	1 — 2147483647	<i>svc-name:</i>	64 characters maximum
<i>service-id:</i>	1 — 2147483647				
<i>svc-name:</i>	64 characters maximum				

**sap** *sap-id* — Evaluates lease state information for the specified SAP.

*sap-id* — Specifies the physical port identifier portion of the SAP definition. See [Common Service Commands on page 2062](#) for *sap-id* command syntax.

**subscriber** *sub-ident-string* — Evaluates lease state information for the specified subscriber identification string.

**ip** *ip-address* — Evaluates lease state information for the specified IP address.

## re-ident-sub

<b>Syntax</b>	<b>re-ident-sub</b> <i>old-sub-ident-string</i> <b>to</b> <i>new-sub-ident-string</i>
<b>Context</b>	tools>perform>subscr-mgmt
<b>Description</b>	This command renames a subscriber identification string.
<b>Parameters</b>	<i>old-sub-ident-string</i> — Specifies the existing subscriber identification string to be renamed. <i>new-sub-ident-string</i> — Specifies the new subscriber identification string name.

## redundancy

<b>Syntax</b>	<b>redundancy</b>
<b>Context</b>	tools>dump
<b>Description</b>	This command enables the context to dump redundancy parameters.

## multi-chassis

<b>Syntax</b>	<b>multi-chassis</b>
<b>Context</b>	tools>dump>redundancy
<b>Description</b>	This command enables the context to dump multi-chassis parameters.

## mc-ipsec

<b>Syntax</b>	<b>mc-ipsec</b>
<b>Context</b>	tools>perform>redundancy>multi-chassis
<b>Description</b>	This command enters the mc-ipsec context.

## force-switchover

<b>Syntax</b>	<b>force-switchover tunnel-group</b> <i>local-group-id</i>
<b>Context</b>	tools>perform>redundancy>multi-chassis>mc-ipsec
<b>Description</b>	This command manually switches over mc-ipsec mastership of the specified tunnel-group.
<b>Parameters</b>	<i>local-group-id</i> — Specifies the local tunnel-group ID configured under config>redundancy.multi-chassis>peer>mc-ipsec.

## mc-ring

<b>Syntax</b>	<b>mc-ring</b>
<b>Context</b>	tools>dump>redundancy>multi-chassis
<b>Description</b>	This command dumps multi-chassis ring data.

## sync-database

<b>Syntax</b>	<b>sync-database</b> [ <b>peer</b> <i>ip-address</i> ] [ <b>port</b> <i>port-id</i>   <i>lag-id</i> ] [ <b>sync-tag</b> <i>sync-tag</i> ] [ <b>application</b> <i>application</i> ] [ <b>detail</b> ] [ <b>type</b> <i>type</i> ]
<b>Context</b>	tools>dump>redundancy>multi-chassis
<b>Description</b>	This command dumps multi-chassis sync database information.
<b>Parameters</b>	<p><b>peer</b> <i>ip-address</i> — Dumps the specified address of the multi-chassis peer.</p> <p><b>port</b> <i>port-id</i> — Dumps the specified port ID of the multi-chassis peer.</p> <p><b>port</b> <i>lag-id</i> — Dumps the specified Link Aggregation Group (LAG) on this system.</p> <p><b>sync-tag</b> <i>sync-tag</i> — Dumps the synchronization tag used while synchronizing this port with the multi-chassis peer.</p> <p><b>application</b> — Dumps the specified application information that was synchronized with the multi-chassis peer.</p> <p><b>Values</b> dhcps, igmp, igmp-snooping, mc-ring, srrp, sub-mgmt, mld-snooping, all</p> <p><b>detail</b> — Displays detailed information.</p> <p><b>type</b> <i>type</i> — Filters by the specified entry type.</p> <p><b>Values</b> alarm-deleted, local-deleted</p>

## srrp-sync-data

<b>Syntax</b>	<b>srrp-sync-database</b> [ <b>instance</b> <i>instance-id</i> ] [ <b>peer</b> <i>ip-address</i> ]
---------------	--

<b>Context</b>	tools>dump>redundancy>multi-chassis
<b>Description</b>	This command dumps multi-chassis SRRP sync database information.
<b>Parameters</b>	<i>instance-id</i> — Specifies the instance ID. <b>Values</b> 1 —4294967295 <i>ip-address</i> — Dumps the specified address (in the form of a.b.c.d).

## route-downloader

<b>Syntax</b>	<b>route-downloader start [force]</b>
<b>Context</b>	tools>perform>aaa
<b>Description</b>	This command causes the download process to start immediately. If an ongoing download is already in progress then no further action is needed, except if the <b>force</b> keyword is added. In case the <b>force</b> keyword is added, then the current download is aborted and a new one is immediately restarted. If aborting the current download, the internal route table should not be emptied or cleared.
<b>Parameters</b>	<b>start</b> — Starts the download process immediately. <b>force</b> — Causes the current download to be aborted and a new one is immediately restarted.

---

## Debug Commands

### arp-host

<b>Syntax</b>	<b>[no]</b> arp-host
<b>Context</b>	debug>service>id
<b>Description</b>	This command enables and configures ARP host debugging. The no form of the command disables ARP host debugging.

### one-time-http-redirection

<b>Syntax</b>	one-time-http-redirection
<b>Context</b>	debug>service>id
<b>Description</b>	This command produces one-time http redirection debug output.

### ppp

<b>Syntax</b>	<b>[no]</b> ppp
<b>Context</b>	debug>service>id>
<b>Description</b>	This command enables the PPP debug context. event

### event

<b>Syntax</b>	<b>[no]</b> event
<b>Context</b>	debug>service>id>ppp
<b>Description</b>	This command enables the PPP event debug context.

### dhcp-client

<b>Syntax</b>	<b>dhcp-client [terminate-only]</b> <b>no dhcp-client</b>
<b>Context</b>	debug>service>id>ppp>event
<b>Description</b>	This command enable PPP event debug for DHCP client.

**Parameters**     **terminate-only** — Enables debug for local terminated PPP session

## l2tp

**Syntax**     **l2tp [terminate-only]**  
**no l2tp**

**Context**     debug>service>id>ppp>event

**Description**     This command enables PPP L2TP event debug.

**Parameters**     **terminate-only** — Enables debug for local terminated PPP session

## ppp

**Syntax**     **ppp [terminate-only]**  
**no ppp**

**Context**     debug>service>id>ppp>event

**Description**     This command enables PPP event debug.

**Parameters**     **terminate-only** — Enables debug for local terminated PPP session

## packet

**Syntax**     **[no] packet**

**Context**     debug>service>id>ppp

**Description**     This command enables the PPP packet debug context.

## detail-level

**Syntax**     **detail-level {low | medium | high}**  
**no detail-level**

**Context**     debug>service>id>ppp>packet

**Description**     This command specify the detail level of PPP packet debug output.

## dhcp-client

**Syntax**     **[no] dhcp-client**

**Context**     debug>service>id>ppp>packet

This command enables packet debug output for DHCP client of the PPP session

## discovery

<b>Syntax</b>	<b>discovery [padi] [pado] [padr] [pads] [padt]</b> <b>no discovery</b>
<b>Context</b>	debug>service>id>ppp>packet
<b>Description</b>	This command enables PPP discovery packet debug output.
<b>Parameters</b>	<b>padi/pado/padr/pads/padt</b> — Enables the corresponding type of PPP discovery packet.

## mode

<b>Syntax</b>	<b>mode {dropped-only   ingr-and-dropped   egr-ingr-and-dropped}</b> <b>no mode</b>
<b>Context</b>	debug>service>id>ppp>packet
<b>Description</b>	This command specifies PPP packet debug mode
<b>Parameters</b>	<b>dropped-only</b> — Only displays dropped packet. <b>ingr-and-dropped</b> — Only displays ingress packet and dropped packet. <b>egr-ingr-and-dropped</b> — Displays ingress, egress and dropped packet.

## ppp

<b>Syntax</b>	<b>ppp [lcp] [pap] [chap] [ipcp] [ipv6cp]</b> <b>no ppp</b>
<b>Context</b>	debug>service>id>ppp>packet
<b>Description</b>	This command enables PPP discovery packet debug output for the specified PPP protocol.
<b>Parameters</b>	<b>lcp/pap/chap/ipcp/ipv6cp</b> — Enables debug for the specified protocol.

## sap

<b>Syntax</b>	<b>[no] sap sap-id</b>
<b>Context</b>	debug>service>id>ppp>packet
<b>Description</b>	This command enables PPP debug output for the specified SAP, this command allow multiple instances.
<b>Parameters</b>	<b>sap-id</b> — Specifies the SAP ID.



## username

<b>Syntax</b>	<b>[no] username</b> <i>username</i>
<b>Context</b>	debug>service>id>ppp
<b>Description</b>	This command enable PPP debug for the specified username. since not all PPP packets contain username, so a mac debug filter will be created automatically when system sees a PPP packet contain the specified username.  Multiple username filters can be specified in the same debug command.
<b>Parameters</b>	<i>user-name</i> — Specifies the ppp username.

## circuit-id

<b>Syntax</b>	<b>[no] circuit-id</b> <i>circuit-id</i>
<b>Context</b>	debug>service>id>ppp
<b>Description</b>	This command enable PPP debug for the specified circuit-id.  Multiple circuit-id filters can be specified in the same debug command.
<b>Parameters</b>	<i>circuit-id</i> — Specifies the circuit-id in PADI.

## remote-id

<b>Syntax</b>	<b>[no]remote-id</b> <i>remote-id</i>
<b>Context</b>	debug>service>id>ppp
<b>Description</b>	This command enable PPP debug for the specified remote-id.  Multiple remote-id filters could be specified in the same debug command.
<b>Parameters</b>	<i>remote-id</i> — Specifies the remote-id in PADI.

## msap

<b>Syntax</b>	<b>[no]msap</b> <i>msap-id</i>
<b>Context</b>	debug>service>id>ppp
<b>Description</b>	This command enable PPP debug for the specified managed SAP.  Multiple msap filters could be specified in the same debug command.
<b>Parameters</b>	<i>msap-id</i> — Specifies the managed SAP ID.

## authentication

<b>Syntax</b>	<b>authentication</b> [ <b>policy</b> <i>policy-name</i> ] [ <b>mac-addr</b> <i>ieee-address</i> ] [ <b>circuit-id</b> <i>circuit-id</i> ]
<b>Context</b>	debug>subscr-mgmt
<b>Description</b>	This command debugs subscriber authentication.
<b>Parameters</b>	<p><b>policy</b> <i>policy-name</i> — Specify an existing subscriber management authentication policy name.</p> <p><b>mac-addr</b> <i>ieee-address</i> — Specifies the 48-bit MAC address xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx.</p> <p><b>circuit-id</b> <i>circuit-id</i> — Specify the circuit-id, up to 256 characters.</p>

## sub-ident-policy

<b>Syntax</b>	<b>[no] sub-ident-policy</b> <i>policy-name</i>
<b>Context</b>	debug>subscr-mgmt
<b>Description</b>	This command debugs subscriber identification policies.
<b>Parameters</b>	<i>policy-name</i> — Specifies the subscriber identification policy to debug.

## script-compile-error

<b>Syntax</b>	<b>[no] script-compile-error</b>
<b>Context</b>	debug>subscr-mgmt>sub-ident-plcy
<b>Description</b>	This command send the traceback of the compile error to the logger. The traceback contains detailed information about where and why the compilation fails. The compilation takes place when the CLI user changes the admin state of the Python URL from shutdown to no-shutdown.

## script-export-variables

<b>Syntax</b>	<b>[no] script-export-variables</b>
<b>Context</b>	debug>subscr-mgmt>sub-ident-plcy
<b>Description</b>	This command sends the result (the three output variables) of the Python script to the logger when the script ran successfully.

## script-output

<b>Syntax</b>	<b>[no] script-output</b>
---------------	---------------------------

<b>Context</b>	debug>subscr-mgmt>sub-ident-plcy
<b>Description</b>	This command sends the output (such as from 'print' statements) of the Python script to the logger.

## script-output-on-error

<b>Syntax</b>	[no] script-output-on-error
<b>Context</b>	debug>subscr-mgmt>sub-ident-plcy
<b>Description</b>	This command sends the output (such as from 'print' statements) of the Python script to the logger, but only when the script fails.

## script-runtime-error

<b>Syntax</b>	[no] script-runtime-error
<b>Context</b>	debug>subscr-mgmt>sub-ident-plcy
<b>Description</b>	This command sends the traceback of the Python script failure to the logger.

## script-all-info

<b>Syntax</b>	script-all-info
<b>Context</b>	debug>subscr-mgmt>sub-ident-plcy
<b>Description</b>	This command enables the script-compile-error, script-export-variables, script-output, script-output-on-error, and script-runtime-error functionalities.

## srrp

<b>Syntax</b>	[no] srrp
<b>Context</b>	debug>router
<b>Description</b>	This command enables debugging for SRRP packets. The <b>no</b> form of the command disables debugging.

## events

<b>Syntax</b>	[no] events [interface <i>ip-int-name</i> ]
<b>Context</b>	debug>router>srrp
<b>Description</b>	This command enables debugging for SRRP packets.

The **no** form of the command disables debugging.

### packets

<b>Syntax</b>	<b>[no] packets [interface <i>ip-int-name</i>]</b>
<b>Context</b>	debug>router>srrp
<b>Description</b>	This command enables debugging for SRRP packets. The <b>no</b> form of the command disables debugging.

### radius

<b>Syntax</b>	<b>[no] radius</b>
<b>Context</b>	debug>router
<b>Description</b>	This command enables the debug router RADIUS context.

### detail-level

<b>Syntax</b>	<b>detail-level {low medium high}</b> <b>no detail-level</b>
<b>Context</b>	debug>router>radius
<b>Description</b>	This command specifies the output detail level of command <b>debug router radius</b> .
<b>Default</b>	medium
<b>Parameters</b>	<b>low</b> — Output includes packet type, server address, length, radius-server-policy name <b>medium</b> — All output in low level plus RADIUS attributes in the packet <b>high</b> — All output in medium level plus hex packet dump

### packet-type

<b>Syntax</b>	<b>packet-type [authentication] [accounting] [coa]</b> <b>no packet-type</b>
<b>Context</b>	debug>router>radius
<b>Description</b>	This command specifies the RADIUS packet type filter of command <b>debug router radius</b>
<b>Default</b>	authentication accounting coa
<b>Parameters</b>	<b>authentication</b> — RADIUS authentication packet.

**accounting** — RADIUS accounting packet.

**coa** — RADIUS change of authorization packet.

## radius-attr

<b>Syntax</b>	<b>radius-attr type attribute-type [transaction]</b> <b>radius-attr type attribute-type [transaction] {address hex integer string} value attribute-value</b> <b>radius-attr vendor vendor-id type attribute-type [transaction] [encoding encoding-type]</b> <b>radius-attr vendor vendor-id type attribute-type [transaction] [encoding encoding-type] {address hex integer string} value attribute-value</b> <b>no radius-attr type attribute-type</b> <b>no radius-attr type attribute-type {address hex integer string} value attribute-value</b> <b>no radius-attr vendor vendor-id type attribute-type</b> <b>no radius-attr vendor vendor-id type attribute-type {address hex integer string} [0..16777215] attribute-value</b>		
<b>Context</b>	debug>router>radius		
<b>Description</b>	This command specifies the RADIUS attribute filter of command <b>debug router radius</b> .		
<b>Default</b>	none		
<b>Parameters</b>	<b>type attribute-type</b> — Specifies the RADIUS attribute type. <div> <b>Values</b> 1 — 255 </div> <b>address</b> — Specifies the value is a IPv4 or IPv6 address/prefix/subnet <b>string</b> — Specifies the value is a ASCII string <b>integer</b> — Specifies the value is a integer <b>hex</b> — Specifies the value is a binary string in hex format, e.g: “\0xAB01FE” <b>value attribute-value</b> — Specifies the value of the RADIUS attribute. <div> <b>Values</b> <div> address &lt;ipv4-address&gt; &lt;ipv6-address&gt;  &lt;ipv6-prefix/prefix-length&gt;  ipv4-address a.b.c.d  ipv6-address x:x:x:x:x:x:x (eight 16-bit pieces)  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]  hex - [0x0..0xFFFFFFFF...(max 506 hex nibbles)]  integer - [0..4294967295]  string - ascii-string (max 253 chars) </div> </div> <b>transaction</b> — With this parameter, system will output both request and response packets in the same session even in case response packet doesn't include the filter attribute. <b>vendor vendor-id</b> — Specifies the vendor id for the vendor specific attribute. <div> <b>Values</b> 0 — 16777215 </div>		

**encoding** *encoding-type* — Specifies the size of vendor-type and vendor-length in bytes. It is a two digitals string: “xy”, x is the size of vendor-type, range from 1-4; y is the size of vendor-length of vendor-length, range from 0-2; it is “11” by default.

**Values** [type-size:1..4][length-size:0..2]

### wpp

**Syntax** [no] wpp

**Context** debug>router

**Description** This command enables the context to configure WPP debugging parameters.

### packet

**Syntax** [no] packet

**Context** debug>router>wpp

**Description** This command enables WPP packet debugging.

### detail-level

**Syntax** detail-level *detail-level*

**Context** debug>router>wpp  
debug>router>wpp>packet

**Description** This command specifies the detail level of WPP packet debugging.

**Parameters** *detail-level* — specifies the detail level of WPP packet debugging

**Values** high — Specifies a high detail level for WPP packet debugging.  
low — Specifies a low detail for WPP packet debugging.

### portal

**Syntax** [no] portal *wpp-portal-name*

**Context** debug>router>wpp

**Description** This command enables WPP debugging for the specified WPP portal.

**Parameters** *wpp-portal-name* — Specifies the WPP portal name.

## Monitor Commands

### subscriber

<b>Syntax</b>	<b>subscriber</b> <i>sub-ident-string</i> <b>sap</b> <i>sap-id</i> <b>sla-profile</b> <i>sla-profile-name</i> [ <b>base</b>   <b>ingress-queue-id</b> <i>ingress-queue-id</i>   <b>egress-queue-id</b> <i>egress-queue-id</i> ] [ <b>interval</b> <i>seconds</i> ] [ <b>repeat</b> <i>repeat</i> ] [ <b>absolute</b>   <b>rate</b> ]
<b>Context</b>	monitor>service
<b>Description</b>	This command monitors statistics for a subscriber.
<b>Parameters</b>	<p><b>sub-ident-string</b> — Specifies an existing subscriber identification profile to monitor.</p> <p><b>sap</b> <i>sap-id</i> — Specifies the physical port identifier portion of the SAP definition. See <a href="#">Common Service Commands on page 2062</a> for <i>sap-id</i> command syntax.</p> <p><b>sla-profile</b> <i>sla-profile-name</i> — Specifies an existing SLA profile.</p> <p><b>interval</b> <i>seconds</i> — Configures the interval for each display in seconds.</p> <p><b>Default</b> 11</p> <p><b>Values</b> 11 — 60</p> <p><b>repeat</b> <i>repeat</i> — Configures how many times the command is repeated.</p> <p><b>Default</b> 10</p> <p><b>Values</b> 1 — 999</p> <p><b>absolute</b> — When the <b>absolute</b> keyword is specified, the raw statistics are displayed, without processing. No calculations are performed on the delta or rate statistics.</p> <p><b>Default</b> mode delta</p> <p><b>rate</b> — When the <b>rate</b> keyword is specified, the rate-per-second for each statistic is displayed instead of the delta.</p> <p><b>base</b> — Monitor base statistics.</p> <p><b>ingress-queue-id</b> <i>ingress-queue-id</i> — Monitors statistics for this queue.</p> <p><b>Values</b> 1 — 32</p> <p><b>egress-queue-id</b> <i>egress-queue-id</i> — Monitors statistics for this queue.</p> <p><b>Values</b> 1 — 8</p>

#### Sample Output

```
A:Dut-A# monitor service subscriber alcatel_100 sap 1/2/1:101 sla-profile sla_default
=====
Monitor statistics for Subscriber alcatel_100
=====
At time t = 0 sec (Base Statistics)
-----
```

## SLA Profile Instance statistics

```

-----
                Packets                Octets
Off. HiPrio      : 0                    0
Off. LowPrio     : 94531                30704535
Off. Uncolor     : 0                    0

Queueing Stats (Ingress QoS Policy 1000)
Dro. HiPrio      : 0                    0
Dro. LowPrio     : 7332                2510859
For. InProf      : 0                    0
For. OutProf     : 87067                28152288

Queueing Stats (Egress QoS Policy 1000)
Dro. InProf      : 880                  127660
Dro. OutProf     : 0                    0
For. InProf      : 90862                12995616
For. OutProf     : 0                    0
-----

```

## SLA Profile Instance per Queue statistics

```

-----
                Packets                Octets
Ingress Queue 1 (Unicast) (Priority)
Off. HiPrio      : 0                    0
Off. LowPrio     : 0                    0
Off. Uncolor     : 0                    0
Dro. HiPrio      : 0                    0
Dro. LowPrio     : 0                    0
For. InProf      : 0                    0
For. OutProf     : 0                    0

Ingress Queue 2 (Unicast) (Priority)
Off. HiPrio      : 0                    0
Off. LowPrio     : 94531                30704535
Off. Uncolor     : 0                    0
Dro. HiPrio      : 0                    0
Dro. LowPrio     : 7332                2510859
For. InProf      : 0                    0
For. OutProf     : 87067                28152288

Ingress Queue 3 (Unicast) (Priority)
Off. HiPrio      : 0                    0
Off. LowPrio     : 0                    0
Off. Uncolor     : 0                    0
Dro. HiPrio      : 0                    0
Dro. LowPrio     : 0                    0
For. InProf      : 0                    0
For. OutProf     : 0                    0

Ingress Queue 11 (Multipoint) (Priority)
Off. HiPrio      : 0                    0
Off. LowPrio     : 0                    0
Off. Uncolor     : 0                    0
Dro. HiPrio      : 0                    0
Dro. LowPrio     : 0                    0
For. InProf      : 0                    0
For. OutProf     : 0                    0

Egress Queue 1
Dro. InProf      : 880                  127660
Dro. OutProf     : 0                    0
-----

```



## Triple Play Service Delivery Architecture

```
For. InProf      : 90862      12995616
For. OutProf     : 0          0
```

### Egress Queue 2

```
Dro. InProf      : 0          0
Dro. OutProf     : 0          0
For. InProf      : 0          0
For. OutProf     : 0          0
```

### Egress Queue 3

```
Dro. InProf      : 0          0
Dro. OutProf     : 0          0
For. InProf      : 0          0
For. OutProf     : 0          0
```

=====

A:Dut-A#

A:Dut-A# monitor service subscriber alcatel\_100 sap 1/2/1:101 sla-profile sla\_default base rate

=====

Monitor statistics for Subscriber alcatel\_100

=====

At time t = 0 sec (Base Statistics)

-----

SLA Profile Instance statistics

```
-----
Packets      Octets
Off. HiPrio   : 0          0
Off. LowPrio  : 109099    35427060
Off. Uncolor  : 0          0
```

Queueing Stats (Ingress QoS Policy 1000)

```
Dro. HiPrio   : 0          0
Dro. LowPrio  : 8449      2894798
For. InProf   : 0          0
For. OutProf  : 100523    32489663
```

Queueing Stats (Egress QoS Policy 1000)

```
Dro. InProf   : 880      127660
Dro. OutProf  : 0          0
For. InProf   : 105578    15104553
For. OutProf  : 0          0
```

-----

At time t = 11 sec (Mode: Rate)

-----

SLA Profile Instance statistics

```
-----
Packets      Octets      % Port
                        Util.
Off. HiPrio   : 0          0          0.00
Off. LowPrio  : 1469      477795    0.38
Off. Uncolor  : 0          0          0.00
```

Queueing Stats (Ingress QoS Policy 1000)

```
Dro. HiPrio   : 0          0          0.00
Dro. LowPrio  : 119      40691    0.03
For. InProf   : 0          0          0.00
For. OutProf  : 1349     437350    0.34
```

Queueing Stats (Egress QoS Policy 1000)

```

Dro. InProf      : 0                0                0.00
Dro. OutProf     : 0                0                0.00
For. InProf      : 1469             209129            0.16
For. OutProf     : 0                0                0.00
=====
A:Dut-A#
A:Dut-A# monitor service subscriber alcatel_100 sap 1/2/1:101 sla-profile sla_default
ingress-queue-id 1
=====
Monitor statistics for Subscriber alcatel_100
=====
At time t = 0 sec (Base Statistics)
-----
                Packets                Octets
Ingress Queue 1 (Unicast) (Priority)
Off. HiPrio      : 0                0
Off. LowPrio     : 0                0
Off. Uncolor     : 0                0
Dro. HiPrio      : 0                0
Dro. LowPrio     : 0                0
For. InProf      : 0                0
For. OutProf     : 0                0
=====
A:Dut-A#

A:Dut-A# monitor service subscriber alcatel_100 sap 1/2/1:101 sla-profile sla_default
egress-queue-id 1
=====
Monitor statistics for Subscriber alcatel_100
=====
At time t = 0 sec (Base Statistics)
-----
                Packets                Octets
Egress Queue 1
Dro. InProf      : 880             127660
Dro. OutProf     : 0                0
For. InProf      : 164366          23506178
For. OutProf     : 0                0
=====
A:Dut-A#

```

## host

<b>Syntax</b>	<b>host [sap <i>sap-id</i>] [wholesaler <i>service-id</i>] [port <i>port-id</i>] [inter-dest-id <i>intermediate-destination-id</i>] [detail]</b> <b>host [sap <i>sap-id</i>] [wholesaler <i>service-id</i>] [port <i>port-id</i>] no-inter-dest-id [detail]</b> <b>host summary</b>
<b>Context</b>	show>service>id
<b>Description</b>	This command displays static host information configured on this service.
<b>Parameters</b>	<b>sap <i>sap-id</i></b> — Displays SAP information for the specified SAP ID. Refer to <a href="#">Common Service Commands on page 2062</a> for <i>sap-id</i> command syntax.

*intermediate-destination-id* — Specifies the intermediate destination identifier which is encoded in the identification strings.

**Values** Up to 32 characters maximum

**summary** — Displays summary host information.

**Values** *service-id:* 1 — 2147483647



# RADIUS Triggered Dynamic Data Services

---

## In This Section

Topics in this section include:

- [Introduction to RADIUS Triggered Dynamic Data Services on page 1442](#)

## Introduction to RADIUS Triggered Dynamic Data Services

RADIUS triggered Dynamic Data Services enables a zero touch, single ended provisioning model for business services. Triggered by the authentication of a single or dual stack PPPoE session or single stack IPv4 host as business CPE control channel, parameters are passed in a RADIUS Access Accept or CoA message to setup a Layer 2 or Layer 3 data service. Dynamic Data Services supported in this release include local Epipe VLL services, Epipe VLL services with dynamic MS-PWs (FEC 129), vpls services with BGP-AD pseudowire, IES and VPRN services. Dynamic Data Service SAPs have to be located on dot1q or qinq encapsulated Ethernet ports and can be part of a LAG.

A Python script interface adds a flexible abstraction layer reducing the OSS integration cost: only the business user specific service parameters (service type, IP address, QoS, and filter parameters) are required from RADIUS and then used in a CLI template to setup the target service. Both XML and RADIUS accounting to up to two different RADIUS destinations can be activated on a dynamic data service SAP.

# RADIUS Triggered Dynamic Data Services Command Reference

## Configuration Commands

### RADIUS Triggered Dynamic Data Services Commands

```

config
  — service
    — dynamic-services
      — dynamic-services-policy dynsrv-policy-name [create]
      — no dynamic-services-policy dynsrv-policy-name
        — accounting-1
          — server-policy policy-name
          — no server-policy
          — stats-type {time | volume-time}
          — no stats-type
          — update-interval [hrs hours] [min minutes] [days days]
          — no update-interval
          — update-interval-jitter absolute seconds
          — no update-interval-jitter
        — accounting-2
          — server-policy policy-name
          — no server-policy
          — stats-type {time|volume-time}
          — no stats-type
          — update-interval [hrs hours] [min minutes] [days days]
          — no update-interval
          — update-interval-jitter absolute seconds
          — no update-interval-jitter
        — cli-user name
        — no cli-user
        — description description-string
        — no description
        — sap-limit [0..131072]
        — no sap-limit
        — script-policy name
        — no script-policy
      — service-range service-id service-id
      — no service-range
      — timers
        — setup-timeout access-accept timeout
        — no setup-timeout

config
  — system

```

## Configuration Commands

```
— security
    — password
        — dynsvc-password password [hash|hash2]
        — no dynsvc-password

<global>
— [no] enable-dynamic-services-config
```



## Show Commands

```

show
  — service
    — dynamic-services
      — dynamic-services-policy [policy-name]
      — root-objects
      — saps [control-session acct-session-id] [port port-id] [dynsvc-policy policy-name]
        [summary] [orphaned] [sap sap-id] [svc-id service-id]
      — script
        — snippets [detail]
        — snippets name snippet-name [instance snippet-instance] [detail]
        — statistics
      — summary
    — sap-using [msap] [dyn-script] [description]

```

## Configuration Commands

## Tools Commands

```
tools
  — perfrom
    — service
      — dynamic-services
        — evaluate-script sap sap-id control-session acct-session-id action script-  
action [dynsvc-policy name] [param-string string]

  — dump
    — service
      — dynamic-services
```

## Clear Commands

```
clear
  — services
    — statistics
      — dynamic-services
```

## Debug Commands

```
debug
  — [no] dynamic-services
    — scripts
      — [no] event
        — [no] cli
        — [no] errors
        — [no] executed-cmd
        — [no] state-change
        — [no] warnings
      — instance instance
        — [no] event
          — [no] cli
          — [no] errors
          — [no] executed-cmd
          — [no] state-change
          — [no] warnings
      — script script
        — [no] event
          — [no] cli
          — [no] errors
          — [no] executed-cmd
          — [no] state-change
          — [no] warnings
```

---

# WIFI Aggregation and Offload Command Reference

---

## Command Hierarchies

### dynamic-services

<b>Syntax</b>	<b>dynamic-services</b>
<b>Context</b>	config>service
<b>Description</b>	This command enables the context to configure dynamic data services. Only available on systems with multi-core CPM (CPM3 or up).

### dynamic-services-policy

<b>Syntax</b>	<b>dynamic-services-policy</b> <i>dynsrv-policy-name</i> [create] <b>no dynamic-services-policy</b> <i>dynsrv-policy-name</i>
<b>Context</b>	config>service>dynsvc
<b>Description</b>	<p>This command creates a new dynamic services policy that can be used to create dynamic data services.</p> <p>The <b>no</b> form of the command removes the dynamic services policy from the configuration. This is only allowed when there are no active dynamic data services referencing this policy.</p>
<b>Default</b>	none
<b>Parameters</b>	<i>dynsrv-policy-name</i> — Specifies a unique name of a dynamic services policy up to 32 characters in length.

### accounting-1

<b>Syntax</b>	<b>accounting-1</b>
<b>Context</b>	config>service>dynsvc>policy
<b>Description</b>	This command enables the context to configure the first RADIUS accounting destination and corresponding RADIUS accounting parameters for dynamic data services.

### accounting-2

<b>Syntax</b>	<b>accounting-2</b>
<b>Context</b>	config>service>dynsvc>policy

**Description** This command enables the context to configure the second RADIUS accounting destination and corresponding RADIUS accounting parameters for dynamic data services.

### server-policy

**Syntax** **server-policy** *policy-name*  
**no server-policy**

**Context** config>service>dynsvc>acct-1  
config>service>dynsvc>acct-2

**Description** This command configures the radius server policy to be used for dynamic data services RADIUS accounting.

The **no** form of the command removes the radius server policy from the configuration. This is only allowed when there are no active dynamic data services referencing this policy.

**Default** no server-policy

**Parameters** *policy-name* — Specifies the nname of the radius server policy.

**Values** max length = 32 characters.

### stats-type

**Syntax** **stats-type** {*time*|*volume-time*}  
**no stats-type**

**Context** config>service>dynsvc>acct-1  
config>service>dynsvc>acct-2

**Description** This command configures the type of statistics to be reported in dynamic data services RADIUS accounting. A RADIUS specified Stats Type overrides the CLI configured value.

The no form of the command resets the default value.

**Default** volume-time

**Parameters** *time* — Only report Session-Time in the RADIUS Accounting Interim-Update and Stop message.

*volume-time* — Report both Session-Time and Volume counter attributes in the RADIUS. Accounting Interim-Update and Stop messages.

### update-interval

**Syntax** **update-interval** [*hrs hours*] [*min minutes*] [*days days*]  
**no update-interval**

**Context** config>service>dynsvc>acct-1  
config>service>dynsvc>acct-2

**Description** This command specifies the interval between each RADIUS Accounting Interim-Update message (minimum 5 minutes; maximum 180 days).

The **no** form of the command disables the sending of Accounting Interim-Update messages.

A RADIUS specified Accounting Interim Interval overrides the CLI configured value.

<b>Default</b>	no update-interval (do not send Accounting Interim-Update messages)
<b>Parameters</b>	<i>hrs</i> — specifies the interval in hours. <div> <b>Values</b> 1 — 23 </div> <i>min</i> — Specifies the interval in minutes. <div> <b>Values</b> 1 — 59 </div> <i>days</i> — specifies the interval in days. <div> <b>Values</b> 1 — 180 </div>

## update-interval-jitter

<b>Syntax</b>	<b>update-interval-jitter absolute seconds</b> <b>no update-interval-jitter</b>
<b>Context</b>	config>service>dynsvc>acct-1 config>service>dynsvc>acct-2
<b>Description</b>	This command specifies the absolute maximum random delay introduced on the update interval between two RADIUS Accounting Interim Update messages. The effective maximum random delay value is the minimum of the configured absolute jitter value and 10% of the configured update-interval.  A value of zero will send the accounting interim update message without introducing an additional random delay.  The <b>no</b> form of the command sets the default to 10% of the configured update-interval.
<b>Default</b>	no update-interval-jitter (10% of the configured update-interval)
<b>Parameters</b>	<i>seconds</i> — Specifies the absolute maximum jitter value in seconds. <div> <b>Values</b> 0 — 3600 </div>

## cli-user

<b>Syntax</b>	<b>cli-user name</b> <b>no cli-user</b>
<b>Context</b>	config>service>dynsvc>policy
<b>Description</b>	This command specifies the CLI user to be used to execute the dynamic data services CLI scripts. Via the specified user's profile, it is possible to further restrict the internal list of allowed commands to be executed via dynamic data service CLI scripts.  The <b>no</b> form of the command sets the CLI user to an internal user with all configuration rights.
<b>Default</b>	no cli-user
<b>Parameters</b>	<i>name</i> — Specifies the CLI user name that must exist in the >config>system>security CLI context.

## description

<b>Syntax</b>	<b>description</b> <i>description-string</i> <b>no description</b>
<b>Context</b>	config>service>dynsvc>policy
<b>Description</b>	The <b>description</b> command associates a text string with a configuration context to help identify the content in the configuration file.  The <b>no</b> form of this command removes the string from the configuration.
<b>Default</b>	no description
<b>Parameters</b>	<i>description-string</i> — The description character string. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

## sap-limit

<b>Syntax</b>	<b>sap-limit</b> [0..131072] <b>no sap-limit</b>
<b>Context</b>	config>service>dynsvc>policy
<b>Description</b>	This command specifies a limit for the number of dynamic data service instances (SAPs) that can be setup simultaneously using a given dynamic services policy.  A value of zero (0) means the policy will be drained: existing dynamic data services can be modified and torn down but no new dynamic data services can be setup.
<b>Default</b>	sap-limit 1
<b>Parameters</b>	[0..131072] — Specifies the number of dynamic data service SAPs that can be setup simultaneously using this dynamic services policy.  <b>Values</b> 0 — 131072

## script-policy

<b>Syntax</b>	<b>script-policy</b> <i>name</i> <b>no script-policy</b>
<b>Context</b>	config>service>dynsvc>policy
<b>Description</b>	This command specifies the radius script policy to be used to setup the dynamic data services. The script-policy configuration cannot be changed when there are active dynamic data services referencing the policy.  The <b>no</b> form of this command removes the script-policy from the configuration. This is only allowed when there are no active dynamic data services referencing this policy.
<b>Default</b>	no script-policy
<b>Parameters</b>	<i>name</i> — Specifies the radius script policy name.

## service-range

<b>Syntax</b>	<b>service-range</b> <i>service-id service-id</i> <b>no service-range</b>
<b>Context</b>	config>service>dynsvc
<b>Description</b>	<p>This command specifies the service id range that is reserved for dynamic data service creation. The range cannot overlap with existing static configured services. Once configured with active dynamic services in the range, the service-range can only be extended at the end.</p> <p>The <b>no</b> form of this command removes the service-range from the configuration. This is only allowed when there are no active dynamic data services.</p> <p>When <b>no service-range</b> is specified, the setup of dynamic data services will fail.</p>
<b>Default</b>	no service-range
<b>Parameters</b>	<i>service-id</i> — Specifies the start and end service-id to define the service-range for dynamic services.
<b>Values</b>	1 — 2147483647

## timers

<b>Syntax</b>	<b>timers</b>
<b>Context</b>	config>service>dynsvc
<b>Description</b>	This command enables the context to configure dynamic data services related timers.

## setup-timeout

<b>Syntax</b>	<b>setup-timeout</b> <i>access-accept timeout</i> <b>no setup-timeout</b>
<b>Context</b>	config>service>dynsvc>timers
<b>Description</b>	<p>This command specifies the time that dynamic data services setup requests from a RADIUS Access-Accept are hold in an internal work queue waiting to be processed. If after the timeout, the dynamic data service setup request is still in the queue (meaning it is not setup), then the dynamic service setup request will be removed from the queue and the setup fails.</p> <p>The <b>no</b> form of this command resets the timeout to 2 seconds.</p>
<b>Default</b>	no setup-timeout (30 seconds)
<b>Parameters</b>	<i>timeout</i> — Specifies the setup-timeout in seconds for setup requests of dynamic services received via Access-Accept.
<b>Values</b>	2 — 3600 seconds

## dynsvc-password

<b>Syntax</b>	<b>dynsvc-password</b> <i>password</i> [ <b>hash</b>   <b>hash2</b> ] <b>no dynsvc-password</b>
<b>Context</b>	config>system>security>password
<b>Description</b>	<p>NOTE: See also the description for the <b>enable-dynamic-services-config</b> command.</p> <p>This command allows a user with admin permissions to configure a system wide password which enables a user to enter a special dynamic services configuration mode.</p> <p>The minimum length of the password is determined by the minimum-length command. The complexity requirements for the password are determined by the complexity command.</p> <p>The <b>no</b> form of the command removes the dynsvc password from the configuration</p>
<b>Default</b>	no dynsvc-password
<b>Parameters</b>	<p><i>password</i> — Configures the password which enables a user to enter a special dynamic services configuration mode. The maximum length can be up to 20 characters if unhashed, 32 characters if hashed, 54 characters if the hash2 keyword is specified.</p> <p><b>hash</b> — Specifies the key is entered in an encrypted form. If the hash parameter is not used, the key is assumed to be in a non-encrypted, clear text form.</p> <p><b>hash2</b> — — Specifies the key is entered in a more complex encrypted form. If the hash2 parameter is not used, the key is assumed to be in a non-encrypted, clear text form.</p>

## enable-dynamic-services-config

<b>Syntax</b>	<b>[no] enable-dynamic-services-config</b>
<b>Context</b>	<global>
<b>Description</b>	<p>NOTE: See also the description for the <b>dynsvc-password</b> command.</p> <p>If the dynsvc-password is configured in the <b>config&gt;system&gt;security&gt;password</b> context, then any user can enter a special dynamic services configuration mode by entering the <b>enable-dynamic-services-config</b> command.</p> <p>The <b>enable-dynamic-services-config</b> command is not in the default profile. To give access to this command, the user must belong to the administrative profile or a new profile should be created.</p> <p>Once the <b>enable-dynamic-services-config</b> command is entered, the user is prompted for a password. If the password matches, the user is given access to the dynamic services configuration. Access to static configuration is in this case prohibited.</p> <p>To verify that a user is in the <b>enable-dynamic-services-config</b> mode, use the <b>show users</b> command. Users in the <b>enable-dynamic-services-config</b> mode lists the letter “D” next to the user’s CLI session.</p> <p>The <b>no</b> form of this command disables the dynamic services configuration mode for this user.</p>
<b>Default</b>	no enable-dynamic-services-config



## Show Commands

### dynamic-services

<b>Syntax</b>	<b>dynamic-services</b>
<b>Context</b>	show>service
<b>Description</b>	This command enables the context to show dynamic services information.

### dynamic-services-policy

<b>Syntax</b>	<b>dynamic-services-policy</b> [ <i>policy-name</i> ]
<b>Context</b>	show>service>dynsvc
<b>Description</b>	This command displays the dynamic services policy information.
<b>Parameters</b>	<i>policy-name</i> — specifies for which dynamic services policy the information is requested.

#### Sample Output

```
# show service dynamic-services dynamic-services-policy "dynsvc-policy-1"
=====
Dynamic Service Policies
=====
Dynamic Service Policy      : dynsvc-policy-1
-----
cli-user                    : (Not Specified)
description                  : Dynamic Services Policy 1
script-policy                : script-policy-5
sap-limit                    : 2000

Accounting instance 1
Stats type                   : volume-time
server policy                 : aaa-server-policy-1
Update interval (minutes)    : 30
Update interval jitter       : 180s

Accounting instance 2
Stats type                   : time
server policy                 : aaa-server-policy-2
Update interval (minutes)    : 0
Update interval jitter       : 10%
-----
No. of Services-policies: 1
=====
```

## root-objects

<b>Syntax</b>	<b>root-objects</b>
<b>Context</b>	show>service>dynsvc
<b>Description</b>	<p>This command displays the root objects created by dynamic data services.</p> <p>OID prefix and index — The corresponding SNMP OID prefix and index for this root object.</p> <p>Snippet name — The name of the python function that created this root object. The name is set to N/A when the root-object is orphaned.</p> <p>Snippet instance — The instance for which the python function with “Snippet name” created this root object. If the snippet is a result from a dynamic reference, then the snippet instance is the reference-id string passed in the dyn.reference(). If the snippet is not the result from a dynamic reference, then the snippet instance is the dynamic data service SAP-ID. The instance is set to N/A when the root object is orphaned.</p> <p>Orphan time — Shows the timestamp when the root-object became orphaned (root-object not deleted when corresponding teardown function is called) or N/A if the root-object is not orphaned.</p>

## Sample Output

```
# show service dynamic-services root-objects
=====
Dynamic Service Root Objects
=====
OID prefix           : svcRowStatus
OID index            : .100000
Snippet name         : vprn
Snippet instance     : VRF-1
Orphan time          : N/A
-----
No. of Root Objects: 1
=====
```

## saps

<b>Syntax</b>	<b>saps [control-session <i>acct-session-id</i>] [port <i>port-id</i>] [dynsvc-policy <i>policy-name</i>] [summary] [orphaned] [sap <i>sap-id</i>] [svc-id <i>service-id</i>]</b>
<b>Context</b>	show>service>dynsvc
<b>Description</b>	<p>This command displays the dynamic services SAPs (instances) details:</p> <ul style="list-style-type: none"> <li>• SAP — The dynamic service SAP id.</li> <li>• Acct session-ID — The dynamic service accounting session id.</li> <li>• Acct session-ID control — The control channel accounting session id.</li> <li>• Service — The dynamic service id.</li> <li>• Orphaned — Yes/No – orphaned state is when the SAP is not deleted when corresponding dynamic service teardown function was called.</li> <li>• Dynamic Services policy — The policy referenced to setup the dynamic service.</li> </ul>

- Number of scripts executed — The number of times the script was executed for this dynamic service (setup, modify, revert or teardown).
- Number of scripts w success — The number of times the script was executed successfully for this dynamic service.
- Last script action — The setup, modify, revert, teardown.
- Time of last script action — The timestamp.
- Parameters of last action — The content of the Dynamic Services Script Parameters attribute corresponding with the last action.
- For each of the two accounting instances:
  - Status — RADIUS accounting enabled or disabled.
  - Stats type — the type of statistics reported in accounting.
- Update interval (minutes) — the interval between Accounting Interim Update messages.

**Parameters**

**summary** — Displays a summary view only.

**orphaned** — Displays only SAPs in the orphaned state.

Filtering options, display SAPs that belong to the specified:

**control-session** *acct-session-id* — Specifies control session accounting session id

**port** *port-id* — Specifies Ethernet port.

**dynsvc-policy** *policy-name* — Specifies dynamic services policy.

**sap sap-id** — Specifies dynamic services SAP id.

**svc-id** *service-id* — - service ID of the dynamic service.

**Sample Output**

```
# show service dynamic-services saps
=====
Dynamic Services SAP's
=====
SAP                               : 1/1/1:1.901
-----
Acct session-ID                   : 242FFF000001AE512CE4B6
Acct session-ID control           : 242FFF000001AB512CE4B6
Service                           : [100000]
Orphaned                          : no
Dynamic Services policy           : dynsvc-policy-1
Number of scripts executed        : 1
Number of scripts w success       : 1
Last script action                 : setup
Time of last script action         : 2013/02/26 16:37:10
Parameters of last action          : vprn_1={ 't': ('VRF-1',65000,1000,'cpe-int-1','192.
                                     : 168.20.1/24','2001:db8:cafe::1/64',901,901,910,92
                                     : 0,'192.168.20.0/24','192.168.20.2') }

Accounting instance 1
Status                           : enabled
Stats type                       : volume-time
Update interval (minutes)        : 30

Accounting instance 2
Status                           : enabled
```

## Show Commands

```
Stats type           : time
Update interval (minutes) : 0

-----
No. of SAP's: 1
=====

# show service dynamic-services saps summary
=====
Dynamic Services SAP's summary
=====
SAP                      Acct-Session-ID      Acct-Session-ID-Ctrl
-----
1/1/1:1.901              242FFF000001AE512CE4B6  242FFF000001AB512CE4B6
-----
No. of SAP's: 1
=====
```

## script

<b>Syntax</b>	<b>script</b>
<b>Context</b>	show>service>dynsvc
<b>Description</b>	This command enables the context to show dynamic services script information.

## snippets

<b>Syntax</b>	<b>snippets [detail]</b> <b>snippets name snippet-name [instance snippet-instance] [detail]</b>
<b>Context</b>	show>service>dynsvc>script
<b>Description</b>	<p>This command displays the dynamic services snippets information.</p> <p>The CLI output generated by a single dynamic service python function call is a snippet instance.</p> <p>The name of the snippet instance is the function key in the dyn.action() dictionary that caused this function to be called. This name is the dictionary name passed via RADIUS for top-level snippets or the first parameter to dyn.reference() for the others.</p> <p>The snippet instance is either the dynamic data service SAP id or if the function is called via dynamic reference, the reference-id (that is, second parameter) provided in the dyn.reference() call.</p>
<b>Parameters</b>	<p><b>detail</b> — display detailed dynamic services snippet information.</p> <p>Filtering options, display dynamic services snippets information that matches:</p> <p><b>name snippet-name</b> — Specifies the snippet name.</p> <p><b>instance snippet-instance</b> — Specifies the snippet instance.</p>

### Sample Output

```
# show service dynamic-services script snippets
```

```

=====
Dynamic Services Snippets
=====
Name                      Instance                      Ref-count  Dict-len
-----
vprn                      VRF-1                      1          75
vprn_1                    1/1/1:1.901                0          190
-----
No. of Snippets: 2
=====

# show service dynamic-services script snippets detail
=====
Dynamic Service Snippets
=====
Snippet                   : vprn:VRF-1
-----
reference-count           : 1
dictionary-length        : 75

Root-object
-----
oid                       : svcRowStatus.100000

Reserved-id
-----
id                        : service-id:100000
-----
Snippet                   : vprn_1:1/1/1:1.901
-----
reference-count           : 0
dictionary-length        : 190

Referenced-snippet
-----
snippet                   : vprn:VRF-1
-----
No. of Snippets: 2
=====

```

## statistics

<b>Syntax</b>	<b>statistics</b>
<b>Context</b>	show>service>dynsvc>script
<b>Description</b>	This command displays dynamic service script statistics. Only non-zero values are shown. The script statistics can be cleared with the “clear service statistics dynamic-services” command.

### Sample Output

```

# show service dynamic-services script statistics
=====
Dynamic Services Script Statistics
=====
Description                      Counter

```

## Show Commands

```
-----
python scripts with 0 retries due to timeout          46
setup    - jobs launched                             16
setup    - jobs handled                              16
setup    - success                                   13
setup    - syntax error                             1
setup    - execution failed                          2
teardown - jobs launched                             15
teardown - jobs handled                              15
teardown - success                                   14
teardown - syntax error                             1
-----
No. of Script Statistics: 10
-----
Last Cleared Time: 02/26/2013 09:59:07
=====
```

## summary

<b>Syntax</b>	<b>summary</b>
<b>Context</b>	show>service>dynsvc
<b>Description</b>	This command displays the global configuration summary for dynamic services: <ul style="list-style-type: none"><li>• Service range</li><li>• Timers</li></ul>

### Sample Output

```
# show service dynamic-services summary
=====
Dynamic Services Summary
=====
range start           : 100000
range end             : 200001
setup timeout Access Accept : 30
=====
```

## sap-using

<b>Syntax</b>	<b>sap-using [msap] [dyn-script] [description]</b>
<b>Context</b>	show>service
<b>Description</b>	This command displays SAP information.
<b>Parameters</b>	<b>dyn-script</b> — Displays dynamic service SAPs information.

### Sample Output

```
# show service sap-using dyn-script
```

```

=====
Service Access Points
=====
PortId                      SvcId      Ing.  Ing.  Egr.  Egr.  Adm  Opr
                        QoS      Fltr  QoS   Fltr
-----
[1/1/1:1.901]              [100000]   901   ip4    901   ip4    Up   Up
-----
Number of SAPs : 1
-----
Number of Dynamic Service SAPs : 1, indicated by [<sap-id>] [<svc-id>]
-----
=====

# show service sap-using dyn-script description
=====
Service Access Points
=====
PortId                      SvcId      Adm  Opr  Description
-----
[1/1/1:1.901]              [100000]    Up   Up   This is a dynamic
                                service SAP
-----
Number of SAPs : 1
-----
Number of Dynamic Service SAPs : 1, indicated by [<sap-id>] [<svc-id>]
-----
=====

```

---

## Clear Command

### dynamic-services

<b>Syntax</b>	<b>dynamic-services</b>
<b>Context</b>	clear>service>stats
<b>Description</b>	This command resets the dynamic services script statistics. See also <b>show service dynamic-services script statistics</b> .



---

## Debug Commands

### dynamic-services

<b>Syntax</b>	<b>[no] dynamic-services</b>
<b>Context</b>	debug
<b>Description</b>	This command enables the context to configure dynamic services debugging.

### scripts

<b>Syntax</b>	<b>scripts</b>
<b>Context</b>	debug>dynsvc
<b>Description</b>	This command enables the context to configure dynamic services script debugging.

### event

<b>Syntax</b>	<b>[no] event</b>
<b>Context</b>	debug>dynsvc>scripts debug>dynsvc>scripts>inst debug>dynsvc>scripts>script
<b>Description</b>	This command enables/disables the generation of all dynamic data service script debugging events output: cli, errors, executed-cmd, warnings, state-change.

### cli

<b>Syntax</b>	<b>[no] cli</b>
<b>Context</b>	debug>dynsvc>scripts>event debug>dynsvc>scripts>inst>event debug>dynsvc>scripts>script>event
<b>Description</b>	This command enables/disables the generation of a specific dynamic data service script debugging event output: cli.

### errors

<b>Syntax</b>	<b>[no] errors</b>
<b>Context</b>	debug>dynsvc>scripts>event

```
debug>dynsvc>scripts>inst>event  
debug>dynsvc>scripts>script>event
```

**Description** This command enables/disables the generation of a specific dynamic data service script debugging event output: errors.

### executed-cmd

**Syntax** [no] **executed-cmd**

**Context** debug>dynsvc>scripts>event  
debug>dynsvc>scripts>inst>event  
debug>dynsvc>scripts>script>event

**Description** This command enables/disables the generation of a specific dynamic data service script debugging event output: executed-cmd.

### state-change

**Syntax** [no] **state-change**

**Context** debug>dynsvc>scripts>event  
debug>dynsvc>scripts>inst>event  
debug>dynsvc>scripts>script>event

**Description** This command enables/disables the generation of a specific dynamic data service script debugging event output: state-change.

### warnings

**Syntax** [no] **warnings**

**Context** debug>dynsvc>scripts>event  
debug>dynsvc>scripts>inst>event  
debug>dynsvc>scripts>script>event

**Description** This command enables/disables the generation of a specific dynamic data service script debugging event output: warnings.

### instance

**Syntax** **instance** *instance*

**Context** debug>dynsvc>scripts

**Description** This command enables the context to configure dynamic services script debugging for a specific instance.

**Parameters** *instance* — Specifies the instance name.

## script

<b>Syntax</b>	<b>script</b> <i>script</i>
<b>Context</b>	debug>dynsvc>scripts
<b>Description</b>	This command enables the context to configure dynamic services script debugging for a specific script.
<b>Parameters</b>	<i>script</i> — Specifies the script name.

---

## Tools Commands

### dynamic-services

<b>Syntax</b>	<b>dynamic-services</b>
<b>Context</b>	tools>perform>service
<b>Description</b>	This command enables the context to execute dynamic services tools perform commands.

### evaluate-script

<b>Syntax</b>	<b>evaluate-script sap</b> <i>sap-id</i> <b>control-session</b> <i>acct-session-id</i> <b>action</b> <i>script-action</i> [ <b>dynsvc-policy</b> <i>name</i> ] [ <b>param-string</b> <i>string</i> ]
<b>Context</b>	tools>perform>service>dynamic-services
<b>Description</b>	This tools command performs the execution of a dynamic service script action as if the corresponding RADIUS attributes were received from RADIUS. It is possible to setup, modify or teardown a dynamic service associated with the specified control channel.
<b>Parameters</b>	<p><b>sap</b> <i>sap-id</i> — specifies the dynamic service SAP id.</p> <p><b>control-session</b> <i>acct-session-id</i> — Specifies the accounting session id of the control channel associated with this dynamic service</p> <p><b>action</b> <i>script-action</i> — Specifies the requested action: setup, modify or teardown.</p> <p><b>dynsvc-policy</b> <i>name</i> — Specifies the dynamic services policy to use for this action. Mandatory parameter for setup and modify actions. In case of a modify action, the dynamic services policy must be the same as the policy used at setup.</p> <p><b>param-string</b> <i>string</i> — Specifies the dynamic service parameter list. Mandatory parameter for setup and modify actions.</p>

### dynamic-services

<b>Syntax</b>	<b>dynamic-services command-list</b>
<b>Context</b>	tools>dump>service
<b>Description</b>	<p>This command displays the list of supported commands that are allowed to be used in dynamic service CLI scripts.</p> <p>There are two types of CLI nodes in this list:</p> <ul style="list-style-type: none"> <li>• Pass through nodes: navigation is allowed but attributes creation or attribute changes are prohibited in this node.</li> <li>• Allowed nodes: navigation, attribute creation and attribute changes are allowed in this node (available from 11.0R2 onward).</li> </ul>

# Python Script Support for ESM

---

## In This Chapter

This section describes the Python script support for Enhanced Subscriber Management (ESM).

The following topics are included:

- [Python Script Support for ESM on page 1466](#)
- [Python in SR-OS Overview on page 1467](#)
  - [Python Changes on page 1467](#)
- [Python Support in sub-ident-policy on page 1468](#)
  - [Configuration on page 1470](#)
  - [Operator Debugging on page 1472](#)
  - [Python Scripts on page 1473](#)
  - [Sample Python Scripts on page 1474](#)
  - [Limitations on page 1478](#)
- [RADIUS Script Policy Overview on page 1479](#)
  - [Python RADIUS API on page 1480](#)
  - [Sample Script on page 1480](#)
- [Python Policy Overview on page 1481](#)
  - [Python Policy – RADIUS API on page 1482](#)
  - [Python Policy – DHCPv4 API on page 1482](#)
  - [Python Policy – DHCPv6 API on page 1486](#)
  - [Python Policy – Diameter API on page 1493](#)
  - [Python Policy – DHCP Transaction Cache API on page 1500](#)
  - [Applying a Python Policy on page 1504](#)
  - [Python Script Protection on page 1504](#)
- [Tips and Tricks on page 1505](#)

## Python Script Support for ESM

In order to provide programmable flexibility in ESM applications, the SR OS provides the following features with Python script support:

- sub-ident-policy
- radius-script-policy
- python-policy

## Python in SR-OS Overview

The SR-OS python script support is based on Python version 2.4.2. Python has a set of language features (such as functions, lists and dictionaries) and a very large set of packages which provide most of the Python functionality. By keeping the language features intact and drastically reducing the number of packages available, the operator is provided with a flexible, although small, scripting language.

The only feature removed from the Python language is unicode support. The only packages provided to the operator are:

- **alc** — The SR OS-provided packages provide access to various ESM objects such as DHCPv4, DHCPv6 or RADIUS packets.
  - **binascii** — Common ASCII decoding like base64.
  - **re** — Regular expression support.
  - **struct** — Parses and manipulates binary strings.
  - **md5** — MD5 message digest algorithm.
- 

## Python Changes

Some changes have been made to Python in order to run on an embedded system:

- No files or sockets can be opened from inside Python scripts.
- No system calls can be made from inside Python scripts nor is the posix package available.
- The maximum recursion depth is fixed to twenty.
- The total amount of dynamic memory available for Python itself and Python scripts is capped at 2MB.
- The size of the script source file must be less than 16KB.

## Python Support in sub-ident-policy

A Python script can be configured in sub-ident-policy to return following ESM attributes:

- sub-id
- sla-profile name
- sub-profile name
- ancp-string

The system will run the Python script configured in the sub-ident-policy against the received DHCPv4 ACK message. This is used as the input of the script. Within the script, the user can set the value with the above ESM attributes.



The alc package contains a DHCP object, and has the following members ([Table 19](#)).

**Table 19: DHCP Object Members**

Name	Read	Write	Class
htype	X		integer
hlen	X		integer
hops	X		integer
flags	X		integer
ciaddr	X		integer
yiaddr	X		integer
siaddr	X		integer
giaddr	X		integer
chaddr	X		string
sname	X		string
file	X		string
options	X		TLV
sub_ident		X	string
sub_profile_string		X	string
sla_profile_string		X	string
ancp_string		X	string
app_profile_string		X	string
category_map_name		X	string
int_dest_id		X	string

The TLV type provides easy access to the value part of a stream of type-length-value variables, as is the case for the DHCP option field. In the example on [page 1470](#), the circuit-ID is accessed as `alc.dhcp.options[82][1]`.

Some DHCP servers do not echo the relay agent option (option 82) when the DHCP message was snooped instead of relayed. For the convenience of the operator, the relay agent option from the request message is returned when `alc.dhcp.options[82]` is called.

## Configuration

As an example consider script `us5.py` on [page 1473](#) which sets the `sub_ident` variable based on the circuit ID of three different DSLAMs:

```
import re
import alcatel
import struct
# ASAM DSLAM circuit ID comes in three flavours:
# FENT string "TLV1: ATM:3/0:100.33"
# GELT octet-stream 0x01010000A0A0A0A0000000640022
# GENT string "ASAM11 atm 1/1/01:100.35"
#
# Script sets output ('subscriber') to 'sub-vpi.vci', e.g.: 'sub-100.33'. circuitid =
str(alcatel.dhcp.options[82][1])
m = re.search(r'(\d+\.\d+)$', circuitid)
if m:
# FENT and GENT
alcatel.dhcp.sub_ident = "sub-" + m.group()
elif len(circuitid) >= 3:
# GELT
# Note: what byte order does GELT use for the VCI?
# Assume network byte (big endian) order for now. vpi = struct.unpack('B', circuitid[-3:-
2])[0]
vci = struct.unpack('>H', circuitid[-2:])[0]
alcatel.dhcp.sub_ident = "sub-%d.%d" % (vpi, vci)
```

Configure the url to this script in a sub-ident-policy as follows:

```
-----
sub-ident-policy "DSLAM" create
description "Parse circuit IDs from different DSLAMs"
primary

script-url "ftp://xxx.xxx.xxx.xx/py/us5.py"
no shutdown
exit
exit
-----
```

And attach this sub-ident-policy to the sub-sla-mgmt from a SAP:

```
A:dut-A>config>service>vpls>sap# info
```

```
-----
dhcp
description "client side"
lease-populate 50

no shutdown
exit
anti-spoof ip-mac
sub-sla-mgmt
sub-ident-policy "DSLAM"
no shutdown
```

```
exit
```

---

Note that DHCP snooping/relaying should be configured properly in order for this to work.

## Operator Debugging

Verbose debug output is sent to debug-trace on compile errors, execution errors, execution output and the exported result variables.

```
A:dut-A>config>subscr-mgmt>sub-ident-pol>primary# script-url "ftp://xxx.xxx.xx.xx/py/
parsefail1.py"
A:dut-A>config>subscr-mgmt>sub-ident-pol>primary# no shutdown

1 2006/07/30 01:17:33.14 UTC MINOR: DEBUG #2001 - Python Compile Error
"Python Compile Error: parsefail1.py
  File "ftp://xxx.xxx.xx.xx/py/parsefail1.py", line 2 def invalid_function():
    ^
IndentationError: expected an indented block
"

A:dut-A>config>subscr-mgmt>sub-ident-pol>primary# script-url "ftp://xxx.xxx.xx.xx/py/
dump.py"

2 2006/07/30 01:24:55.50 UTC MINOR: DEBUG #2001 - Python Output
"Python Output: dump.py htype    = 0
hlen    = 0 hops    = 0 flags    = 0
ciaddr  = '0.0.0.0' yiaddr  = '0.0.0.0' siaddr  = '0.0.0.0' giaddr  = '0.0.0.0' chaddr  =
''
sname   = ''
file    = ''
options = '5\x01\x056\x04\n\x01\x07\n3\x04\x00\x00\x00\xb4\x01\x04\xff\xff\xff
\x00\x1c\x04\n\x02\x02\xffR\x0f\x01\r dut-A|1|1/1/1\xff' "

3 2006/07/30 01:24:55.50 UTC MINOR: DEBUG #2001 - Python Result
"Python Result: dump.py
"

A:dut-A>config>subscr-mgmt>sub-ident-pol>primary# script-url "ftp://xxx.xxx.xx.xx/py/end-
less.py"

4 2006/07/30 01:30:17.27 UTC MINOR: DEBUG #2001 - Python Output
"Python Output: endless.py
"

5 2006/07/30 01:30:17.27 UTC MINOR: DEBUG #2001 - Python Error
"Python Error: endless.py
Traceback (most recent call last):
File "ftp://xxx.xxx.xx.xx/py/endless.py", line 2, in ? FatalError: script interrupted
(timeout)
"
```

Note that all the Python Result events are empty because none of the scripts set any of the output variables.

## Python Scripts

Note that the scripts in this section are test scripts and not scripts which the operator would normally use.

**dump.py** — `from alc import dhcp`

```
def print_field(key, value):
    print '%-8s = %r' % (key, value)

def ipaddr2a(ipaddr):
    return '%d.%d.%d.%d' % (
        (ipaddr & 0xFF000000) >> 24, (ipaddr & 0x00FF0000) >> 16, (ipaddr & 0x0000FF00) >> 8,
        (ipaddr & 0x000000FF))

print_field('htype', dhcp.htype) print_field('hlen', dhcp.hlen) print_
field('hops', dhcp.hops) print_field('flags', dhcp.flags) print_field('ciaddr',
ipaddr2a(dhcp.ciaddr)) print_field('yiaddr', ipaddr2a(dhcp.yiaddr)) print_field('siaddr',
ipaddr2a(dhcp.siaddr)) print_field('giaddr', ipaddr2a(dhcp.giaddr)) print_field('chaddr',
dhcp.chaddr) print_field('sname', dhcp.sname) print_field('file', dhcp.file) print_
field('options', str(dhcp.options))
```

**us5.py** — `import re import alc import struct`

```
# ASAM DSLAM circuit ID comes in three flavors:
#      FENT    string          "TLV1: ATM:3/0:100.33"
#      GELT    octet-stream    0x01010000A0A0A0A00000000640022
#      GENT    string          "ASAM11 atm 1/1/01:100.35"
#
# Script sets output ('subscriber') to 'sub-vpi.vci', e.g.: 'sub-100.33'. circuitid =
str(alc.dhcp.options[82][1])
m = re.search(r'(\d+\.\d+)$', circuitid)
if m:
    # FENT and GENT
    alc.dhcp.sub_ident = "sub-" + m.group()
elif len(circuitid) >= 3:
    # GELT
    # Note: what byte order does GELT use for the VCI?
    # Assume network byte (big endian) order for now. vpi = struct.unpack('B', circuitid[-3:-
2])[0]
    vci = struct.unpack('>H', circuitid[-2:])[0]
    alc.dhcp.sub_ident = "sub-%d.%d" % (vpi, vci)
```

## Sample Python Scripts

This section provides examples to show how the script can be used in the context of Enhanced Subscriber Management.

Note that these scripts are included for informational purposes only. The operator must customize the script to match their own network and processes.

---

### Example

This script uses the IP address assigned by the DHCP server to derive both *sub\_ident* and *sla\_profile\_string*.

Script:

```
1. import alc
2. yiaddr = alc.dhcp.yiaddr
3. # Subscriber ID equals full client IP address.
4. # Note: IP address 10.10.10.10 yields 'sub-168430090'
5. # and not 'sub-10.10.10.10'
6. alc.dhcp.sub_ident = 'sub-' + str(yiaddr)
7. # DHCP server is configured such that the third byte (field) of the IP
8. # address indicates the session Profile ID.
9. alc.dhcp.sla_profile_string = 'sp-' + str((yiaddr & 0x0000FF00) >> 8)
```

Explanation:

**Line 1** — Imports the library “alc” — Library imports can reside anywhere in the script as long as the items are imported before they are used.

**Line 2** — Assigns the decimal value of the host’s IP address to a temporary variable “yiaddr”. Line 6: The text “sub\_” followed by yiaddr is assigned to “sub\_ident” string.

**Line 9** — The text “sp-” followed with the third byte of the IP address is assigned to the “sla-profile” string.

If this script is run, for example, with a DHCP-assigned IP address of:

```
yiaddr = 10.10.0.2
```

The following variables are returned:

```
sub_ident: sub-168427522 (hex = A0A00002 = 10.10.0.2)
sla_ident: sp-0
```

## Example

This script returns the *sub\_profile\_string* and *sla\_profile\_string*, which are coded directly in the Option 82 string.

Script:

```
1. import re
2. import alc
3. # option 82 formatted as follows:
4. # "<subscriber Profile>-<sla-profile>"
5. ident = str(alc.dhcp.options[82][1])
6. alc.dhcp.sub_ident = ident
7. tmp = re.match("(?P<sub>.+)-(?P<sla>.+)", str(ident))
8. alc.dhcp.sub_profile_string = tmp.group("sub")
9. alc.dhcp.sla_profile_string = tmp.group("sla")
```

Explanation:

**Line 1-2** — Import the libraries “re” and “alc”. Library imports can reside anywhere in the script as long as the items are imported before they are used.

**Line 6** — Assigns the full contents of the DHCP Option 82 field to the “sub\_ident” variable.

**Line 7** — Splits the options 82 string into two parts, separated by “-”.

**Line 8** — Assigns the first part of the string to the variable “sub\_profile\_string”.

**Line 9** — Assigns the second part of the string to the variable “sla\_profile\_string”.

If this script is run, for example, with DHCP option field:

```
options = \x52\x0D\x01\x0Bmydsl-video
```

The following variables are returned:

```
sub_ident: mydsl-video
sub_profile_string: mydsl
sla_profile_string: video
```

## Example

This script parses the Option82 “circuit-id” info inserted in the DHCP packet by a DSLAM, and returns the *sub\_ident* string.

Script:

```
1. import re
2. import alc
3. import struct
4. # Alcatel 7300 ASAM circuit ID comes in three flavors:
5. #      FENT    string      "TLV1: ATM:3/0:100.33"
6. #      GELT    octet-stream 0x01010000A0A0A0A0000000640022
7. #      GENT    string      "ASAM11 atm 1/1/01:100.35"
8. #
9. # Script sets output ('subscriber') to 'sub-vpi.vci',
10. # e.g.: 'sub- 100.33'.
11. circuitid = str(alc.dhcp.options[82][1])
12. m = re.search(r'(\d+\.\d+)$', circuitid)
13. if m:
14.     # FENT and GENT
15.     alc.dhcp.sub_ident = "sub-" + m.group()
16. elif len(circuitid) >= 3:
17.     # GELT
18.     # Note: GELT uses network byte (big endian) order for the VCI
19.     vpi = struct.unpack('B', circuitid[-3:-2])[0]
20.     vci = struct.unpack('>H', circuitid[-2:])[0]
21.     alc.dhcp.sub_ident = "sub-%d.%d" % (vpi, vci)
```

Explanation:

**Line 1-2** — Import the libraries “re” and “alc” – Library imports can reside anywhere in the script as long as the items are imported before they are used. Needed if regular expressions are used.

**Line 3** — Imports the “struct” library, needed if regular expressions are used.

**Line 11** — Assigns the contents of the DHCP Option 82 Circuit-ID field to a temporary variable called “circuitid”.

**Line 12** — Parses the “circuitid” and checks for the existence of the regular expression “digit.digit” at the end of the string.

**Line 15** — If found, a string containing the text “sub-” followed by these two digits is assigned to the variable “sub-ident”.

**Line 16** — If not found, and the length of circuit-id is at least 3.

**Line 19** — Parses the “circuitid” and assigns the third-last byte to the temporary variable “vpi”.

**Line 20** — Parses the “circuitid” and assigns the last two bytes to the temporary variable “vci”.

**Line 21** — Assigns a string containing the text “sub-” followed by vpi and vci to the variable “sub-ident”.



If this script is run, for example, with DHCP option field (assigned by an ASAM with FENT card) containing:

```
options = \x52\x16\x01\x14TLV1: ATM:3/0:100.33
```

(in decimal: 80, 22, 1, 20TLV...)

The following variables are returned:

```
sub_ident: sub-100.33
```

If the above script is run, for example, with a DHCP option field (assigned by an ASAM with GELT card) containing:

```
options = \x52\x10\x01\x0E\x01\x01\x00\x00\xA0\xA0\xA0\xA0\x00\x00\x00\x64 \x00\x22
```

(in decimal: 82, 16, 1, 15, 1, 1, 0, 0, 160, 160, 160, 160, 0, 0, 0, 100, 0, 34; corresponding to VPI 100, VCI 34)

Python returns the following variables:

```
sub_ident: sub-100.34
```

If the above script is run, for example, with a DHCP option field (assigned by an ASAM with GENT card) containing:

```
options = \x52\x1A\x01\x18ASAM11 atm 1/1/01:100.35
```

The following variables are returned:

```
sub_ident: sub-100.35
```

## Limitations

**'%' operator** — While %f is supported, %g and %e are not supported.

**Floating Point Arithmetic** — The floating point arithmetic precision on the box is less than the precision required by the regression suites of Python. For example, `pow(2., 30)` equals to

`1024.*1024.*1024.` until five numbers after the point instead of seven and `sqrt(9)` equals to 3. for the first seven numbers after the point.

Using the round operator fixes these problems. For example, `round(pow(2., 30))` equals `round(1024.*1024.*1024.)` and `round(sqrt(9))` equals 3.

## RADIUS Script Policy Overview

Python scripts for RADIUS AAA packets support manipulation in subscriber management application. This feature is supported on 7750 SRs and 7450 ESSes in mixed mode. A Python script can be executed in following cases:

- Before the system sends an access-request packet.
- After the system receives an access-accept packet.
- After the system receives an CoA-request packet.
- Before the system sends an accounting-request packet.

The input of the script is the corresponding original packet; and the output of packet will be used as the new corresponding packet for further ESM AAA process.

The **radius-script-policy** contains URLs of a primary and optionally a secondary Python script, which could be a local CF file path or a FTP URL. The configured radius-script-policy could be used in different ESM policies like authentication-policy or radius-accounting-policy.

The following operations are supported within the script:

- Obtain the value of an existing attribute or VSA.
- Modify the value of an existing attribute or VSA.
- Add a new attribute or VSA.
- Remove an existing attribute or VSA.

Note that the following RADIUS attributes or VSA are read-only to Python script:

- Message-Authenticator
- Alc-LI-Action
- Alc-LI-Direction
- Alc-LI-Destination
- Alc-LI-FC
- Alc-LI-Intercept-Id
- Alc-LI-Session-Id

Since R12.0R1, users should use a Python policy (instead of a RADIUS script policy) for RADIUS packet manipulation.

## Python RADIUS API

The following new Python objects, `alc.radius.attributes`, have the following methods:

- `drop()`: drop the packet
  - `header()`: Return the a dictionary object includes RADIUS header information
  - `get(type)`: Return the first attribute with specified type as a string.
  - `getTuple(type)`: The same as above but returns a tuple of strings.
  - `getVSA(vendor, type)`: Return the first VSA as a string.
  - `getVSATuple(vendor, type)`: The same as above but returns a tuple of strings.
  - `set(type, value)`: Set the specified attribute to the value. The value must be either a string or a tuple of strings.
  - `setVSA(vendor, type, value)`: Set the specified VSA to the value. The value must be either a string or a tuple of strings.
  - `clear(type)`: Remove the specified attribute.
  - `clearVSA(vendor, type)`: Remove the specified VSA.
- 

## Sample Script

From `alc` import `radius`

#1. Get the value of an existing Attribute

```
Username=radius.attributes.get(1)
```

#2. Modify an existing attribute

```
radius.attributes.set(1, 'Tom')
```

#3. Remove an existing attribute

```
radius.attributes.clear(1)
```

#4. Add a new attribute

```
radius.attributes.set(126,"WIFI-operator")
```

## Python Policy Overview

The Python policy represents a general framework to support all existing and new python features. A Python policy allows users to configure a Python script for specified ESM packet type (such as DHCP, RADIUS, etc.) in a specified direction (ingress/egress). The system will execute the configured script when sending or receiving the specified type of packet.

Within the script, the corresponding original packet will be used as input. The user can use the system-provided API to manipulate the input packet (such as add/change/remove option/attribute) and the changed packet is the output for further ESM processing. And in case of a DHCP transaction cache, the script could also return ESM attributes.

Python policies support following ESM packet types and application:

- RADIUS
- DHCPv4
- DHCPv6
- DHCP Transaction Cache
- Diameter
- Python cache

The following is an example configuration on a specified group interface. The system will execute `cfl:/dhcpv4.py` after received DHCPv4 discovery and before system forward DHCPv4 request packet.

```
config>python# info
-----
python-script "dhcpv4" create
    primary-url "cfl:/dhcpv4.py"
    no shutdown
exit
python-policy "dhcp" create
    dhcp discover direction ingress script "dhcpv4"
    dhcp request direction egress script "dhcpv4"
exit
-----
config>service>vprn>sub-if>grp-if>dhcp# info
-----
python-policy "dhcp"
server 9.9.9.9
lease-populate 100
gi-address 192.168.100.1
no shutdown
-----
```

## Python Policy – RADIUS API

The RADIUS API in Python policy uses the same API of the radius-script-policy.

## Python Policy – DHCPv4 API

- The system will provide a Python object for input DHCPv4 packet: `alc.dhcpv4`.
- `alc.dhcpv4` has following attributes to represent the DHCPv4 header fields:

**Table 20: alc.dhcpv4 Attributes**

Class Attrs	DHCPv4 Header Field	Access
<code>alc.dhcpv4.pkt_len</code>	int, Total length of original DHCPv4 packet(UDP/IP header excluded, including pad option) in bytes	read
<code>alc.dhcpv4.pkt_netlen</code>	int, Total length of original DHCPv4 packet(UDP/IP header excluded, pad option in the “options” field excluded) in bytes	read
<code>alc.dhcpv4.op</code>	op	read
<code>alc.dhcpv4.htype</code>	htype	read/write
<code>alc.dhcpv4.hlen</code>	hlen	read/write
<code>alc.dhcpv4.hops</code>	hops	read/write
<code>alc.dhcpv4.xid</code>	xid	read
<code>alc.dhcpv4.secs</code>	secs	read/write
<code>alc.dhcpv4.flags</code>	flags	read/write
<code>alc.dhcpv4.ciaddr</code>	ciaddr	read/write
<code>alc.dhcpv4.yiaddr</code>	yiaddr	read/write
<code>alc.dhcpv4.siaddr</code>	siaddr	read/write
<code>alc.dhcpv4.giaddr</code>	giaddr	read/write
<code>alc.dhcpv4.chaddr</code>	chaddr	read/write
<code>alc.dhcpv4.sname</code>	sname	read/write
<code>alc.dhcpv4.file</code>	file	read/write

All attributes, except `alc.dhcpv4.pkt_len` and `alc.dhcpv4.pkt_netlen`, are string with value of the actual bytes in the header.

The following is a list all functions of `alc.dhcpv4`:

- `alc.dhcpv4.drop()`
- `alc.dhcpv4.getOptionList()`
- `alc.dhcpv4.pad(min_size=300)`
- `alc.dhcpv4.get(op_code)`
- `alc.dhcpv4.set(op_code, valTuple)`
- `alc.dhcpv4.clear(op_code)`
- `alc.dhcpv4.get_relayagent() / alc.dhcpv4.set_relayagent(D4OL)`
- `alc.dhcpv4.get_vendorspecific() / alc.dhcpv4.set_vendorspecific (D4OL)`

DHCPv4 allows using `sname` and `file` fields to store options. However all DHCPv4 functions will only operate with the “options” field. If a customer wants to manipulate options in the `sname/file` field, they need to do the parsing work in the script. (extended `string.tlvxy` method could help here)

- `alc.dhcpv4.drop()`: The system will drop the result packet
- `alc.dhcpv4.getOptionList()`: Returns a tuple that includes the option-code of the existing top level DHCPv4 options in the packet.
  - The order of the element in the tuple is as same as the options that appear in the packet.
  - If there are multiple instances of the same option, then each instance is one element in the tuple.
  - Pad option(0) is excluded.
  - End option(255) is included
  - Example: A DHCP discovery packet with `msg-type/lease-time/request-addr/parameter-request-list/agent-info/end` will return (53,51,50,55,82,255)
- `alc.dhcpv4.pad(min_size=300)`: This function will pad the resulting DHCPv4 packet to the specified `min_size` with pad option(0) after executing the whole script. Padding will not be added if the result packet is already `>=min_size`. The default value of `min_size` is 300. Although not defined in DHCPv4 RFC, many DHCPv4 implementation has a minimal length requirement of 300 bytes. So this function could pad the result packet to the specified `min_size`.

- `alc.dhcpv4.get(op_code)`: Returns a tuple that includes all instances of the specified top level option as a string. The value of this string is the exact bytes of the option as it appears in packet(excludes option-code and option-len).
  - If the specified option does not exist, then the function will return ()
  - If the certain instance of specified option does not have the value (len=0 or doesn't have len and value part), then the function will return "" for that instance in the tuple.
  - Example: There is an address lease time option(51) in the packet, with value 60, then `alc.dhcpv4.get(51)` should return: (`'\x00\x00\x00\x3c'`,)
- `alc.dhcpv4.set(op-code,valTuple)`: This function will in fact remove the all existing instances of specified top level option and insert a list of new options. Each element in `valTuple` is a string, representing one instance of the new option to be inserted; For each new option, the option-code is specified in `op-code`. The option-len is the length of the element. The rest of option is the element itself.
  - Example: To insert an address lease time option(51) in the packet, with the value 60; use `alc.dhcpv4.set(51, ('\x00\x00\x00\x3c',))`
- `alc.dhcpv4.clear(op-code)`: This function will remove the all existing instances of specified top level option.
- Although `alc.dhcpv4.get()` and `alc.dhcpv4.set()` provide a generic way to manipulate DHCPv4 top level options, but some DHCPv4 options have a complex/hierarchical structure like option82 and option43. To provide a friendly access to these kinds of options, the system provides the following options' specific functions:
  - `alc.dhcpv4.get_relayagent()` / `alc.dhcpv4.set_relayagent(D4OL)`
  - `alc.dhcpv4.get_vendorspecific()` / `alc.dhcpv4.set_vendorspecific(D4OL)`

All `alc.dhcpv4.get_XXX()` will return a data structure:"D4OL" (DHCPv4 Option List)

→ D4OL is a list. Each element in the list represents an instance of that option. For example, if there are 3 option82 in the "options" field of packet, then `get_relayagent()` will return a list of 3 elements. Each element represents one instance of the option in the packet.

→ Each element in D4OL is a dict (called dict as "D4O" in this example):

- The key of D4O is the sub-option- code. The value is a list of strings of sub-option-value of all instance of the sub-option.

All `alc.dhcpv4.set_XX(OPDL)` will accept D4OL as the parameter. Remove all existing instances of the corresponding options and then insert the new options represented by specified D4OL.



Examples:

For a packet with an option82 like following

```
Option: (82) Agent Information Option
  Length: 22
  Option (82) Suboption: (1) Agent Circuit ID
    Length: 8
    Agent Circuit ID: 4a616e737656e73
  Option 82 Suboption: (2) Agent Remote ID
    Length: 10
    Agent Remote ID 62617369632364617461
```

The option-data is (hex formatted)

“01:08:4a:61:6e:73:73:65:6e:73:02:0a:62:61:73:69:63:23:64:61:74:61”

The following is an example script:

```
import alc

option82_list=alc.dhcpv4.get_relayagent()
#option82_list will be
[{1:['\x4a\x61\x6e:73\x73\x65\x6e\x73',],2:['\x62\x61\x73\x69\x63\x23\x64\x61\x74\x61',]}
,]
Option82_list[0][2][0]='basic#video' #change remote-id to 'basic#video'
alc.dhcpv4.set_relayagent(option82_list)#update the option82
```

## Python Policy – DHCPv6 API

- The system provides a Python object for input DHCPv6 packet: `alc.dhcpv6`
- `alc.dhcpv6` has following attributes to represent the DHCPv6 header fields:

**Table 21: DHCPv6 Header Fields**

Class Attrs	DHCPv6 Header Field	Client/ Server Msg	Relay Msg	Access
<code>alc.dhcpv6.pkt_len</code>	int, Total length of original DHCPv4 packet(UDP/IP header excluded) in bytes	•	•	Read
<code>alc.dhcpv6.msg_type</code>	msg-type	•	•	Read
<code>alc.dhcpv6.transaction_id</code>	transaction-id	•		Read
<code>alc.dhcpv6.hop_count</code>	hop-count		•	read/write
<code>alc.dhcpv6.link_addr</code>	link-address		•	read/write
<code>alc.dhcpv6.peer_addr</code>	peer-address		•	read/write

All header fields (as the attribute of `alc.dhcpv6` class) are strings(except `pkt_len` ) with exact bytes as it appears in the packet.

If certain attribute does not exist in the given msg-type, for example if the `link_attr` does not exist in client/server message(C/S msg), then its value should be `None`.

- The following is a list of all functions in the class:
  - `alc.dhcpv6.drop()`
  - `alc.dhcpv6.getOptionList()`
  - `alc.dhcpv6.get(op-code)`
  - `alc.dhcpv6.set(op-code,valTuple)`
  - `alc.dhcpv6.clear(op-code)`
  - `alc.dhcpv6.get_iana() / alc.dhcpv6.set_iana(OPDL)`
  - `alc.dhcpv6.get_iata() / alc.dhcpv6.set_iata(OPDL)`
  - `alc.dhcpv6.get_vendoropts() / alc.dhcpv6.set_vendoropts(OPDL)`
  - `alc.dhcpv6.get_iapd() / alc.dhcpv6.set_iapd(OPDL)`
  - `alc.dhcpv6.get_relaymsg()`
  - `alc.dhcpv6.set_relaymsg(packet)`
- `alc.dhcpv6.drop()`: The system will drop the resulting packet.

- `alc.dhcpv6. getOptionList()`: Returns a tuple that includes the option-code of the existing top level DHCPv6 options in the packet. The order of the element in the tuple is as same as the options appear in the packet. If there are multiple instances of same option, then each instance is one element in the tuple. For example:
  - A C/S Msg with Elapsed Time/Client Identifier/IANA/FQDN/Vendor Class/Option Request will return (8,1,3,39,16,6).
  - A Relay Msg with Relay Message option only will return (9)
- `alc.dhcpv6. get(op-code)`: Returns a tuple that includes all instances of the specified top level option as string. The value of this string is the exact bytes of the option as it appears in packet(excludes option-code and option-len). If the specified option does not exist in the input packet, then it will return ().

Examples:

- If there is a Status Code option in the packet, status-code 0 and status-msg:”Address Assigned”; then `alc.dhcpv6.get(13)` should return: (`‘\x00\x00Address Assigned’`,)
- `alc.dhcpv6. set(op-code,valTuple)`: This function will remove the all existing instances of the specified top level option and insert a list of new options. Each element in `valTuple` is a string, representing one instance of the new option to be inserted. For each new option, the option-code is specified in `op-code`, the option-len is the length of the element, reset of option is the element itself.
  - To insert a Status Code options with status-code 0 and status-msg:”Address Assigned”; use `alc.dhcpv6.set(13, (‘\x00\x00Address Assigned’,))`
- `alc.dhcpv6. clear(op-code)`: This function will remove the all existing instances of specified top level option.
- Although `alc.dhcpv6.get()` and `alc.dhcpv6.set()` provide a generic way to manipulate DHCPv6 top level options, but some DHCPv6 options have more complex/hierarchical structure like IA\_NA/IA\_TA, etc. To provide a friendly access to these kinds of options, the system provides the following options specific functions:
  - `alc.dhcpv6.get_iana()` / `alc.dhcpv6.set_iana(OPDL)`
  - `alc.dhcpv6.get_iata()` / `alc.dhcpv6.set_iata(OPDL)`
  - `alc.dhcpv6.get_vendoropts()` / `alc.dhcpv6.set_vendoropts(OPDL)`
  - `alc.dhcpv6.get_iapd()` / `alc.dhcpv6.set_iapd(OPDL)`

All `alc.dhcpv6.get_XXX()` will return a data structure:“OPDL” (Option Data structure List)

- OPDL is a list. Each element in the list represents an instance of that option. For example, if there are 3 IANA in the packet, then `get_iana()` will return a list of 3 elements, each element represent one instance of IANA option in the packet.
- Each element in OPDL is a list, referred to as “OPD” in this list.
  - Each element in OPD represent one field in the option(option-code and option-len are not included), the order of the element is as same as the fields appear in the option
  - For field that could be parsed into sub-option by RFC, then the element is a dict, the key of this dict is the sub-option type, if sub-option is one of following supported-sub-option, the value to the key is a sub-option\_OPDL represent the list of that specific sub-option
    - IAADDR(5)
    - Status Code(13)
    - IAPREFIX(26)

Else, if the sub-option is **not** one of above, then the value to the key is a list of string of sub-option-data, each string represent one instance of the sub-option.

- The structure of sub-option\_OPDL of IAADDR is: `[[v6_addr,prefer_lifetime, valid_lifetime,sub-option_OPDL], etc.]`
- The structure of sub-option\_OPDL of Status Code is: `[[status-code,status-msg], etc.]`
- The structure of sub-option\_OPDL of IAPREFIX is: `[[prefer_lifetime,valid_lifetime,prefix-len,v6prefix,sub-option_OPDL],..]`
- For the field (by RFC definition) could be parsed into sub-options, but it does not actually exist, then the dict will be empty `{}`
- For field that cannot be parsed into sub-option by RFC, the element is a string of exact bytes of that field

All `alc.dhcpv6.set_XX(OPDL)` will accept an OPDL as the parameter. Remove all existing instances of the corresponding options and then insert new options represented by the specified OPDL.

- `alc.dhcpv6.get_iana()/alc.dhcpv6.get_iana(OPDL)`
- The general OPDL structure for these two functions is: `[[IAID_val,T1_val,T2_val, sub-option_dict]]`
- The structure of sub-option\_dict is: `{sub-option-type:sub-option_val}`
- If sub-option is supported-sub-option, then sub-option\_val is a sub-option\_OPDL
- For all other sub-options, the sub-option\_val is a list of string of sub-option-data

Examples:

For a packet with an IANA option like following:

```
-----
  Identity Association for Non-temporary Address
    Option: Identity Association for Non-temporary Address (3)
    Length: 40
    Value: 0ff0def10002a30000043800000500182001055860450047...
    IAID: 0ff0def1
    T1: 172800
    T2: 276480
  IA Address: 2001:558:6045:47:45cc:d9f2:5727:eae0
    Option: IA Address (5)
    Length: 24
    Value: 200105586045004745ccd9f25727eae00005460000054600
    IPv6 address: 2001:558:6045:47:45cc:d9f2:5727:eae0
    Preferred lifetime: 345600
    valid lifetime: 345600
```

The option-data is (hex formatted)

```
"0f:f0:de:f1:00:02:a3:00:00:04:38:00:00:05:00:18:20:01:05:58:60:45:00:47:45:cc:d9:f2:57:27:ea:e0:00:05:46:00:00:05:46:00"
```

The following is an example script:

```
import alc

iana_list=alc.dhcpv6.get_iana()
#iana_list will be
[['\x0f\x00\xde\xfa', '\x00\x02\xa3\x00', '\x00\x04\x38\x00', {5: [['\x20\x01\x05\x58\x60\x45\x00\x47\x45\xcc\x00', '\x00\x05\x46\x00', '\x00\x05\x46\x00', {}]]}],
['\x00\x00\x04\xb0', '\x00\x05\x46\x00', '\x00\x05\x46\x00', {}]]]
iana_list[0][1]='\x00\x00\x04\xb0' #change T1 to 1200
alc.dhcpv6.set_iana(iana_list)#update the iana
```

- `alc.dhcpv6.get_iata()/alc.dhcpv6.get_iata(OPDL)`
  - The general OPDL structure for these two functions is: `[[IAID_val, sub-option_dict]]`
  - The structure of `sub-option_dict` is: `{sub-option-type:sub-option_val}`
  - If sub-option is supported-sub-option, then `sub-option_val` is a `sub-option_OPDL`
  - For all other sub-options, the `sub-option_val` is a list of string of sub-option-data

Examples: These two function are very similar with IANA, so the examples are skipped here.

- `alc.dhcpv6.get_iapd()/alc.dhcpv6.get_iapd(OPDL)`
- The general OPDL structure for these two functions is: `[[IAID_val, T1_val, T2_val, sub-option_dict]]`
- The structure of `sub-option_dict` is: `{sub-option-type:sub-option_val}`
- If sub-option is supported-sub-option, then `sub-option_val` is a `sub-option_OPDL`

- For all other sub-options, the sub-option\_val is a list of string of sub-option-data

Examples: For a packet with IA\_PD like following:

```

Identity Association for Prefix Delegation
Option: Identity Association for Prefix Delegation (25)
Length: 41
Value: 000000010000070800000b40001a001900000e1000015180...
IAID: 00000001
T1: 1800
T2: 2880
IA Prefix
Option: IA Prefix (26)
Length: 25
Value: 00000e10000151803820010db8000200000000000000000...
Preferred lifetime: 3600
Valid lifetime: 86400
Prefix length: 56
Prefix address: 2001:db8:2::
    
```

The option-data is (hex formatted)

```
"00:00:00:01:00:00:07:08:00:00:0b:40:00:1a:00:19:00:00:0e:10:00:01:51:80:38:20:01:0d:b8:00:02:00:00:00:00:00:00:00:00:00:00"
```

Following is an example script:

```

import alc

iapd_list=alc.dhcpv6.get_iapd()
#iapd_list will be
[['\x00\x00\x00\x01', '\x00\x00\x07\x08', '\x00\x00\x0b\x40', {26: [['\x00\x00\x0e\x10', '\x00\x01\x51\x80', '\x38', '\x20\x01\x0d\xb8\x00\x02\x00\x00\x00\x00\x00\x00\x00\x00', {'}]]}]]
iapd_list[0][1]='\x00\x00\x04\xb0' #change T1 to 1200
alc.dhcpv6.set_iapd(iapd_list)#update the iapd
    
```

- alc.dhcpv6.get\_vendoropts()/alc.dhcpv6.get\_vendoropts (OPDL)
  - The general OPDL structure for these two functions is: [[enterpriseid\_val, sub-option\_dict]]
  - The structure of sub-option\_dict is: {sub-option-type:sub-option\_val}
  - If sub-option is supported-sub-option, then sub-option\_val is a sub-option\_OPDL
  - For all other sub-options, the sub-option\_val is a list of string of sub-option-data

Examples: For a packet with vendor options like following:

```

❑ Vendor-specific Information
  Option: Vendor-specific Information (17)
  Length: 40
  Value: 0000197f0001000969612d6e615f3030310002000969612d...
  Enterprise ID: Panthera Networks, Inc. (6527)
❑ option
  option code: 1
  option length: 9
  option-data
❑ option
  option code: 2
  option length: 9
  option-data
❑ option
  option code: 3
  option length: 1
  option-data
❑ option
  option code: 4
  option length: 1
  option-data

```

The option-data is (hex formatted)

```

"00:00:19:7f:00:01:00:09:69:61:2d:6e:61:5f:30:30:31:00:02:00:09:69:61:2d:70:64:5f:30:30:31:00:03:00:01:38:00:04:00:01:40"

```

The following is an example script:

```

import alc

vendoropts_list=alc.dhcpv6.get_vendoropts()
# vendoropts_list will be [['\x00\x00\x19\x7f',
{1:['\x69\x61\x2d\x6e\x61\x5f\x30\x30\x31'],2:['\x69\x61\x2d\x70\x64\x5f\x30\x30\x31'],3:['\x38'], 4:['\x40']}]
iapid_list[0][1][4][0]='\x60' #change sub-option 4's value to 0x60
alc.dhcpv6.set_vendoropts(vendoropts_list)#update the vendor options

```

- For DHCPv6 relay message, the “Relay Message” option embedded a full DHCPv6 packet and the embedded packet could itself have a “Relay Message” option which embedded another DHCPv6 packet.

To provide direct access to embedded DHCPv6 packet, the system provides following functions:

```

→ alc.dhcpv6.get_relaymsg()
→ alc.dhcpv6.set_relaymsg(packet)

```

- `alc.dhcpv6.get_relaymsg()`: This function will return a populated `alc.dhcpv6` object, which means the returned object was initialized with the DHCPv6 packet embedded in “Relay Message” option as the input.
- `alc.dhcpv6.set_relaymsg(packet)`: This function will accept an `alc.dhcpv6` object as a parameter. This object will replace existing “Relay Message” option.

- Example-1 script for single relay message:

```
import alc
#input packet is a relay-reply msg
embed_dhcpv6_packet=alc.dhcpv6.get_relaymsg()
iana_list=embed_dhcpv6_packet.get_iana()
iana_list[0][1]='\x00\x00\x04\xb0' #change T1 to 1200
embed_dhcpv6_packet.set_iana(iana_list)#update the iana of the embedded packet
alc.dhcpv6.set_relaymsg(embed_dhcpv6_packet)#update the Relay Message option
```

- Example-2 script for double relay message (relay of relay):

```
import alc
#input packet is a relay-reply msg
embed_lv1_packet=alc.dhcpv6.get_relaymsg() #get the level=1 embedded packet
embed_lv2_packet= embed_lv1_packet.get_relaymsg()#get level-2 packet embedded in level-1
packet
iana_list=embed_dhcpv6_packet.get_iana()
iana_list[0][1]='\x00\x00\x04\xb0' #change T1 to 1200
embed_dhcpv6_packet.set_iana(iana_list)#update the iana
embed_lv1_packet.set_relaymsg(embed_lv2_packet)#update the Relay Message option of lv1 msg
alc.dhcpv6.set_relaymsg(embed_lv1_packet)#update the Relay Message option of the top level
msg
```



## Python Policy – Diameter API

The `alc.diameter` Python module provides an API for Diameter message manipulation.

Terminology used in the API description:

- **top-level-AVP** — AVP appearing at the top level in a Diameter message, i.e. not embedded in the Data field of a grouped AVP
- **embedded-AVP** — AVP embedded in the Data field of a grouped AVP. An embedded AVP can be a grouped AVP. This is called nesting.
- **AVP-tuple** — Python tuple with following values:  
(AVP code, vendor id, flags)  
AVP code : integer, AVP header field  
vendor id : integer, AVP header field  
flags : string, AVP header field
- **AVP-value-tuple** — Python tuple with following values:  
(flags, data)  
flags: string, AVP header field  
data: string, AVP data field
- **AVP-key-tuple** — Python tuple with following values:  
(AVP code, vendor id)  
AVP code: integer, AVP header field  
vendor id : integer, AVP header field
- **grouped-AVP-value-tuple** — Python tuple with following values:  
(flags, grouped-AVP-dictionary)  
flags: string, AVP header field
- **grouped-AVP-dictionary** - Python dictionary with following key:value pairs:  
{AVP-key-tuple : [AVP-value-tuple or grouped-AVP-value-tuple, ...], ... }  
key = AVP-key-tuple  
value = list of AVP-value-tuples or list of grouped-AVP-value-tuples
- **grouped-AVP-decode-tuple** — Python tuple with following values:  
(AVP-key-tuple, ...)  
tuple of AVP-key-tuples
- **AVP-order-tuple** — Python tuple with following values:  
(AVP-key-tuple, ...)  
tuple of AVP-key-tuples

Table 22 displays attributes available in **alc.diameter** module providing access to the Diameter message header:

**Table 22: Diameter Message Header alc.diameter Attributes**

Attribute	Description	Type	Access
application_id	Diameter message header field	string	Read/Write
code	Diameter message header field	string	Read/Write
end_to_end_id	Diameter message header field	string	Read/Write
flags	Diameter message header field	string	Read/Write
hop_by_hop_id	Diameter message header field	string	Read/Write
msg_length	Diameter message header field. The value is the message length of the original diameter message.	string	Read
version	Diameter message header field	string	Read/Write

Table 23 displays methods available in **alc.diameter** module providing message and AVP manipulation functionality:

**Table 23: Message and AVP Manipulation Functionality alc.diameter Methods**

Method	Description
clear_avps(AVP code, vendor id) AVP code, vendor id = top-level-AVP	Remove all instances of the specified AVP from the message. Applies to top-level-AVP's only. If the specified AVP is not present, no python error is generated. Vendor id value zero matches top-level-AVP's without Vendor Id field. Return value: None For example: diameter.clear_avps(256, 12645)
drop()	Drop the Diameter message. Packet is consumed at TCP level (ack send). A drop will trigger retransmits on Diameter level. Return value: None For example: diameter.drop()

**Table 23: Message and AVP Manipulation Functionality alc.diameter Methods (Continued)**

Method	Description
<code>get_avps(AVP code, vendor id)</code> AVP code, vendor id = top-level-AVP	<p>Returns a list of AVP-value-tuples. Each AVP-value-tuple represents an instance of the specified AVP in the message.</p> <p>Applies to top-level-AVP's only. The position in the list corresponds with the position of the AVP instance in the message at that stage in the script. When executed before any clear or set AVP method, the list order corresponds with the AVP order in the received message. If the specified AVP is a grouped AVP, then the data will contain all the embedded-AVP's. An empty list is returned if the specified AVP is not present. Vendor id value zero matches top-level-AVP's without Vendor Id field. For example:</p> <pre>diameter.get_avps(263, 0) [('@', 'bng.alcatel-lucent.com;1398156449;28')]</pre>
<code>get_avps_list()</code>	<p>Returns a list of AVP-tuples. Each AVP-tuple represents an instance of an AVP in the message. Applies to top-level-AVP's only. The position in the list corresponds with the position of the AVP in the message at that stage in the script. When executed before any clear or set AVP method, the list order corresponds with the AVP order in the received message. When multiple instances of an AVP are present in the message, then there will be multiple instances in the list. The Vendor Id has value zero when not present. Grouped AVPs cannot be distinguished from other AVPs in the list. For example:</p> <pre>diameter.get_avps_list() [(263, 0, '@'), (264, 0, '@'), (296, 0, '@'), (258, 0, '@'), (416, 0, '@'), (415, 0, '@'), (268, 0, '@'), (55, 0, '@'), (456, 0, '@'), (456, 0, '@'), (456, 0, '@'), (293, 0, '@'), (256, 12645, 'x80')]</pre>

**Table 23: Message and AVP Manipulation Functionality alc.diameter Methods (Continued)**

Method	Description
get_grouped_avps(AVP code, vendor_id, grouped-AVP-decode-tuple) AVP code, vendor id = top-level-AVP	<p>Returns a list of grouped-AVP-value-tuples with each grouped-AVP-dictionary entry representing an embedded AVP. Each grouped-AVP-value-tuple represents an instance of the specified AVP in the message. Applies to top-level-AVP's of type grouped only. The position in the list corresponds with the position of the grouped AVP instance in the message at that stage in the script. When executed before any clear or set AVP method, the list order corresponds with the AVP order in the received message.</p> <p>If the grouped-AVP-decode-tuple is empty, only the specified top-level-AVP is expanded in a grouped-AVP-value-tuple, with each grouped-AVP-dictionary entry representing an embedded AVP and all dictionary values of type "list of AVP-value-tuples"</p> <p>To expand nested AVPs (grouped AVPs embedded in a grouped AVP), the grouped top-level-AVP and grouped embedded-AVP to expand must be added to the grouped-AVP-decode-tuple. All grouped AVP's in the grouped-AVP-decode-tuple will be expanded in a list of grouped-AVP-value-tuples provided that their embedding AVP is also in the list. The position of the embedded AVPs in the grouped-AVP-dictionary does not correspond with the position in the grouped AVP.</p> <p>If the specified top-level-AVP is not a grouped AVP, then a Python error is generated: 'ValueError: malformed diameter message'.</p> <p>For example: To expand the Multiple Services Credit Control (456) grouped top level AVP:</p> <pre>diameter.get_grouped_avps(456,0,())</pre> <pre>[('@', {(432, 0): [('@', '\x00\x00\x00\x01')], (431, 0): [('@', '\x00\x00\x01\xa4@\x00\x00\x0c\x00\x00\x00d')], (448, 0): [('@', '\x00\x00\x04\xb0')], (268, 0): [('@', '\x00\x00\x07\xd1')}]}, ('@', {(432, 0): [('@', '\x00\x00\x00\x02')], (431, 0): [('@', '\x00\x00\x01\xa4@\x00\x00\x0c\x00\x00\x03\x84')], (448, 0): [('@', '\x00\x00\x04\xb0')], (268, 0): [('@', '\x00\x00\x07\xd1')}]}, ('@', {(432, 0): [('@', '\x00\x00\x00\x03')], (431, 0): [('@', '\x00\x00\x01\xa4@\x00\x00\x0c\x00\x00\x00&lt;')], (448, 0): [('@', '\x00\x00\x04\xb0')], (268, 0): [('@', '\x00\x00\x07\xd1')}]})]</pre>

**Table 23: Message and AVP Manipulation Functionality alc.diameter Methods (Continued)**

Method	Description
	<p>To expand the nested Granted-Service-Unit AVP (code 431) in the grouped Multiple Services Credit Control top level AVP (code 456):</p> <pre>diameter.get_grouped_avps(456,0,((456,0),(431,0))) [('@', {(432, 0): [('@', '\x00\x00\x00\x01')], (431, 0): [('@', {(420, 0): [('@', '\x00\x00\x00d')]}), (448, 0): [('@', '\x00\x00\x04\xb0')], (268, 0): [('@', '\x00\x00\x07\xd1')]}), ('@', {(432, 0): [('@', '\x00\x00\x00\x02')], (431, 0): [('@', {(420, 0): [('@', '\x00\x00\x03\x84')]}), (448, 0): [('@', '\x00\x00\x04\xb0')], (268, 0): [('@', '\x00\x00\x07\xd1')]}), ('@', {(432, 0): [('@', '\x00\x00\x00\x03')], (431, 0): [('@', {(420, 0): [('@', '\x00\x00\x00&lt;')]}), (448, 0): [('@', '\x00\x00\x04\xb0')], (268, 0): [('@', '\x00\x00\x07\xd1')]}))]</pre>
set_avps(AVP code, vendor id, list of AVP-value-tuples) AVP code, vendor id = top-level-AVP	<p>Remove all instances of the specified top-level-AVP from the message. For each entry in the AVP-value-tuple list, a top-level-AVP instance is inserted.</p> <p>If the specified vendor id value is zero, then no vendor id field is inserted and setting the Vendor-Specific bit in the flags field of the AVP value tuple will then result in a Python error: “ValueError: no vendor ID but vendor flag set”.</p> <p>If the specified vendor id value is non-zero, then a vendor id field is inserted and not setting the Vendor-Specific bit in the flags field of the AVP value tuple will result in a Python error: “ValueError: vendor ID but vendor flag not set”.</p> <p>Padding between AVPs, AVP length and Diameter message length are adapted accordingly by the system.</p> <p>Return value is None.</p> <p>For example:</p> <pre>diameter.set_avps(461,0,[('\x40', 'Python-1'), ('\x40', 'Python-2')])</pre>
set_fixed_position_avps(AVP-order-tuple)	<p>Put the specified top-level-AVPs at the beginning of the message in the order as specified in the AVP-order-tuple.</p> <p>This method overrides the order of the top-level-AVPs in the resulting Diameter message. If for example the session-id AVP must appear as first in the message, then the corresponding AVP-key-tuple must be included in the first position of the AVP-order-tuple.</p> <p>AVPs not present in the message but specified in the AVP-order-tuple are ignored.</p>

**Table 23: Message and AVP Manipulation Functionality alc.diameter Methods (Continued)**

Method	Description
	<p>AVPs present in the message and not specified in the AVP-order-tuple will be included in the final message after the AVPs listed in the AVP-order-tuple. The order is deterministic but implementation specific.</p> <p>This method can appear at any point in the script. The last call will override the previous one.</p> <p>From a black box viewpoint, this method is executed at the end of the script: the result of the call is not reflected in the list returned by a subsequent <code>get_avp_list()</code> call.</p> <p>Return value: None</p> <p>For example:</p> <pre>diameter.set_fixed_position_avps(((263,0), (264,0), (296,0), (268,0)</pre>
<code>set_grouped_avps(AVP code, vendor id, list of grouped-AVP-value-tuples)</code> AVP code, vendor id = top-level-AVP	<p>Remove all instances of the specified grouped top-level-AVP from the message. For each entry in the grouped-AVP-value-tuples list, a grouped top-level-AVP instance is inserted. The order of the embedded-AVPs in the grouped AVP cannot be specified. If the specified vendor id value is zero, then no vendor id field is inserted and setting the Vendor-Specific bit in the flags field of the AVP value tuple will then result in a Python error: “ValueError: no vendor ID but vendor flag set”. If the specified vendor id value is non-zero, then a vendor id field is inserted and not setting the Vendor-Specific bit in the flags field of the AVP value tuple will result in a Python error: “ValueError: vendor ID but vendor flag not set”.</p> <p>Padding between AVPs, AVP length and Diameter message length are adapted accordingly.</p> <p>Return value is None.</p> <p>For example:</p> <pre>diameter.set_grouped_avps(456,0,[('@', {(432, 0): [('@', '\x00\x00\x00\x01')], (431, 0): [('@', {(420, 0): [('@', '\x00\x00\x00\x2b')}]}, (448, 0): [('@', '\x00\x00\x04\xb0')], (268, 0): [('@', '\x00\x00\x07\xd1')}]}, (@, {(432, 0): [('@', '\x00\x00\x00\x03')], (431, 0): [('@', {(420, 0): [('@', '\x00\x00\x00\x53')}]}, (448, 0): [('@', '\x00\x00\x04\xb0')], (268, 0): [('@', '\x00\x00\x07\xd1')}]})]</pre>

To enable Diameter message manipulation using Python, a python-policy must be configured in the diameter-peer-policy. For example:

```
configure
aaa
    diameter-peer-policy "diameter-peer-policy-1" create
        description "Diameter peer policy"
        applications gy
        origin-host "bng.alcatel-lucent.com"
        origin-realm "alcatel-lucent.com"
        python-policy "py-policy-diam-1"
        source-address 192.0.2.5
        peer "peer-1" create
            address 172.16.1.1
            destination-host "server.alcatel-lucent.com"
            destination-realm "alcatel-lucent.com"
            no shutdown
        exit
    exit
exit
```

The python-policy specifies the python-script to use for each Diameter message type received or transmitted on a Diameter peer. In the ingress direction, the Python script is executed when the corresponding packet type is received on the Diameter peer but prior to further processing in the system. In the egress direction, the Python script is executed prior to sending the corresponding packet type on the Diameter peer. For example:

```
configure
python
    python-policy "py-policy-diam-1" create
        description "Python policy"
        diameter ccr direction egress script "diameter-1"
        diameter cca direction ingress script "diameter-2"
    exit
exit
```

The python-script specifies the location of the script and optional protection mechanism. For example:

```
configure
python
    python-script "diameter-1" create
        primary-url "ftp://usr:pwd@192.0.2.1/./py/diam-1.py"
        no shutdown
    exit
exit
```

As an example, the diam-1.py script, clears the M-bit from the Event-Timestamp AVP (code 55):

```
from alc import diameter
avp55=diameter.get_avps(55,0)
diameter.set_avps(55,0, [('x00', avp55[0][1])])
```

## Python Policy – DHCP Transaction Cache API

A DHCP transaction cache (DTC) is a short-lived cache during DHCPv4/v6 transaction. The cache could be used to store user-chosen information or return ESM attributes via a Python script. The DTC's lifetime is only during a single DHCP transaction (for example, only between discovery-offer, request-reply). This includes both `alc.dtc.store()` data and `alc.dtc.setESM()` data. DTC is also a transaction specific cache, which means the cached information could only be accessed by the Python script running in same DHCP transaction.

The following are the DTC APIs:

- `alc.dtc.derivedId`: A string used as a LUDB lookup key
- `alc.dtc.store(cache-key, cache-value)`: Store the value with the specified cache-key in DTC, both key and value are string.
- `alc.dtc.retrieve(cache-key)`: Returns the cached value string according to the specified cache-key, raise exception if specified key does not exist.
- `alc.dtc.setESM(ESM-key, value)`: Sets the specified ESM attribute, which could be used when system creating the ESM host. Note that due to the short-live nature of DTC, `setESM` should be used in final DHCP transaction before system create ESM host., such as DHCPv4 REQUEST-ACK.



- The following is a list of supported ESM-key and its corresponding python type:

- alc.dtc.accountingPolicy:str
- alc.dtc.ancpString:str
- alc.dtc.appProfileString:str
- alc.dtc.catMapString:str
- alc.dtc.defGw:str
- alc.dtc.dhcpv4GIAAddr:str
- alc.dtc.dhcpv4Pool:str
- alc.dtc.dhcpv4ServerAddr:str
- alc.dtc.dhcpv6LinkAddr:str
- alc.dtc.dhcpv6ServerAddr:str
- alc.dtc.intDestId:str
- alc.dtc.ipAddress:str
- alc.dtc.ipv4LeaseTime:int
- alc.dtc.ipv4PrimDns:str
- alc.dtc.ipv4SecDns:str
- alc.dtc.ipv6Address:str
- alc.dtc.ipv6DelegatedPrefix:str
- alc.dtc.ipv6DelegatedPrefixLength:int
- alc.dtc.ipv6PrefixPool:str
- alc.dtc.ipv6PrimDns:str
- alc.dtc.ipv6SecDns:str
- alc.dtc.ipv6SlaacPrefix:str
- alc.dtc.ipv6WanPool:str
- alc.dtc.msapGroupInterface:str
- alc.dtc.msapPolicy:str
- alc.dtc.msapServiceId:str
- alc.dtc.primNbns:str
- alc.dtc.retailServiceId:str
- alc.dtc.secNbns:str
- alc.dtc.slaProfileString:str
- alc.dtc.subIdent:str
- alc.dtc.subProfileString:str
- alc.dtc.subnetMask:str

## Python Cache Support

Python cache support allows information to be shared across different run times of the same python script or even different python scripts in a programmatic way. It essentially provides a central memory cache and a set of APIs which let the user store and retrieve strings. For example, a DHCP python-script could store a DHCP option into cache and later a RADIUS python-script could retrieve stored string and add it into access-request.

Each cached entry in the cache is a tuple of (**key**, **val**). **key** is used as entry id. **val** is the string to be cached. Both **key** and **val** are strings. The max length of the key is 512 bytes. Future more, the combine length of key+val is limited by the configured value of **entry-size** *size* command in the python-policy.

The Python cache can be enabled per python-policy. Each python policy has its own cache memory which script in other python-policy cannot access. This also implies that the key of a cached entry in different a python policy could overlap.

The user can also specify the max number cache entry per python policy the command **max-entries** command. System has a global limit for python cache memory of 256MB.

The cached entries could be made persistent by saving it to CF card. This can be enabled with the **persistence** command in the python policy.

**Note:** From memory consumption point of view, with MCS enabled, each cached entry will have a corresponding MCS record, so each entry will consume twice amount of memory.

The system also supports syncing the python cache across chassis with MCS. This can be configured per python policy with the **mcs-peer** command in the python policy.

Each cached entry has a remaining lifetime. If it decreases over time, the system will remove the cached entry if its remaining lifetime is 0. The remaining lifetime can be changed using a system-provided API. The initial lifetime of a newly created cache entry is 600 seconds.

The following are the python cache APIs in a module alc.cache:

- `alc.cache.save(val,key)`: Saves the val identified by the key into the cache. If there is an existing cache entry with same key, then it will be overwritten with the new val. An exception will be raised if the save failed (for example, due to exceeding the max number of entries).
- `alc.cache.retrieve(key)`: Returns the stored entry's val identified by the key. A `KeyError` exception will be raised if the specified entry does not exist.
- `alc.cache.clear(key)`: Removes the cached entry identified by the key. Raise `KeyError` exception if the specified entry does not exist.
- `cache.get_lifetime(key)`: The system will return a integer as seconds of remaining lifetime of the specified entry. It will return none if the specified entry does not exist. An exception will be raised for any other error.

- `cache.set_lifetime(key,new_lifetime)`: The `new_lifetime` value is an integer. The system will set the remaining lifetime of the specified entry to the number of seconds of the `new_lifetime`. An exception will be raised for any error including specified entry does not exist. If the `new_lifetime`  $\geq$  `max_lifetime` (configurable using the **max-entry-life** command in the python policy), then the system will set the actual lifetime to the `max_lifetime`.

## Applying a Python Policy

The following is a list of places that a Python policy could be applied:

- Under capture SAP — Apply to the DHCPv4/v6 packets sent/received on the capture SAP.
  - Under group-interface — Apply to DHCPv4/v6 packets sent/received on the group-interface.
  - Under subscriber-interface — Apply to DHCPv4 packets on the retail subscriber interface.
  - In the radius-server-policy — Apply to the RADIUS packets sent/received to/from the RADIUS servers configured in the radius-server-policy.
  - In the radius-proxy-server — Apply to the RADIUS packets on the client side of proxy.
  - In the diameter-peer-policy — Apply to the Diameter packets send/received on the Diameter peers configured in the policy.
- 

## Python Script Protection

In order to protect the Python script from unintended changes, the SR-OS supports a new Python script file format:SRPY. Since 12.0R1, SRPY includes a key based hash(HMAC) of the original script content. When the system loads a script with SRPY format, a hash will be computed by using a configured key and script content. The result hash will be compared to the embedded hash. If it is the same, then this script is considered valid. Otherwise, the system will abort with a warning message.

Users can configure **protection hmac-sha256 key <key>** within a Python script. To mandate, all configured scripts must be in SRPY format.

The system provides a tool command (**tool perform python-script protect**) to convert a Python script into SRPY format.

## Tips and Tricks

- Use xrange() instead of range().
- Avoid too many string operations. The following scripts provide the same output:

```
# This script takes 2.5 seconds. s = ""
for c in 'A'*50000:
    s += str(ord(c)) + ', '
print '[' + s[:-2] + ']'

# This script takes 0.1 seconds. print map(ord, 'A'*50000)
```



# Standards and Protocol Support

Note that the information presented is subject to change without notice.  
Alcatel-Lucent assumes no responsibility for inaccuracies contained herein.

## Ethernet Standards

IEEE 1588 Precision Clock Synchronization Protocol  
IEEE 802.1AB Station and Media Access Control Connectivity Discovery  
IEEE 802.1ad Provider Bridges  
IEEE 802.1ag Connectivity Fault Management  
IEEE 802.1ah Provider Backbone Bridges  
IEEE 802.1ak Multiple Registration Protocol  
IEEE 802.1aq Shortest Path Bridging  
IEEE 802.1ax Link Aggregation  
IEEE 802.1D MAC Bridges  
IEEE 802.1p Traffic Class Expediting  
IEEE 802.1Q Virtual LANs  
IEEE 802.1s Multiple Spanning Trees  
IEEE 802.1w Rapid Reconfiguration of Spanning Tree  
IEEE 802.1X Port Based Network Access Control  
IEEE 802.3ab 1000BASE-T  
IEEE 802.3ac VLAN Tag  
IEEE 802.3ad Link Aggregation  
IEEE 802.3ae 10 Gb/s Ethernet  
IEEE 802.3ah Ethernet in the First Mile  
IEEE 802.3ba 40 Gb/s and 100 Gb/s Ethernet  
IEEE 802.3i Ethernet  
IEEE 802.3u Fast Ethernet  
IEEE 802.3x Ethernet Flow Control  
IEEE 802.3z Gigabit Ethernet  
ITU-T G.8031 Ethernet Linear Protection Switching  
ITU-T G.8032 Ethernet Ring Protection Switching  
ITU-T Y.1731 OAM functions and mechanisms for Ethernet based networks

## OSPF

RFC 1586 Guidelines for Running OSPF Over Frame Relay Networks  
RFC 1765 OSPF Database Overflow  
RFC 2328 OSPF Version 2  
RFC 3101 The OSPF Not-So-Stubby Area (NSSA) Option  
RFC 3509 Alternative Implementations of OSPF Area Border Routers  
RFC 3623 Graceful OSPF Restart (Helper Mode)  
RFC 3630 Traffic Engineering (TE) Extensions to OSPF Version 2  
RFC 4203 OSPF Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)  
RFC 4222 Prioritized Treatment of Specific OSPF Version 2 Packets and Congestion Avoidance  
RFC 4576 Using a Link State Advertisement (LSA) Options Bit to Prevent Looping in BGP/MPLS IP Virtual Private Networks (VPNs)  
RFC 4970 Extensions to OSPF for Advertising Optional Router Capabilities  
RFC 5185 OSPF Multi-Area Adjacency  
RFC 5243 OSPF Database Exchange Summary List Optimization  
RFC 5250 The OSPF Opaque LSA Option  
RFC 5709 OSPFv2 HMAC-SHA Cryptographic Authentication  
RFC 6987 OSPF Stub Router Advertisement

## BGP

RFC 1397 BGP Default Route Advertisement  
RFC 1772 Application of BGP in the Internet  
RFC 1965 Confederations for BGP  
RFC 1997 BGP Communities Attribute  
RFC 2385 Protection of BGP Sessions via MD5  
RFC 2439 BGP Route Flap Dampening

RFC 2858 Multiprotocol Extensions for BGP-4  
RFC 2918 Route Refresh Capability for BGP-4  
RFC 3107 Carrying Label Information in BGP-4  
RFC 3392 Capabilities Advertisement with BGP4  
RFC 4271 BGP-4 (previously RFC 1771)  
RFC 4360 BGP Extended Communities Attribute  
RFC 4364 BGP/MPLS IP Virtual Private Networks (VPNs)(previously RFC 2547bis BGP/MPLS VPNs)  
RFC 4456 BGP Route Reflection: Alternative to Full-mesh IBGP  
RFC 4486 Subcodes for BGP Cease Notification Message  
RFC 4577 OSPF as the Provider/Customer Edge Protocol for BGP/MPLS IP Virtual Private Networks (VPNs)  
RFC 4659 BGP-MPLS IP Virtual Private Network (VPN) Extension for IPv6 VPN  
RFC 4684 Constrained Route Distribution for Border Gateway Protocol/MultiProtocol Label Switching (BGP/MPLS) Internet Protocol (IP) Virtual Private Networks (VPNs)  
RFC 4724 Graceful Restart Mechanism for BGP – GR helper  
RFC 4760 Multi-protocol Extensions for BGP  
RFC 4798 Connecting IPv6 Islands over IPv4 MPLS Using IPv6 Provider Edge Routers (6PE)  
RFC 4893 BGP Support for Four-octet AS Number Space  
RFC 5004 Avoid BGP Best Path Transitions from One External to Another  
RFC 5065 Confederations for BGP (obsoletes 3065)  
RFC 5291 Outbound Route Filtering Capability for BGP-4

RFC 5575 Dissemination of Flow Specification Rules  
RFC 5668 4-Octet AS Specific BGP Extended Community  
draft-ietf-idr-add-paths Advertisement of Multiple Paths in BGP  
draft-ietf-idr-best-external Advertisement of the Best External Route in BGP

### IS-IS

ISO/IEC 10589:2002, Second Edition, Nov. 2002 Intermediate System to Intermediate System Intra-Domain Routeing Information Exchange Protocol  
RFC 1195 Use of OSI IS-IS for Routing in TCP/IP and Dual Environments  
RFC 2973 IS-IS Mesh Groups  
RFC 3359 Reserved Type, Length and Value (TLV) Codepoints in Intermediate System to Intermediate System  
RFC 3719 Recommendations for Interoperable Networks using Intermediate System to Intermediate System (IS-IS)  
RFC 3787 Recommendations for Interoperable IP Networks using Intermediate System to Intermediate System (IS-IS)  
RFC 4971 Intermediate System to Intermediate System (IS-IS) Extensions for Advertising Router Information  
RFC 5120 M-ISIS: Multi Topology (MT) Routing in IS-IS  
RFC 5130 A Policy Control Mechanism in IS-IS Using Administrative Tags  
RFC 5301 Dynamic Hostname Exchange Mechanism for IS-IS  
RFC 5302 Domain-wide Prefix Distribution with Two-Level IS-IS  
RFC 5303 Three-Way Handshake for IS-IS Point-to-Point Adjacencies  
RFC 5304 IS-IS Cryptographic Authentication  
RFC 5305 IS-IS Extensions for Traffic Engineering TE  
RFC 5306 Restart Signaling for IS-IS (Helper Mode)  
RFC 5307 IS-IS Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)

RFC 5309 Point-to-Point Operation over LAN in Link State Routing Protocols  
RFC 5310 IS-IS Generic Cryptographic Authentication  
RFC 6213 IS-IS BFD-Enabled TLV  
RFC 6329 IS-IS Extensions Supporting IEEE 802.1aq Shortest Path Bridging  
draft-ietf-isis-mi-02 IS-IS Multi-Instance

### IP, LDP, and Segment Routing Fast Reroute (FRR)

RFC 5286 Basic Specification for IP Fast Reroute: Loop-Free Alternates  
draft-ietf-isis-segment-routing-extensions-03 IS-IS Extensions for Segment Routing  
draft-ietf-rtgwg-lfa-manageability-07 Operational management of Loop Free Alternates  
draft-ietf-rtgwg-remote-lfa-09 Remote LFA FRR  
draft-kratn-mofrr-02 Multicast only Fast Re-Route

### IPSec

RFC 2401 Security Architecture for the Internet Protocol  
RFC 2406 IP Encapsulating Security Payload (ESP)  
RFC 2409 The Internet Key Exchange (IKE)  
RFC 2560 X.509 Internet Public Key Infrastructure Online Certificate Status Protocol - OCSP  
RFC 3706 IKE Dead Peer Detection  
RFC 3947 Negotiation of NAT-Traversal in the IKE  
RFC 3948 UDP Encapsulation of IPsec ESP Packets  
RFC 4210 Internet X.509 Public Key Infrastructure Certificate Management Protocol (CMP)  
RFC 4211 Internet X.509 Public Key Infrastructure Certificate Request Message Format (CRMF)  
RFC 5996 Internet Key Exchange Protocol Version 2 (IKEv2)  
RFC 5998 An Extension for EAP-Only Authentication in IKEv2

draft-ietf-ipsec-isakmp-xauth-06 Extended Authentication within ISAKMP/Oakley (XAUTH)  
draft-ietf-ipsec-isakmp-modecfg-05 The ISAKMP Configuration Method

### IPv6

RFC 1981 Path MTU Discovery for IPv6  
RFC 2375 IPv6 Multicast Address Assignments  
RFC 2460 Internet Protocol, Version 6 (IPv6) Specification  
RFC 2461 Neighbor Discovery for IPv6  
RFC 2462 IPv6 Stateless Address Auto configuration  
RFC 2464 Transmission of IPv6 Packets over Ethernet Networks  
RFC 2529 Transmission of IPv6 over IPv4 Domains without Explicit Tunnels  
RFC 2545 Use of BGP-4 Multiprotocol Extension for IPv6 Inter-Domain Routing  
RFC 2710 Multicast Listener Discovery (MLD) for IPv6  
RFC 2740 OSPF for IPv6  
RFC 3306 Unicast-Prefix-based IPv6 Multicast Addresses  
RFC 3315 Dynamic Host Configuration Protocol for IPv6  
RFC 3587 IPv6 Global Unicast Address Format  
RFC 3590 Source Address Selection for the Multicast Listener Discovery (MLD) Protocol  
RFC 3810 Multicast Listener Discovery Version 2 (MLDv2) for IPv6  
RFC 3971 SEcure Neighbor Discovery (SEND)  
RFC 3972 Cryptographically Generated Addresses (CGA)  
RFC 4007 IPv6 Scoped Address Architecture  
RFC 4193 Unique Local IPv6 Unicast Addresses  
RFC 4291 IPv6 Addressing Architecture  
RFC 4443 Internet Control Message Protocol (ICMPv6)  
for the Internet Protocol Version 6 (IPv6) Specification  
RFC 4552 Authentication/Confidentiality for OSPFv3



RFC 4659 BGP-MPLS IP Virtual Private Network (VPN) Extension for IPv6 VPN  
 RFC 5072 IP Version 6 over PPP  
 RFC 5095 Deprecation of Type 0 Routing Headers in IPv6  
 RFC 5187 OSPFv3 Graceful Restart (Helper Mode)  
 RFC 5308 Routing IPv6 with IS-IS  
 RFC 5340 OSPF for IPv6  
 RFC 5838 Support of Address Families in OSPFv3

### **Multicast**

RFC 1112 Host Extensions for IP Multicasting (Snooping)  
 RFC 2236 Internet Group Management Protocol, (Snooping)  
 RFC 2362 Protocol Independent Multicast-Sparse Mode (PIMSM)  
 RFC 3376 Internet Group Management Protocol, Version 3 (Snooping)  
 RFC 3446 Anycast Rendezvous Point (RP) mechanism using Protocol Independent Multicast (PIM) and Multicast Source Discovery Protocol (MSDP)  
 RFC 3618 Multicast Source Discovery Protocol (MSDP)  
 RFC 3956 Embedding the Rendezvous Point (RP) Address in an IPv6 Multicast Address  
 RFC 4601 Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)  
 RFC 4604 Using IGMPv3 and MLDv2 for Source-Specific Multicast  
 RFC 4607 Source-Specific Multicast for IP  
 RFC 4608 Source-Specific Protocol Independent Multicast in 232/8  
 RFC 4610 Anycast-RP Using Protocol Independent Multicast (PIM)  
 RFC 4624 Multicast Source Discovery Protocol (MSDP) MIB  
 RFC 5059 Bootstrap Router (BSR) Mechanism for Protocol Independent Multicast (PIM)  
 RFC 5059 Bootstrap Router (BSR) Mechanism for Protocol Independent Multicast (PIM)

RFC 5384 The Protocol Independent Multicast (PIM) Join Attribute Format  
 RFC 5496 The Reverse Path Forwarding (RPF) Vector TLV  
 RFC 6037 Cisco Systems' Solution for Multicast in MPLS/BGP IP VPNs  
 RFC 6513 Multicast in MPLS/BGP IP VPNs  
 RFC 6514 BGP Encodings and Procedures for Multicast in MPLS/ IP VPNs  
 RFC 6515 IPv4 and IPv6 Infrastructure Addresses in BGP Updates for Multicast VPNs  
 RFC 6516 IPv6 Multicast MVPN Support Using PIM Control Plane and Selective Provider Multicast Service Interface (S-PMSI) Join Messages  
 RFC 6625 Wildcards in Multicast VPN Auto-Discover Routes  
 RFC 6826 Multipoint LDP In-Band Signaling for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Path  
 RFC 7246 Multipoint Label Distribution Protocol In-Band Signaling in a Virtual Routing and Forwarding (VRF)  
 RFC 7385 IANA Registry for P-Multicast Service Interface (PMSI) Tunnel Type Code Points  
 draft-dolganow-l3vpn-mvpn-expl-track-00 Explicit tracking in MPLS/BGP IP VPN

### **MPLS — GENERAL**

RFC 2430 A Provider Architecture DiffServ & TE  
 RFC 2474 Definition of the DS Field the IPv4 and IPv6 Headers (Rev)  
 RFC 2597 Assured Forwarding PHB Group (rev3260)  
 RFC 2598 An Expedited Forwarding PHB  
 RFC 3031 MPLS Architecture  
 RFC 3032 MPLS Label Stack Encoding  
 RFC 3140 Per-Hop Behavior Identification Codes  
 RFC 3443 Time To Live (TTL) Processing in Multi-Protocol Label Switching (MPLS) Networks

RFC 4023 Encapsulating MPLS in IP or Generic Routing Encapsulation (GRE)  
 RFC 4182 Removing a Restriction on the use of MPLS Explicit NULL  
 RFC 5332 MPLS Multicast Encapsulations

### **MPLS — LDP**

RFC 3037 LDP Applicability  
 RFC 3478 Graceful Restart Mechanism for LDP – GR helper  
 RFC 5036 LDP Specification  
 RFC 5283 LDP extension for Inter-Area LSP  
 RFC 5443 LDP IGP Synchronization  
 RFC 5561 LDP Capabilities  
 RFC 6388 LDP Extensions for Point-to-Multipoint and Multipoint-to-Multipoint LSP  
 RFC 6826 Multipoint LDP in-band signaling for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths  
 draft-ietf-mpls-ldp-ip-pw-capability-09 Disabling IPoMPLS and P2P PW LDP Application's State Advertisement  
 draft-ietf-mpls-ldp-ipv6-15 Updates to LDP for IPv6  
 draft-pdutta-mpls-ldp-adj-capability-00 LDP Adjacency Capabilities  
 draft-pdutta-mpls-ldp-v2-00 LDP Version 2  
 draft-pdutta-mpls-multi-ldp-instance-00 Multiple LDP Instances  
 draft-pdutta-mpls-tldp-hello-reduce-04 Targeted LDP Hello Reduction

### **MPLS/RSVP — TE**

RFC 2702 Requirements for Traffic Engineering over MPLS  
 RFC2747 RSVP Cryptographic Authentication  
 RFC 2961 RSVP Refresh Overhead Reduction Extensions  
 RFC3097 RSVP Cryptographic Authentication - Updated Message Type Value  
 RFC 3209 Extensions to RSVP for Tunnels

RFC 3473 Generalized Multi-Protocol Label Switching (GMPLS) Signaling

Resource Reservation Protocol-Traffic Engineering (RSVP-TE) Extensions – (support of of IF\_ID RSVP\_HOP object with unnumbered interface and RSVP-TE Graceful Restart Helper Procedures)

RFC 3477 Signalling Unnumbered Links in Resource Reservation Protocol-Traffic Engineering (RSVP-TE)

RFC 3564 Requirements for Diff-Serv-aware TE

RFC 3906 Calculating Interior Gateway Protocol (IGP) Routes Over Traffic Engineering Tunnels

RFC 4090 Fast reroute Extensions to RSVP-TE for LSP Tunnels

RFC 4124 Protocol Extensions for Support of Diffserv-aware MPLS Traffic Engineering

RFC 4125 Maximum Allocation Bandwidth Constraints Model for Diffserv-aware MPLS Traffic Engineering

RFC 4127 Russian Dolls Bandwidth Constraints Model for Diffserv-aware MPLS Traffic Engineering

RFC 4561 Definition of a RRO Node-Id Sub-Object

RFC 4875 Extensions to Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label Switched Paths (LSPs)

RFC 4950 ICMP Extensions for Multiprotocol Label Switching

RFC 5151 Inter-domain MPLS and GMPLS Traffic Engineering – RSVP-TE Extensions

RFC 5712 MPLS Traffic Engineering Soft Preemption

RFC 5817 Graceful Shutdown in GMPLS Traffic Engineering Networks

draft-newton-mpls-te-dynamic-overbooking-00 A Diffserv-TE Implementation Model to dynamically change booking factors during failure events

### **MPLS — OAM**

RFC 4379 Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures

RFC 6424 Mechanism for Performing Label Switched Path Ping (LSP Ping) over MPLS Tunnels

RFC 6425 Detecting Data Plane Failures in Point-to-Multipoint Multiprotocol Label Switching (MPLS) - Extensions to LSP Ping

### **MPLS — TP (7750/7450 only)**

RFC 5586 MPLS Generic Associated Channel

RFC 5921 A Framework for MPLS in Transport Networks

RFC 5960 MPLS Transport Profile Data Plane Architecture

RFC 6370 MPLS-TP Identifiers

RFC 6378 MPLS-TP Linear Protection

RFC 6428 Proactive Connectivity Verification, Continuity Check and Remote Defect indication for MPLS Transport Profile

RFC 6426 MPLS On-Demand Connectivity and Route Tracing

RFC 6478 Pseudowire Status for Static Pseudowires

RFC 7213 MPLS Transport Profile (MPLS-TP) Next-Hop Ethernet Addressing

### **MPLS — GMPLS**

RFC 3471 Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description

RFC 3473 Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource Reservation Protocol-Traffic Engineering (RSVP-TE) Extensions

RFC 4204 Link Management Protocol (LMP)

RFC 4208 Generalized Multiprotocol Label Switching (GMPLS) User-Network Interface (UNI): Resource Reservation Protocol-Traffic Engineering (RSVP-TE) Support for the Overlay Model

RFC 4872 RSVP-TE Extensions in Support of End to End GMPLS recovery

draft-ietf-ccamp-rsvp-te-srlg-collect-04 RSVP-TE Extensions for Collecting SRLG Information

### **RIP**

RFC 1058 RIP Version 1

RFC 2080 RIPng for IPv6

RFC 2082 RIP-2 MD5 Authentication

RFC 2453 RIP Version 2

### **TCP/IP**

RFC 768 UDP

RFC 791 IP

RFC 792 ICMP

RFC 793 TCP

RFC 826 ARP

RFC 854 Telnet

RFC 951 Bootstrap Protocol (BOOTP)

RFC 1350 The Tftp Protocol (revision 2)

RFC 1519 CIDR

RFC 1542 Clarifications and Extensions for the Bootstrap Protocol

RFC 1812 Requirements for IPv4 Routers

RFC 2347 TFTP option Extension

RFC 2328 TFTP Blocksize Option

RFC 2349 TFTP Timeout Interval and Transfer Size option

RFC 2401 Security Architecture for Internet Protocol

RFC 2428 FTP Extensions for IPv6 and NATs

RFC 3596 DNS Extensions to Support IP version 6

RFC 5880 Bidirectional Forwarding Detection

RFC 5881 BFD IPv4 and IPv6 (Single Hop)

RFC 5883 BFD for Multihop Paths

### **VRRP**

RFC 2787 Definitions of Managed Objects for the Virtual Router Redundancy Protocol

RFC 3768 Virtual Router Redundancy Protocol

draft-ietf-vrrp-unified-spec-02 Virtual Router Redundancy Protocol Version 3 for IPv4 and IPv6

**PPP**

RFC 1332 PPP IPCP  
 RFC 1377 PPP OSINLCP  
 RFC 1638/2878 PPP BCP  
 RFC 1661 PPP (rev RFC2151)  
 RFC 1662 PPP in HDLC-like Framing  
 RFC 1877 PPP Internet Protocol Control Protocol Extensions for Name Server Addresses  
 RFC 1989 PPP Link Quality Monitoring  
 RFC 1990 The PPP Multilink Protocol (MP)  
 RFC 1994 PPP Challenge Handshake Authentication Protocol (CHAP)  
 RFC 2516 A Method for Transmitting PPP Over Ethernet  
 RFC 2615 PPP over SONET/SDH  
 RFC 2686 The Multi-Class Extension to Multi-Link PPP

**Frame Relay**

FRF.1.2 - PVC User-to-Network Interface (UNI) Implementation Agreement  
 FRF.5 - Frame Relay/ATM PVC Network Interworking Implementation  
 ANSI T1.617 Annex D, DSS1 — Signalling Specification For Frame Relay Bearer Service.  
 FRF2.2 PVC Network-to- Network Interface (NNI) Implementation Agreement.  
 FRF.12 Frame Relay Fragmentation Implementation Agreement  
 FRF.16.1 Multilink Frame Relay UNI/ NNI Implementation Agreement  
 ITU-T Q.933, Annex A Additional procedures for Permanent Virtual Connection (PVC) status management

**ATM**

RFC 1626 Default IP MTU for use over ATM AAL5  
 RFC 2514 Definitions of Textual Conventions and OBJECT\_IDENTITIES for ATM Management  
 RFC 2515 Definition of Managed Objects for ATM Management  
 RFC 2684 Multiprotocol Encapsulation over ATM Adaptation Layer 5

AF-TM-0121.000 Traffic Management Specification Version 4.1  
 ITU-T Recommendation I.610 B-ISDN Operation and Maintenance Principles and Functions version 11/95  
 ITU-T Recommendation I.432.1 BISDN user-network interface – Physical layer specification: General characteristics  
 GR-1248-CORE Generic Requirements for Operations of ATM Network Elements (NEs). Issue 3  
 GR-1113-CORE Bellcore, Asynchronous Transfer Mode (ATM) and ATM Adaptation Layer (AAL) Protocols Generic Requirements, Issue 1  
 AF-ILMI-0065.000 Integrated Local Management Interface (ILMI) Version 4.0  
 AF-TM-0150.00 Addendum to Traffic Management v4.1 optional minimum desired cell rate indication for UBR  
 AF-PHY-0086.001 Inverse Multiplexing for ATM (IMA) Specification Version 1.1

**DHCP**

RFC 2131 Dynamic Host Configuration Protocol (REV)  
 RFC 3046 DHCP Relay Agent Information Option (Option 82)  
 RFC 1534 Interoperation between DHCP and BOOTP

**Policy Management and Credit Control**

3GPP TS 29.212 Policy and Charging Control (PCC) over Gx/Sd Reference Point (Release 11 and Release 12) - Gx support as it applies to wireline environment (BNG)  
 RFC 3588 Diameter Base Protocol  
 RFC 4006 Diameter Credit Control Application

**NAT**

RFC 5382 NAT Behavioral Requirements for TCP  
 RFC 5508 NAT Behavioral Requirements for ICMP

RFC 6146 Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers  
 RFC 6333 Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion  
 RFC 6334 Dynamic Host Configuration Protocol for IPv6 (DHCPv6) Option for Dual-Stack Lite  
 RFC 6888 Common Requirements For Carrier-Grade NATs (CGNs)

**VPLS**

RFC 4761 Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling  
 RFC 4762 Virtual Private LAN Services Using LDP  
 RFC 5501 Requirements for Multicast Support in Virtual Private LAN Services  
 RFC 6074 Provisioning, Auto-Discovery, and Signaling in Layer 2 Virtual Private Networks (L2VPNs)  
 RFC 7041 Extensions to the Virtual Private LAN Service (VPLS) Provider Edge (PE) Model for Provider Backbone Bridging  
 RFC 7117 Multicast in Virtual Private LAN Service (VPLS)

**Pseudowire**

RFC 3985 Pseudo Wire Emulation Edge-to-Edge (PWE3)  
 RFC 4385 Pseudo Wire Emulation Edge-to-Edge (PWE3) Control Word for Use over an MPLS PSN  
 RFC 3916 Requirements for Pseudo-Wire Emulation Edge-to-Edge (PWE3)  
 RFC 4717 Encapsulation Methods for Transport ATM over MPLS Networks  
 RFC 4816 PWE3 ATM Transparent Cell Transport Service  
 RFC 4448 Encapsulation Methods for Transport of Ethernet over MPLS Networks  
 RFC 4619 Encapsulation Methods for Transport of Frame Relay over MPLS Networks  
 RFC 4446 IANA Allocations for PWE3  
 RFC 4447 Pseudowire Setup and Maintenance Using LDP

## Standards and Protocols

RFC 5085 Pseudowire Virtual Circuit Connectivity Verification (VCCV): A Control Channel for Pseudowires  
RFC 5659 An Architecture for Multi-Segment Pseudowire Emulation Edge-to-Edge  
RFC 5885 Bidirectional Forwarding Detection (BFD) for the Pseudowire Virtual Circuit Connectivity Verification (VCCV)  
RFC 6073 Segmented Pseudowire  
RFC 6310 Pseudowire (PW) OAM Message Mapping  
RFC 6391 Flow Aware Transport of Pseudowires over an MPLS PSN  
RFC 6575 ARP Mediation for IP Interworking of Layer 2 VPN  
RFC 6718 Pseudowire Redundancy  
RFC 6829 Label Switched Path (LSP) Ping for Pseudowire Forwarding Equivalence Classes (FECs) Advertised over IPv6  
RFC 6870 Pseudowire Preferential Forwarding Status bit  
RFC 7023 MPLS and Ethernet OAM Interworking  
RFC 7267 Dynamic Placement of Multi-Segment Pseudowires  
draft-ietf-l2vpn-vpws-iw-oam-04 OAM Procedures for VPWS Interworking  
MFA Forum 9.0.0 The Use of Virtual trunks for ATM/MPLS Control Plane Interworking  
MFA Forum 12.0.0 Multiservice Interworking - Ethernet over MPLS  
MFA Forum 13.0.0 Fault Management for Multiservice Interworking v1.0  
MFA Forum 16.0.0 Multiservice Interworking - IP over MPLS

### ANCP/L2CP

RFC 5851 ANCP framework  
draft-ietf-ancp-protocol-02 ANCP Protocol

### Voice /Video Performance:

ITU-T G.107 The E Model- A computational model for use in planning.  
ETSI TS 101 329-5 Annex E extensions- QoS Measurement for VoIP - Method for determining an

Equipment Impairment Factor using Passive Monitoring  
ITU-T Rec. P.564 Conformance testing for voice over IP transmission quality assessment models  
ITU-T G.1020, Appendix I Performance Parameter Definitions for Quality of Speech and other Voiceband Applications Utilizing IP Networks- Mean Absolute Packet Delay Variation & Markov Models.  
RFC 3550, Appendix A.8 RTP: A Transport Protocol for Real-Time Applications- Estimating the Interarrival Jitter.

### Circuit Emulation

RFC 4553 Structure-Agnostic Time Division Multiplexing (TDM) over Packet (SAToP)  
RFC 5086 Structure-Aware Time Division Multiplexed (TDM) Circuit Emulation Service over Packet Switched Network (CESoPSN)  
MEF-8 Implementation Agreement for the Emulation of PDH Circuits over Metro Ethernet Networks, October 2004  
RFC 5287 Control Protocol Extensions for the Setup of Time-Division Multiplexing (TDM) Pseudowires in MPLS Networks

### SONET/SDH

ITU-G.841 Telecommunication Standardization Section of ITU, Types and Characteristics of SDH Networks Protection Architecture, issued in October 1998 and as augmented by Corrigendum1 issued in July 2002

### AAA

RFC 2865 Remote Authentication Dial In User Service  
RFC 2866 RADIUS Accounting  
draft-grant-tacacs-02 The TACACS+ Protocol

### SSH

RFC 4250 The Secure Shell (SSH) Protocol Assigned Numbers  
RFC 4251 The Secure Shell (SSH) Protocol Architecture

RFC 4254 The Secure Shell (SSH) Connection Protocol

### OpenFlow

ONF OpenFlow Switch Specification Version 1.3.1 (Hybrid-switch/FlowTable)

### Timing

GR-253-CORE SONET Transport Systems: Common Generic Criteria. Issue 3, September 2000  
ITU-T G.781 Telecommunication Standardization Section of ITU, Synchronization layer functions, issued 09/2008  
ITU-T G.813 Telecommunication Standardization Section of ITU, Timing characteristics of SDH equipment slave clocks (SEC), issued 03/2003.  
GR-1244-CORE Clocks for the Synchronized Network: Common Generic Criteria, Issue 3, May 2005  
ITU-T G.8261 Telecommunication Standardization Section of ITU, Timing and synchronization aspects in packet networks, issued 04/2008.  
ITU-T G.8262 Telecommunication Standardization Section of ITU, Timing characteristics of synchronous Ethernet equipment slave clock (EEC), issued 08/2007.  
ITU-T G.8264 Telecommunication Standardization Section of ITU, Distribution of timing information through packet networks, issued 10/2008.  
ITU-T G.8265.1 Telecommunication Standardization Section of ITU, Precision time protocol telecom profile for frequency synchronization, issued 10/2010.  
IEEE 1588-2008 IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems

**Network Management**

ITU-T X.721 Information technology-  
OSI-Structure of Management  
Information

ITU-T X.734 Information technology-  
OSI-Systems Management: Event  
Report Management Function

M.3100/3120 Equipment and Connection  
Models

TMF 509/613 Network Connectivity  
Model

RFC 1157 SNMPv1

RFC 1215 A Convention for Defining  
Traps for use with the SNMP

RFC 1657 BGP4-MIB

RFC 1724 RIPv2-MIB

RFC 1850 OSPF-MIB

RFC 1907 SNMPv2-MIB

RFC 2011 IP-MIB

RFC 2138 RADIUS

RFC 2206 RSVP-MIB

RFC 2452 IPv6 Management Information  
Base for the Transmission Control  
Protocol

RFC 2465 Management Information  
Base for IPv6: Textual Conventions  
and General Group

RFC 2558 SONET-MIB

RFC 2571 SNMP-FRAMEWORKMIB

RFC 2572 SNMP-MPD-MIB

RFC 2573 SNMP-TARGET-&-  
NOTIFICATION-MIB

RFC 2574 SNMP-USER-BASED-  
SMMIB

RFC 2575 SNMP-VIEW-BASED-ACM-  
MIB

RFC 2576 SNMP-COMMUNITY-MIB

RFC 2578 Structure of Management  
Information Version 2 (SMIv2)

RFC 2665 EtherLike-MIB

RFC 2819 RMON-MIB

RFC 2863 IF-MIB

RFC 2864 INVERTED-STACK-MIB

RFC 2987 VRRP-MIB

RFC 3014 NOTIFICATION-LOGMIB

RFC 3019 IP Version 6 Management  
Information Base for The Multicast  
Listener Discovery Protocol

RFC 3164 Syslog

RFC 3273 HCRMON-MIB

RFC 3411 An Architecture for  
Describing Simple Network

Management Protocol (SNMP)  
Management Frameworks

RFC 3412 Message Processing and  
Dispatching for the Simple Network  
Management Protocol (SNMP)

RFC 3413 Simple Network Management  
Protocol (SNMP) Applications

RFC 3414 User-based Security Model  
(USM) for version 3 of the Simple  
Network Management Protocol  
(SNMPv3)

RFC 3418 SNMP MIB

RFC 3826 The Advanced Encryption  
Standard (AES) Cipher Algorithm in  
the SNMP User-based Security  
Model

RFC 4113 Management Information  
Base for the User Datagram Protocol  
(UDP)

RFC 4292 IP Forwarding Table MIB

RFC 4293 MIB for the Internet Protocol

RFC 5101 Specification of the IP Flow  
Information Export (IPFIX)  
Protocol for the Exchange of IP  
Traffic Flow Information

RFC 6241 Network Configuration  
Protocol (NETCONF)

RFC 6242 Using the NETCONF Protocol  
over Secure Shell (SSH)

draft-ietf-bfd-mib-00 Bidirectional  
Forwarding Detection Management  
Information Base

draft-ietf-isis-wg-mib-06 Management  
Information Base for Intermediate  
System to Intermediate System (IS-  
IS)

draft-ietf-ospf-mib-update-04 OSPF  
Version 2 Management Information  
Base

draft-ietf-mboned-msdp-mib-01  
Multicast Source Discovery protocol  
MIB

draft-ietf-mppls-lsr-mib-06 Multiprotocol  
Label Switching (MPLS) Label  
Switching Router (LSR)  
Management Information Base

draft-ietf-mppls-te-mib-04 Multiprotocol  
Label Switching (MPLS) Traffic  
Engineering Management  
Information Base

draft-ietf-mppls-ldp-mib-07 Definitions of  
Managed Objects for the  
Multiprotocol Label Switching,  
Label Distribution Protocol (LDP)

IEEE 802.3ad MIB



# Customer documentation and product support



## Customer documentation

<http://documentation.alcatel-lucent.com>



## Technical support

<http://support.alcatel-lucent.com>



## Documentation feedback

[documentation.feedback@alcatel-lucent.com](mailto:documentation.feedback@alcatel-lucent.com)

