



7750 SR OS Triple Play Guide

Software Version: 7750 SR OS 8.0 Rev. 01

February 2010

Document Part Number: 93-0098-06-01



This document is protected by copyright. Except as specifically permitted herein, no portion of the provided information can be reproduced in any form, or by any means, without prior written permission from Alcatel-Lucent.

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

The information presented is subject to change without notice.

Alcatel-Lucent assumes no responsibility for inaccuracies contained herein.

Copyright 2010 Alcatel-Lucent Alcatel-Lucent. All rights reserved.

Table of Contents

Preface	13
Alcatel-Lucent 7750 SR-Series Services Configuration Process	15
Introduction to Triple Play	
Alcatel-Lucent's Triple Play Service Delivery Architecture	18
Introduction to Triple Play	18
Blueprint for Optimizing Triple Play Service Infrastructures	19
Architectural Foundations	20
Optimizing Triple Play Service Infrastructures	22
Distributed Service Edges	23
Service Differentiation, QoS Enablement	25
Virtual MAC Subnetting for VPLS	29
Services	32
Service Types	33
Service Policies	34
Alcatel-Lucent Service Model	35
Introduction	35
Service Entities	36
Customers	36
Service Access Points (SAPs)	37
SAP Encapsulation Types and Identifiers	38
Ethernet Encapsulations	38
SAP Considerations	39
Service Distribution Points (SDPs)	41
SDP Binding	41
Spoke and Mesh SDPs	42
SDP Encapsulation Types	43
SDP Keepalives	44
Epipe Service Overview	45
VPLS Service Overview	46
Split Horizon SAP Groups and Split Horizon Spoke SDP Groups	46
Residential Split Horizon Groups	47
IES Service Overview	48
IP Interface	49
VPRN Service Overview	50
Service Creation Process Overview	51
Service Components	52
Deploying and Provisioning Services	54
Phase 1: Core Network Construction	54
Phase 2: Service Administration	54
Phase 3: Service Provisioning	54
Configuration Notes	55
General	55
Configuring Triple Play Services with CLI	57
Configuring VPLS Residential Split Horizon Groups	57

Table of Contents

Configuring Static Hosts	58
Triple Play Services Command Reference	61
DHCP Management	
DHCP Principles	302
DHCP Features	304
DHCP Relay	304
Subscriber Identification Using Option 82 Field	304
DHCP Snooping	307
DHCP Lease State Table	309
DHCP and Layer 3 Aggregation	311
Local DHCP Servers	313
DHCPv6	315
DHCPv6 Relay Agent	315
DHCPv6 Prefix Options	315
Neighbor Resolution via DHCPv6 Relay	315
DHCPv6 Lease Persistency	316
Local Proxy Neighbor Discovery	316
Proxy DHCP Server	317
Configuring DHCP with CLI	321
Common Configuration Tasks	322
Enabling DHCP Snooping	322
Configuring Option 82 Handling	324
Enabling DHCP Relay	325
Configuring Local User Database Parameters	326
Triple Play DHCP Command Reference	333
Point-to-Point Protocol over Ethernet (PPPoE) Management	
PPPoE	424
PPPoE Authentication and Authorization	427
General Flow	427
RADIUS	428
Local User Database Directly Assigned to PPPoE Node	429
Local DHCP Server with Local User Database	430
Multiple Sessions Per MAC Address	432
Private Retail Subnets	433
Private Retail Subnets	434
PPPoE Command Reference	437
Triple Play Security	
Triple Play Security Features	466
Anti-Spoofing Filters	466
Anti-spoofing Filter Types	467
Filtering packets	467
Layer 2 Triple Play Security Features	469
MAC Pinning	469
MAC Protection	469
DoS Protection	470
VPLS Redirect Policy	472

ARP Handling	473
ARP Reply Agent	473
Dynamic ARP table population	474
Local Proxy ARP	474
Web Portal Redirect	475
Configuring Triple Play Security with CLI	479
Common Configuration Tasks	480
Configuring Anti-Spoofing Filters	480
Configuring Triple Play Security features	481
Configuring ARP Handling	487
Configuring Web Portal Redirect	495
Triple Play Security Command Reference	497

Triple Play Multicast

Introduction to Multicast	524
Multicast in the BSR	525
Internet Group Management Protocol	525
IGMP Versions and Interoperability Requirements	525
IGMP Version Transition	526
Source Specific Multicast Groups	526
Protocol Independent Multicast Sparse Mode (PIM-SM)	527
Ingress Multicast Path Management (IMPM) Enhancements	528
Multicast in the BSA	529
IGMP Snooping	529
IGMP Message processing	530
IGMP Filtering	531
Multicast VPLS Registration (MVR)	532
Layer 3 Multicast Load Balancing	533
Configuring Triple Play Multicast Services with CLI	535
Configuring IGMP Snooping in the BSA	536
Enabling IGMP Snooping in a VPLS Service	536
With IGMPv3 Multicast Routers	536
With IGMPv1/2 Multicast Routers	537
Modifying IGMP Snooping Parameters	538
Modifying IGMP Snooping Parameters for a SAP or SDP	539
Configuring Static Multicast Groups on a SAP or SDP	541
Enabling IGMP Group Membership Report Filtering	542
Configuring Multicast VPLS Registration (MVR)	545
Configuring IGMP and PIM in the BSR	546
Enabling IGMP	546
Configuring IGMP Interface Parameters	547
Configuring Static Parameters	548
Configuring SSM Translation	549
Configuring PIM	550
Configuring Bootstrap Message Import and Export Policies	556
Triple Play Multicast Command Reference	557

Triple Play Enhanced Subscriber Management

RADIUS Authentication of Subscriber Sessions	605

Table of Contents

RADIUS Authentication Extensions	607
Triple Play Network with RADIUS Authentication	608
RADIUS Authorization Extensions	610
Calling-Station-ID	611
Provisioning of Enhanced Subscriber Management (ESM) Objects	613
RADIUS-Based Accounting	620
RADIUS Per Host Accounting	622
No Host-Accounting	622
Host-Accounting Enabled	622
Class Attribute	625
User Name	625
Enhanced Subscriber Management Overview	626
Enhanced Subscriber Management Basics	626
Standard and Enhanced Subscriber Management	626
ESM for IPv6	633
Models	633
Setup	635
Behavior	636
Dynamic Subscriber Host Processing	637
Dynamic Tables	637
Enhanced Subscriber Management Entities	640
Instantiating a New Host	641
Packet Processing for an Existing Host	642
ANCP and GSMP	643
ANCP	643
General Switch Management Protocol Version 3 (GSMPv3)	646
DHCP Release Messages	647
DHCP Client Mobility	647
DHCP Lease Control	647
Using Scripts for Dynamic Recognition of Subscribers	648
Python Language and Programmable Subscriber Configuration Policy (PSCP)	648
Determining the Subscriber Profile and SLA Profile of a Host	649
Determining the Subscriber Profile	650
Determining the SLA Profile	652
Limiting Subscribers and Hosts on a SAP	655
Static Subscriber Hosts	655
QoS for Subscribers and Hosts	656
QoS Parameters in Different Profiles	656
QoS Policy Overrides	656
Multi-Chassis Synchronization	657
Overview	657
Routed Central Office (CO)	660
Layer 3 Subscriber Interfaces	662
Routed Subscriber Hosts	665
Interaction Notes	667
VPRN Routed COs	671
Dual Homing	677
Dual Homing to Two PEs (Redundant-Pair Nodes) in Triple Play Aggregation	677
Steady-State Operation of Dual-homed Ring	680

Broken-Ring Operation and the Transition to this State	682
Transition from Broken to Closed Ring State	684
Provisioning Aspects and Error Cases	684
Dual Homing to Two BSR Nodes	685
MC Services	686
Routed CO Dual Homing	688
SRRP and Multi-Chassis Synchronization	691
Dual Homing and ANCP	691
Subscriber Router Redundancy Protocol	692
SRRP Messaging	697
SRRP Instance	698
Managed SAP (M-SAP)	701
ESM Identification Process	703
Default-Subscriber	703
Multicast Management	704
Subscriber Mirroring	704
Volume and Time Based Accounting	705
Metering	705
Categories Map and Categories	706
Quota Consumption	707
RADIUS VSA Credit-Control-Quota	707
Credit Negotiation Mechanisms	708
Action on Credit Exhaustion	709
Action on Error-Conditions	709
Applicability of Volume and Time Based Accounting	710
Configuring Enhanced Subscriber Management with CLI	711
Configuring RADIUS Authentication of DHCP Sessions	712
Configuring Enhanced Subscriber Management	713
Basic Configurations	713
Subscriber Interface Configuration	713
Configuring Enhanced Subscriber Management Entities	714
Routed CO with Basic Subscriber Management Features	720
Applying the Profiles and Policies	721
Configuring Dual Homing	723
Configuring Wholesale/Retail Parameters	725
Subscriber Management Command Reference	731

Service Manager Scripting Language

Service Manager Scripting Language	1008
Python Changes	1009
Configuration	1009
Operator Debugging	1011
Python Scripts	1012
Limitations	1013
Tips and Tricks	1014
Sample Python Scripts	1015

Common CLI Command Descriptions

Common Service Commands	1020

Table of Contents

Standards and Protocol Support	1025
Index	1031

List of Tables

Getting Started

Table 1:	Configuration Process	15

Introduction to Triple Play

Table 2:	Alcatel-Lucent's TPSDA	22
Table 3:	Downstream QoS Enablement	26
Table 4:	Upstream QoS Enablement	27
Table 5:	Default QinQ and TopQ SAP Dot1P Evaluation	150
Table 6:	Top Position QinQ and TopQ SAP Dot1P Evaluation	150
Table 7:	Bottom Position QinQ and TopQ SAP Dot1P Evaluation	151

Triple Play Multicast

Table 8:	Multicast Planes Based on Chassis Type	596

Triple Play Enhanced Subscriber Management

Table 9:	SRRP State Effect on Subscriber Hosts Associated with Group IP Interface	693

Service Manager Scripting Language

Table 10:	DHCP Object Members	1008

List of Figures

Introduction to Triple Play

Figure 1:	Triple Play Service Delivery Architecture	21
Figure 2:	Alcatel-Lucent's Triple Play Service Delivery Architecture	23
Figure 3:	Downstream QoS Enablement	25
Figure 4:	Upstream QoS Enablement	26
Figure 5:	Distributed Multicasting in TPSDA	28
Figure 6:	Subnetting Topology	29
Figure 7:	VMAC Subnetting Topology	30
Figure 8:	Alcatel-Lucent's Service Entities	36
Figure 9:	Service Access Point (SAP)	37
Figure 10:	Multiple SAPs on a Single Port/Channel	39
Figure 11:	A GRE Service Distribution Point (SDP) pointing from ALA-A to ALA-B	42
Figure 12:	Epipe/VLL Service	45
Figure 13:	Internet Enhanced Service	48
Figure 14:	Service Creation and Implementation Flow	51
Figure 15:	Subscriber Service Components	52

DHCP Management

Figure 16:	IP Address Assignment with DHCP	302
Figure 17:	DHCP Lease State Table	309
Figure 18:	CMTS/WAC Network Configuration Example	312
Figure 19:	Typical DHCP Deployment Scenarios	317
Figure 20:	Example of Triple Play Aggregation Network With DHCP to RADIUS Authentication	319

Point-to-Point Protocol over Ethernet (PPPoE) Management

Figure 21:	PPPoE VPRN with IP Overlap	434

Triple Play Security

Figure 22:	IP Illustration of Message Flow in Web Portal Redirect	476
Figure 23:	VPLS Redirect Policy Example	484

Triple Play Multicast

Figure 24:	IGMP Message Processing	530
Figure 25:	MVR and MVR by Proxy	532

Triple Play Enhanced Subscriber Management

Figure 26:	Triple Play Aggregation Network with RADIUS-Based DHCP Host Authentication	608
Figure 27:	Enhanced Subscriber Management Dynamic Tables	637
Figure 28:	Relationship Between Enhanced Subscriber Management Entities	640
Figure 29:	Static ANCP Management Example	644
Figure 30:	ESM Dynamic ANCP Example	645
Figure 31:	Data Flow in Determining Subscriber Profile and SLA Profile	649
Figure 32:	Determining the Subscriber Profile	650
Figure 33:	Determining the SLA Profile	653
Figure 34:	Dual-Homing Configuration	657

List of Figures

Figure 35:	7750 SR DSLAM Connection	660
Figure 36:	Subscriber Interface in an IES Service	661
Figure 37:	Details of a Group-Interface	661
Figure 38:	Aggregation Network with Direct DSLAM-BSR Connection	662
Figure 39:	Detailed View of Configurable Objects Related to Layer 3 Subscriber Interfaces	663
Figure 40:	Routed CO Example	672
Figure 41:	Traffic Flow	673
Figure 42:	Forwarding Model Without Wholesale-Retail Type Flag	675
Figure 43:	Forwarding Model With Wholesale-Retail Type Flag	676
Figure 44:	Dual-Homing to Two PEs	677
Figure 45:	Layer 2 CO Dual Homing - Network Diagram	679
Figure 46:	Dual Homing Ring Under Steady-State Condition	680
Figure 47:	Broken Ring State	682
Figure 48:	Ring Topology Using Layer 3 CO Model	685
Figure 49:	MC Services in a Layer 3-Ring Topology (a)	686
Figure 50:	MC Services on a Layer 3-Ring Topology (b)	687
Figure 51:	Dual Homing Example	690
Figure 52:	Threshold Configured / Not Configured	708

Preface

About This Guide

This guide describes details pertaining to Triple Play Services Delivery Architecture (TPSDA) support provided by the 7750 SR OS and presents examples to configure and implement various protocols and services.

This document 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 manual is intended for network administrators who are responsible for configuring the 7750 SR routers. It is assumed that the network administrators have an understanding of networking principles and configurations. Protocols, standards, and services described in this manual include the following:

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

List of Technical Publications

The 7750 SR documentation set is composed of the following books:

- 7750 SR OS Basic System Configuration Guide
This guide describes basic system configurations and operations.
- 7750 SR OS System Management Guide
This guide describes system security and access configurations as well as event logging and accounting logs.

- **7750 SR OS Interface Configuration Guide**
This guide describes card, Media Dependent Adapter (MDA), and port provisioning.
 - **7750 SR OS Router Configuration Guide**
This guide describes logical IP routing interfaces and associated attributes such as an IP address, port, link aggregation group (LAG) as well as IP and MAC-based filtering, and VRRP, and Cflowd.
 - **7750 SR OS Routing Protocols Guide**
This guide provides an overview of routing concepts and provides configuration examples for RIP, OSPF, IS-IS, Multicast, BGP, and route policies.
 - **7750 SR OS MPLS Guide**
This guide describes how to configure Multiprotocol Label Switching (MPLS) and Label Distribution Protocol (LDP).
 - **7750 SR OS Services Guide**
This guide describes how to configure service parameters such as service distribution points (SDPs), customer information, and user services.
 - **7750 SR OS OAM and Diagnostic Guide**
This guide describes how to configure features such as service mirroring and Operations, Administration and Management (OAM) tools.
 - **7750 SR Quality of Service Guide**
This guide describes how to configure Quality of Service (QoS) policy management.
 - **7750 SR Triple Play Guide**
This guide describes Triple Play and subscriber management services and support provided by the 7750 SR and presents examples to configure and implement various protocols and services.
 - **OS Multi-Service ISA Guide**
This guide describes services provided by integrated service adapters such as Application Assurance, IPSec, ad insertion (ADI) and Network Address Translation (NAT).
-

Technical Support

If you purchased a service agreement for your 7750 SR 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, contact your welcome center.

Web: http://www1.alcatel-lucent.com/comps/pages/carrier_support.jhtml

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 7750 SR-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	Introduction to Triple Play on page 17
DHCP	Theory, configuration	DHCP Management on page 301
Security	Anti-spoofing, MAC pinning, MAC protection, DoS protection VPLS redirect policies, Web portal redirect, ARP handling	Triple Play Security on page 465
PPPoE	Theory, configuration	Point-to-Point Protocol over Ethernet (PPPoE) Management on page 423
Multicast	Multicast, IGMP, SSM, PIM	Triple Play Multicast on page 523
Enhanced Subscriber management	RADIUS authentication, ESM entities, Routed CO, M-SAP	Triple Play Enhanced Subscriber Management on page 603

Table 1: Configuration Process

Area	Task	Chapter (Continued)
	Python Scripting	Service Manager Scripting Language on page 1007
Reference	List of IEEE, IETF, and other proprietary entities.	Standards and Protocol Support on page 1025

Introduction to Triple Play

In This Section

This section provides an overview of the 7750 SR 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 7750 SR OS Services Guide.

Topics in this section include:

- [Alcatel-Lucent's Triple Play Service Delivery Architecture on page 18](#)
 - [Introduction to Triple Play on page 18](#)
 - [Blueprint for Optimizing Triple Play Service Infrastructures on page 19](#)
 - [Architectural Foundations on page 20](#)
 - [Optimizing Triple Play Service Infrastructures on page 22](#)
- [Services on page 32](#)
 - [Service Types on page 33](#)
 - [Service Policies on page 34](#)
- [Alcatel-Lucent Service Model on page 35](#)
 - [Service Entities on page 36](#)
 - [Customers on page 36](#)
 - [Service Access Points \(SAPs\) on page 37](#)
 - [Service Distribution Points \(SDPs\) on page 41](#)
- [Epipe Service Overview on page 45](#)
- [VPLS Service Overview on page 46](#)
 - [Split Horizon SAP Groups and Split Horizon Spoke SDP Groups on page 46](#)
- [IES Service Overview on page 48](#)
 - [IP Interface on page 49](#)
- [VPRN Service Overview on page 50](#)
- [Service Creation Process Overview on page 51](#)

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 7750 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 7750 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 7750 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 7750 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 7750/7710 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.



OSSG069A

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 7450 ESS) 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 7750 SR 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 7750 SR router and another service access point (a SAP is where traffic enters and exits the service) on the same (local) or another 7750 SR router (distributed). A distributed service spans more than one router.

Distributed services use service distribution points (SDPs) to direct traffic to another 7750 SR through a service tunnel. SDPs are created on each participating 7750 SR, specifying the origination address (the 7750 SR router participating in the service communication) and the destination address of another SR-Series. SDPs are then bound to a specific customer service. Without the binding process, far-end 7750 SR devices are not able to participate in the service (there is no service without associating an SDP with a service).

Service Types

The 7750 SR 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.
 - ATM VLL (Apipe) — A point-to-point ATM service between users connected to 7750 nodes on an IP/MPLS network.
 - Frame-Relay (Fpipe) — A point-to-point Frame Relay service between users connected to 7750 nodes on the IP/MPLS network.
 - 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.
- Virtual Private Routed Network (VPRN) — Layer 3 IP multipoint-to-multipoint VPN service as defined in RFC 2547bis.

Service Policies

Common to all 7750 SR 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 7750 SR service enhancements. The types of policies that are common to all 7750 SR 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 7750 SR 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 36](#)
 - [Customers on page 36](#)
 - [Service Access Points \(SAPs\) on page 37](#)
 - [Service Distribution Points \(SDPs\) on page 41](#)
-

Introduction

In the 7750 SR service model, the 7750 SR service edge routers are deployed at the provider edge. Services are provisioned on 7750 SRs 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 36)
- [Service Access Points \(SAPs\)](#) (see page 37)
- [Service Distribution Points \(SDPs\)](#) (see page 41) (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 7750 SR router (Figure 9). The SAP configuration requires that slot, MDA, and port/channel information be specified. The slot, MDA, and port/channel parameters must be configured prior to provisioning a service (see the [Cards, MDAs, and Ports](#) section of the 7750 SR OS Interface Configuration Guide).

A SAP is a local entity to the 7750 SR and is uniquely identified by:

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

Depending on the encapsulation, a physical port or channel can have more than one SAP associated with it. SAPs can only be created on ports or channels 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.




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 or TDM channel. The appropriate encapsulation type for the port or channel depends on the requirements to support multiple services on a single port/channel on the associated SAP and the capabilities of the downstream equipment connected to the port/channel. 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 or channel 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/Channel

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 7750 SR.
- 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 7750 SR.
- A port/channel 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/channel is administratively shutdown, all SAPs on that port/channel will be operationally out of service.
- A SAP cannot be deleted until it has been administratively disabled (shutdown).

- 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 7750 SR to another 7750 SR through a uni-directional (one-way) service tunnel. The SDP terminates at the far-end 7750 SR 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 7750 SR. The same SDP ID can appear on other 7750 SR routers.
- An SDP uses the system IP address to identify the far-end 7750 SR 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 7750 SR requires a return path SDP from the far-end 7750 SR back to the local 7750 SR. 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 7750 SR device(s) cannot participate in the service (there is no service).



OSSG004

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 7750 SR 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 7750 SR 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 7750 SR 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/channel on which service traffic enters the 7750 SR 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 (VPRN). 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) 7750 SR 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, we refer to VPLS section in the 7750 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 7750 SR 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)
- Cflowd
- Secondary IP addresses
- ICMP options

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

- Unnumbered interfaces
- NTP broadcast receipt

VPRN Service Overview

RFC2547bis is an extension to the original RFC 2547, which details a method of distributing routing information and forwarding data to provide a Layer 3 Virtual Private Network (VPN) service to end customers.

Each Virtual Private Routed Network (VPRN) consists of a set of customer sites connected to one or more PE routers. Each associated PE router maintains a separate IP forwarding table for each VPRN. Additionally, the PE routers exchange the routing information configured or learned from all customer sites via MP-BGP peering. Each route exchanged via the MP-BGP protocol includes a Route Distinguisher (RD), which identifies the VPRN association.

The service provider uses BGP to exchange the routes of a particular VPN among the PE routers that are attached to that VPN. This is done in a way which ensures that routes from different VPNs remain distinct and separate, even if two VPNs have an overlapping address space. The PE routers distribute routes from other CE routers in that VPN to the CE routers in a particular VPN. Since the CE routers do not peer with each other there is no overlay visible to the VPN's routing algorithm.

When BGP distributes a VPN route, it also distributes an MPLS label for that route. On a SR-Series, a single label is assigned to all routes in a VPN.

Before a customer data packet travels across the service provider's backbone, it is encapsulated with the MPLS label that corresponds, in the customer's VPN, to the route which best matches the packet's destination address. The MPLS packet is further encapsulated with either another MPLS label or GRE tunnel header, so that it gets tunneled across the backbone to the proper PE router. Each route exchanged by the MP-BGP protocol includes a route distinguisher (RD), which identifies the VPRN association. Thus the backbone core routers do not need to know the VPN routes.

Service Creation Process Overview

Figure 14 displays the overall process to provision core and subscriber services.




Figure 14: Service Creation and Implementation Flow

Service Components

Figure 15 displays the basic components of a subscriber service.




Figure 15: Subscriber Service Components

- Customer ID — This is the subscriber ID used to associate information with a particular customer.
- Service type — The service type specifies the connectivity type such as Epipe, VPLS, or IES.
- Service ID — Each service is uniquely identified by a Service ID within a service area.
- SAP — The subscriber-side service entry (access) and exit point for a service on an Alcatel-Lucent 7750 SR router.
- Interface — Configure a logical IP routing interface for Internet Ethernet Service (IES), and Virtual Private Routed Network (VPRN) services.
- Port ID — The physical port identifier portion of the SAP definition.
- QoS policy — The QoS policy associated with an ingress or egress SAP or IP interface. QoS policy ID 1 is the default.
- Filter policy — Optional. The IP or MAC filter policy associated with an ingress or egress SAP or IP interface.
- Scheduler policy — Optional. A scheduler policy defines the hierarchy and operating parameters for virtual schedulers. This implementation of hierarchical scheduling is a method that defines a bounded operation for a group of queues. One or more queues are mapped to a given scheduler with strict and weighted metrics controlling access to the

scheduler. The scheduler has an optional prescribed maximum operating rate that limits the aggregate rate of the child queues. This scheduler may then feed into another virtual scheduler in a higher tier.

- SDP — An SDP is a logical mechanism that ties a far-end 7750 SR to a specific service without having to specifically define the far-end SAPs. Each SDP, identified by a local SDP ID, represents a method for reaching a far-end 7750 SR service edge router.

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 7750 SR 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, IES, or VPRN 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, IES and VPRN services.

Topics in this section include:

- [Configuring VPLS Residential Split Horizon Groups on page 57](#)
- [Configuring Static Hosts on page 58](#)
- [Configuring Static Hosts on an IES SAP on page 59](#)
- [Configuring Static Hosts on a VPRN SAP on page 60](#)

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

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 or VPRN, 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.

Configuring Static Hosts on a 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# info
-----
  interface "test2" create
    address 102.22.1.1/24
    sap 7/1/5 create
      description "IES with static host"
      static-host ip 10.1.1.1
    exit
  exit
-----
*A:ALA-49>config>service>ies#
```

Configuring Static Hosts on a VPRN SAP

The following displays a static host on a VPRN SAP:

```
*A:ALA-49>config>service>vprn# info
-----
      description "VPRN service with static host"
      interface "to_Dslam" create
        sap 2/1/6:0 create
          description "SAP with static host"
          static-host ip 10.1.1.1
        exit
      exit
    no shutdown
-----
*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 7750 SR OS Services Guide guide.

- [Generic VPLS Commands on page 61](#)
- [Generic IES Commands on page 65](#)
- [DHCP VPRN Commands on page 68](#)
- [ARP VPRN Commands on page 69](#)
- [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]
      — local-age aging-timer
      — no local-age
      — mac-protect
        — [no] mac ieee-address
      — mac-subnet-length subnet-length
      — no mac-subnet-length
      — 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**
 - **calling-station-id** *calling-station-id*
 - **no calling-station-id**
 - **description** *description-string*
 - **no description**
 - **[no] collect-stats**
 - **[no] disable-aging**
 - **[no] disable-learning**
 - **egress**
 - **agg-rate-limit** *agg-rate*
 - **no agg-rate-limit** *agg-r*
 - **filter ip** *ip-filter-id*
 - **filter ipv6** *ipv6-filter-id*
 - **filter mac** *mac-filter-id*
 - **no filter**
 - **no filter** [*ip ip-filter-id*] [*mac mac-filter-id*] [*ipv6 ipv6-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*] [*anccp-string anccp-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 ipv6** *ipv6-filter-id*
 - **filter mac** *mac-filter-id*

```

— no filter
— no filter [ip ip-filter-id] [mac mac-filter-id] [ipv6 ipv6-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
            — agg-rate-limit agg-rate
            — no agg-rate-limit agg-rate
            — 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
    — managed-routes
        — route {ip-prefix/length | ip-prefix netmask} [create]
        — no route {ip-prefix/length | ip-prefix netmask}
    — [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
— spoke-sdp sdp-id:vc-id
    — egress
        — filter {ip ip-filter-id}
        — no filter
        — vc-label egress-vc-label
        — no vc-label [egress-vc-label]
    — ingress
        — filter {ip ip-filter-id}
        — no filter
        — vc-label ingress-vc-label
        — no vc-label [ingress-vc-label]

```

Triple Play 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]
— no sap sap-id
— anti-spoof {ip | mac | ip-mac}
— no anti-spoof
— authentication-policy auth-plcy-name
— no authentication-policy
— dhcp
— description description-string
— no description
— lease-populate [number-of-entries]
— 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
— server server1 [server2...(up to 8 max)]
— no server
— [no] shutdown
— [no] trusted

config
— service
    — vpn
        — [no] interface ip-int-name
            — dhcp
                — description description-string
                — no description
                — gi-address ip-address [src-ip-addr]
                — no gi-address
                — lease-populate [number-of-entries]
                — no lease-populate
                — [no] option
                    — action {replace | drop | keep}
                    — no action
                    — circuit-id [ascii-tuple | ifindex]
                    — no circuit-id
                    — [no] remote-id
                — 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

```

```

config
  — service
    — vrpn
      — [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
        — static-arp ip-address ieee-address
        — no static-arp ip-address [ieee-address]

```

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] 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-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

```


7750 SR-Only Router IGMP and PIM Command CLI Reference

NOTE: The following commands are listed for informational purposes. The command descriptions are documented in the Multicast Command Reference section of the *7750 SR OS Routing Protocols Guide*.

```

config
  — router
    — [no] igmp
      — [no] interface ip-int-name
        — import policy-name
        — no import
        — max-groups value
        — no max-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
            — [no] shutdown
          — policy policy-name
          — no policy
          — unconstrained-bw bandwidth mandatory-bw mandatory-bw
          — no unconstrained-bw
        — [no] shutdown
      — ssm-translate
        — grp-range start end
        — no grp-range start end
          — [no] source ip-address
      — static
        — [no] group grp-ip-address
          — [no] source ip-address
          — [no] starg
        — [no] subnet-check
        — version version
        — no version
      — query-interval seconds
      — no query-interval
      — query-last-member-interval seconds
      — no query-last-member-interval
      — query-response-interval seconds
      — no query-response-interval
      — robust-count robust-count
      — no robust-count
      — ssm-translate
        — [no] grp-range start end
          — [no] source ip-address
      — [no] shutdown
config
  — router
    — [no] pim
      — apply-to {ies | non-ies | all | none}
      — import {join-policy | register-policy} policy-name [.. policy-name]
      — no import {join-policy | register-policy}
      — [no] interface ip-int-name

```

```

— assert-period assert-period
— no assert-period
— [no] bfd-enable
— [no] bsm-check-rtr-alert
— hello-interval hello-interval
— no hello-interval
— hello-multiplier deci-units
— no hello-multiplier
— [no] improved-assert
— [no] ipv4-multicast-disable
— [no] ipv6-multicast-disable
— max-groups value
— no max-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
        — [no] shutdown
    — policy policy-name
    — no policy
    — unconstrained-bw bandwidth mandatory-bw mandatory-bw
    — no unconstrained-bw
— multicast-senders {auto | always | never}
— no multicast-senders
— priority dr-priority
— no priority
— [no] shutdown
— sticky-dr [priority dr-priority]
— no sticky-dr
— three-way-hello [compatibility-mode]
— no three-way-hello
— [no] tracking-support
— [no] ipv4-multicast-disable
— [no] ipv6-multicast-disable
— [no] non-dr-attract-traffic
— rp
    — [no] anycast rp-ip-address
        — [no] rp-set-peer ip-address
    — bootstrap-export policy-name [.. policy-name]
    — no bootstrap-export
    — bootstrap-import policy-name [.. policy-name]
    — no bootstrap-import
    — bsr-candidate
        — address ip-address
        — no address
        — hash-mask-len hash-mask-length
        — no hash-mask-len
        — priority bootstrap-priority
        — no priority
        — [no] shutdown
    — rp-candidate
        — address ip-address
        — no address
        — [no] group-range {grp-ip-address/mask | grp-ip-address net-mask}

```

```

— holdtime holdtime
— no holdtime
— priority priority
— no priority
— [no] shutdown
— static
— [no] address ip-address
— [no] group-prefix {grp-ip-address/mask | grp-ip-address netmask}
— [no] override
— rpf-table {rtable-m | rtable-u | both}
— no rpf-table
— [no] shutdown
— spt-switchover-threshold {grp-ip-address/mask | grp-ip-address netmask} spt-threshold
— spt-switchover-threshold {grp-ipv4-prefix/ipv4-prefix-length | grp-ipv4-prefix netmask | grp-ipv6-prefix/ipv6-prefix-length} spt-threshold
— no spt-switchover-threshold grp-ipv4-prefix/ipv4-prefix-length | grp-ipv4-prefix netmask | grp-ipv6-prefix/ipv6-prefix-length}
— no spt-switchover-threshold {grp-ip-address/mask | grp-ip-address netmask}
— [no] ssm-groups
— [no] group-range {grp-ip-address/mask | grp-ip-address netmask}

show
— router
— igmp
— group [grp-ip-address]
— interface [ip-int-name | ip-address] [group] [grp-address] [detail]
— ssm-translate
— static [ip-int-name | ip-addr]
— statistics [ip-int-name | ip-address]
— status
— pim
— anycast [detail]
— crp [ip-address]
— group [grp-ip-address] [source ip-address] [type {starstarrp|starg|sg}] [detail]
— interface [ip-int-name | mt-int-name | ip-address] [group [grp-ip-address] source ip-address] [type {starstarrp | starg | sg}] [detail]
— mc-ecmp-balance [detail]
— neighbor [ip-address | ip-int-name] [address ip-address] [detail]
— rp [family] | ip-address]
— rp-hash ip-address
— s-pmsi [mdSrcAddr] [mdGrpAddr] [detail]
— statistics [ip-int-name | mt-int-name | ip-address]
— status [detail] [family]

clear
— router
— igmp
— database [interface ip-int-name|ip-address] [group ip-address] [source ip-address]
— statistics [interface ip-int-name|ip-address]
— version [interface ip-int-name|ip-address]

```

Show Commands

```

show
  — service
    — 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
        — mrouter [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

Syntax	[no] shutdown
Context	<pre> config>service>vpls>sap>dhcp config>service>vpls>igmp-snooping>mvr config>service>vpls config>service>vpls>sap config>service>vpls>sap>arp-host config>service>vpls>mesh-sdp config>service>vpls>spoke-sdp config>service>ies>if>sap config>service>ies>if>dhcp config>service>ies>sub-if>grp-if>srrp config>service>vprn config>service>vprn>sub-if config>service>vprn>sub-if>dhcp config>service>vprn>sub-if>dhcp>proxy-server </pre>
Description	This command administratively disables an entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics. 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. Services are created in the administratively down (shutdown) state. When a no shutdown 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. The no form of this command places the entity into an administratively enabled state.
Special Cases	Service Admin State — 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. Service Operational State — A service is regarded as operational providing that two SAPs or if one SDP are operational. SDP (global) — 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. SDP (service level) — 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.

SDP Keepalives — 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.

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

Syntax	description <i>description-string</i> no description
Context	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>vprn config>service>vprn>subscriber-interface config>service>vprn>subscriber-interface>group-interface config>service>vprn>subscriber-interface>grp-if>dhcp config>service>vprn>sub-if>grp-if>srrp config>service>vpls>sap>dhcp config>service>vpls>mld-snooping>mvr
Description	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 no form of this command removes the string from the configuration.
Default	No description associated with the configuration context.
Parameters	<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

Syntax	vpls <i>service-id</i> customer <i>customer-id</i> vpn <i>vpn-id</i> [m-vpls] vpls <i>service-id</i> no vpls <i>service-id</i>
Context	config>service
Description	This command creates or edits a Virtual Private LAN Services (VPLS) instance. The vpls 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. 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. When a service is created, the create keyword must be specified if the create command is enabled in the environment context. When a service is created, the customer 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 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. To create a management VPLS, the m-vpls keyword must be specified. See section Hierarchical VPLS Redundancy for an introduction to the concept of management VPLS. Once a service is created, the use of the customer <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. More than one VPLS service may be created for a single customer ID. By default, no VPLS instances exist until they are explicitly created. The no 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. <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 7750 SR on which this service is defined. Values service-id: 1 — 214748364 svc-name: A string up to 64 characters in length. customer <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. Values 1 — 2147483647

Triple Play Service Configuration Commands

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

Syntax	service-name <i>service-name</i> no service-name
Context	config>service>epipe config>service>ies config>service>vpls config>service>vpn
Description	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, 7450 ESS and 7710 SR 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.
Parameters	<i>service-name</i> — Specifies a unique service name to identify the service. Service names may not begin with an integer (0-9).

ies

Syntax	ies <i>service-id</i> customer <i>customer-id</i> vpn <i>vpn-id</i> ies <i>service-id</i> no ies <i>service-id</i>
Context	config>service
Description	This command creates or edits an IES service instance. The ies 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 **config router service-prefix** 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.

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 — 214748364
 svc-name: A string up to 64 characters in length.

customer *customer-id* — 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.

Values 1 — 2147483647

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)

vprn

Syntax	vprn <i>service-id</i> [customer <i>customer-id</i>] no vprn <i>service-id</i>
Context	config>service
Description	This command creates or edits a Virtual Private Routed Network (VPRN) service instance.

If the *service-id* does not exist, a context for the service is created. If the *service-id* exists, the context for editing the service is entered.

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 you try 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	<i>service-id</i> — The unique service identification number identifies the service in the service domain. The ID must be unique to this service and cannot be used for any other service of any type. The <i>service-id</i> must be the same number used for every 7750 SR on which this service is defined. Values service-id: 1 — 214748364 svc-name: A string up to 64 characters in length. customer <i>customer-id</i> — The 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 — The customer identification number must already have been created.

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 disabled, new source MAC addresses will be learned and entered into the VPLS forwarding database. 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

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

fdb-table-low-wmark

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

fdb-table-size

Syntax	fdb-table-size <i>table-size</i> no fdb-table-size [<i>table-size</i>]
Context	config>service>vpls
Description	This command specifies the maximum number of MAC entries in the forwarding database (FDB) for the VPLS instance on this node. The fdb-table-size specifies the maximum number of forwarding database entries for both learned and static MAC addresses for the VPLS instance. The no form of this command returns the maximum FDB table size to default.
Default	250 — Forwarding table of 250 MAC entries.
Parameters	<i>table-size</i> — The number of entries permitted in the forwarding database for this VPLS instance.
Values	Chassis-mode A or B limit: 131071 Chassis-mode D limit: 511999

local-age

Syntax	local-age <i>seconds</i> no local-age
Context	config>service>vpls
Description	Specifies the aging time for locally learned MAC addresses in the forwarding database (FDB) for the Virtual Private LAN Service (VPLS) instance. 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. 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 local-age timer specifies the aging time for local learned MAC addresses. The no form of this command returns the local aging timer to the default value.
Default	local age 300 — Local MACs aged after 300 seconds.
Parameters	<i>seconds</i> — The aging time for local MACs expressed in seconds. Values 60 — 86400

remote-age

Syntax	remote-age <i>seconds</i> no remote-age
Context	config>service>vpls
Description	Specifies the aging time for remotely learned MAC addresses in the forwarding database (FDB) for the Virtual Private LAN Service (VPLS) instance. 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. 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 remote-age 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 local-age timer. The no form of this command returns the remote aging timer to the default value.
Default	remote age 900 — Remote MACs aged after 900 seconds.
Parameters	<i>seconds</i> — The aging time for remote MACs expressed in seconds. Values 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

Syntax	[no] split-horizon-group [<i>group-name</i>] [<i>residential-group</i>]
Context	config>service>vpls
Description	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. A split horizon group must be created before SAPs and spoke SDPs can be assigned to the group. 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. Up to 30 split horizon groups can be defined per VPLS instance. The no form of the command removes the group name from the configuration.
Parameters	<i>group-name</i> — Specifies the name of the split horizon group to which the SDP belongs. <i>residential-group</i> — Defines a split horizon group as a residential split horizon group (RSHG). Doing so entails that: a) SAPs which are members of this Residential Split Horizon Group will have: – Double-pass queuing at ingress as default setting (can be disabled) – STP disabled (can <u>not</u> be enabled) – ARP reply agent enabled by default (can be disabled) – MAC pinning enabled by default (can be disabled) – broadcast packets are discarded, blocking unknown, flooded traffic b) Spoke SDPs which are members of this Residential Split Horizon Group will have: – Downstream multicast traffic supported – Double-pass queuing is not applicable – STP is disabled (can be enabled) – ARP reply agent is not applicable (dhcp-lease-states are not supported on spoke SDPs) – MAC pinning enabled per default (can be disabled)
Default	A split horizon group is by default not created as a residential-group.

sap

Syntax	sap <i>sap-id</i> [split-horizon-group <i>group-name</i>] [create] no sap <i>sap-id</i>
Context	config>service>vpls
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

7750. 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. For Internet Ethernet Service (IES), the IP interface must be shutdown before the SAP on that interface may be removed.

Default	No SAPs are defined.
Special Cases	A VPLS SAP can be defined with Ethernet ports, SONET/SDH or TDM channels. 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).
Parameters	<i>sap-id</i> — Specifies the physical port identifier portion of the SAP definition. See Common Service Commands on page 1020 for <i>sap-id</i> command syntax. <i>port-id</i> — Specifies the physical port ID in the <i>slot/mda/port</i> format. 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. The <i>port-id</i> must reference a valid port type. When the <i>port-id</i> parameter represents SONET/SDH and TDM channels, 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. If the SONET/SDH port is configured as clear-channel then only the port is specified. split-horizon-group <i>group-name</i> — Specifies an existing split horizon group name. create — Keyword used to create a SAP instance. The create keyword requirement can be enabled/disabled in the environment>create context.

accounting-policy

Syntax	accounting-policy acct-policy-id no accounting-policy
Context	config>service>vpls>sap config>service>vpls>spoke-sdp config>service>vpls>mesh-sdp

```
config>service>ies>if>sap
config>service>ies>sub-if>grp-if>sap
```

Description	This command creates the accounting policy context that can be applied to a SAP or SDP. 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. A maximum of one accounting policy can be associated with a SAP or SDP at one time. Accounting policies are configured in the config>log context. The no form of this command removes the accounting policy association from the SAP or SDP, and the accounting policy reverts to the default.
Default	Default accounting policy.
Parameters	<i>acct-policy-id</i> — Enter the accounting <i>policy-id</i> as configured in the config>log>accounting-policy context.
Values	1 — 99

collect-stats

Syntax	[no] collect-stats
Context	<pre>config>service>vpls>sap config>service>vpls>spoke-sdp config>service>vpls>mesh-sdp config>service>ies>if>sap</pre>
Description	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. When the no collect-stats 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 collect-stats command is issued then the counters written to the billing file include all the traffic while the no collect-stats command was in effect.
Default	collect-stats

interface

Syntax	interface <i>ip-int-name</i> no interface <i>ip-int-name</i>
Context	config>service>vprn
Description	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.

The **interface** command, under the context of services, is used to create and maintain IP routing interfaces within VPRN service IDs. The **interface** 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.

Interface names are case sensitive and must be unique within the group of defined IP interfaces defined for **config router interface** and **config service vprn interface**. 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.

The available IP address space for local subnets and routes is controlled with the **config router service-prefix** command. The **service-prefix** 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 **config router** and **config service** domains.

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 VPRN IP interfaces must be explicitly defined. Interfaces are created in an enabled state.

The **no** form of this command removes the interface and all the associated configuration. The interface must be administratively shutdown before issuing the **no interface** command.

For VPRN services, the IP interface must be shutdown before the SAP on that interface may be removed. VPRN services do not have the **shutdown** command in the SAP CLI context. VPRN service SAPs rely on the interface status to enable and disable them.

Parameters

ip-int-name — Specifies the name of the IP interface. Interface names must be unique within the group of defined IP interfaces for **config router interface** and **config service vprn interface** 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.

cflowd

Syntax	cflowd {acl interface} no cflowd
Context	config>service>ies>interface
Description	This command enables cflowd to collect traffic flow samples through a router for analysis.

cflowd is used for network planning and traffic engineering, capacity planning, security, application and user profiling, performance monitoring, usage-based billing, and SLA measurement. When **cflowd** is enabled at the interface level, all packets forwarded by the interface are subjected to analysis according to the **cflowd** configuration.

Default	no cflowd
Parameters	<i>acl</i> — <i>cflowd</i> configuration associated with a filter. <i>interface</i> — <i>cflowd</i> configuration associated with an IP interface.

limit-mac-move

Syntax	limit-mac-move [blockable non-blockable] no limit-mac-move
Context	config>service>vpls>sap config>service>vpls>spoke-sdp
Description	This command indicates whether or not the mac-move agent, when enabled using config>service>vpls>mac-move or config>service>epipe>mac-move , will limit the MAC re-learn (move) rate on this SAP.
Default	SAPs and spoke SDPs are blockable
Parameters	blockable — the agent will monitor the MAC re-learn rate on the SAP, and it will block it when the re-learn rate is exceeded. non-blockable — When specified, this SAP will not be blocked, and another blockable SAP will be blocked instead.

mac-pinning

Syntax	[no] mac-pinning
Context	config>service>vpls>sap config>service>vpls>spoke-sdp config>service>vpls>mesh-sdp
Description	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. The age of the MAC address entry in the FIB is set by the age timer. If mac-aging is disabled on a given VPLS service, any MAC address learned on a SAP/SDP with mac-pinning enabled will remain in the FIB on this SAP/SDP forever. Every event that would otherwise result in re-learning will be logged (MAC address, original-SAP, new-SAP). 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.

Default 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

Syntax **managed-vlan-list**

Context config>service>vpls>sap

Description 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.

This command is only valid when the VPLS in which it is entered was created as a management VPLS.

range

Syntax [no] **range** *vlan-range*

Context config>service>vpls>sap>managed-vlan-list

Description This command configures a range of VLANs on an access port that are to be managed by an existing management VPLS.

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.

To modify the range of VLANs, first the new range should be entered and afterwards the old range removed.

Default None

Parameters *vlan-range* — 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>

Values start-vlan: 1 — 4094
end-vlan: 1 — 4094

VPLS SAP ATM Commands

atm

Syntax	atm
Context	config>service>vpls>sap
Description	This command enables access to the context to configure ATM-related attributes. This command can only be used when a given context (for example, a channel or SAP) supports ATM functionality such as: • Configuring ATM port or ATM port-related functionality on MDAs supporting ATM functionality • Configuring ATM-related configuration for ATM-based SAPs that exist on MDAs supporting ATM functionality. If ATM functionality is not supported for a given context, the command returns an error.

egress

Syntax	egress
Context	config>service>vpls>sap>atm
Description	This command enables the context to configure egress ATM attributes for the SAP.

encapsulation

Syntax	encapsulation <i>atm-encap-type</i>		
Context	config>service>vpls>sap>atm		
Description	This command specifies the data encapsulation for an ATM PVCC delimited SAP. The definition references RFC 2684, <i>Multiprotocol Encapsulation over ATM AAL5</i>, and to the ATM Forum LAN Emulation specification. Ingress traffic that does not match the configured encapsulation will be dropped.		
Default	The encapsulation is driven by the services for which the SAP is configured. For IES and VPRN service SAPs, the default is aal5snap-routed.		
Parameters	<i>atm-encap-type</i> — Specify the encapsulation type. <table> <tr> <td>Values</td><td> aal5snap-routed — Routed encapsulation for LLC encapsulated circuit (LLC/ SNAP precedes protocol datagram) as defined in RFC 2684. aal5mux-ip — Routed IP encapsulation for VC multiplexed circuit as defined in RFC 2684 </td></tr> </table>	Values	aal5snap-routed — Routed encapsulation for LLC encapsulated circuit (LLC/ SNAP precedes protocol datagram) as defined in RFC 2684. aal5mux-ip — Routed IP encapsulation for VC multiplexed circuit as defined in RFC 2684
Values	aal5snap-routed — Routed encapsulation for LLC encapsulated circuit (LLC/ SNAP precedes protocol datagram) as defined in RFC 2684. aal5mux-ip — Routed IP encapsulation for VC multiplexed circuit as defined in RFC 2684		

ingress

Syntax	ingress
Context	config>service>vpls>sap>atm
Description	This command enables the context to configure ingress ATM attributes for the SAP.

traffic-desc

Syntax	traffic-desc <i>traffic-desc-profile-id</i> no traffic-desc
Context	config>service>vpls>sap>atm>ingress config>service>vpls>sap>atm>egress
Description	This command assigns an ATM traffic descriptor profile to a given context (for example, a SAP). When configured under the ingress context, the specified traffic descriptor profile defines the traffic contract in the forward direction. When configured under the egress context, the specified traffic descriptor profile defines the traffic contract in the backward direction. The no form of the command reverts the traffic descriptor to the default traffic descriptor profile.
Default	The default traffic descriptor (trafficDescProfileId. = 1) is associated with newly created PVCC-delimited SAPs.
Parameters	<i>traffic-desc-profile-id</i> — Specify a defined traffic descriptor profile (see the QoS atm-td-profile command).

oam

Syntax	oam
Context	config>service>vpls>sap>atm
Description	This command enables the context to configure OAM functionality for a PVCC delimiting a SAP. The ATM-capable MDAs support F5 end-to-end OAM functionality (AIS, RDI, loopback): • ITU-T Recommendation I.610 - B-ISDN Operation and Maintenance Principles and Functions version 11/95 • GR-1248-CORE - Generic Requirements for Operations of ATM Network Elements (NEs). Issue 3 June 1996 • GR-1113-CORE - Bellcore, Asynchronous Transfer Mode (ATM) and ATM Adaptation Layer (AAL) Protocols Generic Requirements, Issue 1, July 1994

alarm-cells

Syntax	[no] alarm-cells
Context	config>service>vpls>sap>atm
Description	This command configures AIS/RDI fault management on a PVCC. Fault management allows PVCC termination to monitor and report the status of their connection by propagating fault information through the network and by driving PVCC's operational status. When alarm-cells functionality is enabled, PVCC's operational status is affected when a PVCC goes into AIS or RDI state because of an AIS/RDI processing (i.e. assuming nothing else affects PVCC's operational status, PVCC goes DOWN, when it enters a fault state and comes back UP, when it exits that fault state) and RDI cells are generated when PVCC is operationally DOWN. No OAM-specific SNMP trap is raised whenever an endpoint enters/exits an AIS or RDI state; however, if as result of an OAM state change, the PVCC changes operational status; then a trap is expected from an entity the PVCC is associated with (for example a SAP). The no command disables alarm-cells functionality for a PVCC. When alarm-cells functionality is disabled, PVCC's operational status is no longer affected by PVCC's OAM state changes due to AIS/RDI processing (Note that when alarm-cells is disabled, a PVCC will change operational status to UP, if it was DOWN because of the alarm-cell processing) and RDI cells are not generated as result of PVCC going into an AIS or RDI state; however, PVCC's OAM status will record OAM faults as described above.
Default	Enabled for PVCCs delimiting VPLS SAPs

Service Billing Commands

authentication-policy

Syntax	authentication-policy <i>name</i> no authentication-policy
Context	config>service>vpls>sap config>service>ies>sub-if>grp-if
Description	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.
Parameters	<i>name</i> — Specifies a unique authentication policy name.

root-guard

Syntax	[no] root-guard
Context	config>service>vpls>sap>stp
Description	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.
Default	no root-guard

SAP Subscriber Management Commands

Subscriber management commands are also described in the [Triple Play Services Command Reference on page 61](#) section.

sub-sla-mgmt

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

def-sla-profile

Syntax	def-sla-profile default-sla-profile-name no def-sla-profile
Context	config>service>vpls>sap>sub-sla-mgmt config>service>ies>if>sap>sub-sla-mgmt config>service>ies>sub-if>grp-if>sap>sub-sla-mgmt
Description	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 config>subscriber-mgmt>sla-profile context. 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. The no form of the command removes the default SLA profile from the SAP configuration.
Default	no def-sla-profile
Parameters	<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 config>subscriber-mgmt>sla-profile context.

def-sub-profile

Syntax	def-sub-profile <i>default-subscriber-profile-name</i>
Context	config>service>vpls>sap>sub-sla-mgmt config>service>ies>if>sap>sub-sla-mgmt config>service>ies>sub-if>grp-if>sap>sub-sla-mgmt
Description	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 config>subscriber-mgmt>sub-profile context. 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. The no form of the command removes the default SLA profile from the SAP configuration.
Parameters	<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 config>subscriber-mgmt>sub-profile context.

sub-ident-policy

Syntax	sub-ident-policy <i>sub-ident-policy-name</i>
Context	config>service>vpls>sap>sub-sla-mgmt config>service>ies>if>sap>sub-sla-mgmt config>service>ies>sub-if>grp-if>sap>sub-sla-mgmt
Description	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 config>subscriber-mgmt>sub-ident-policy context. 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. 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. 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. The no form of the command removes the default subscriber identification policy from the SAP configuration.
Default	no sub-ident-policy

Parameters	<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 config>subscriber-mgmt>sub-ident-policy context. subscriber <i>sub-ident-string</i> — 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 <i>sub-ident</i> 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 <i>sub-ident</i> 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 <i>sub-ident</i> 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

Syntax	[no] profiled-traffic-only
Context	config>service>vpls>sap>sub-sla-mgmt config>service>ies>if>sap>sub-sla-mgmt>single-sub config>service>ies>sub-if>grp-if>sap>sub-sla-mgmt
Description	This command enables profiled traffic only for this SAP. The no form of the command disables the command.

non-sub-traffic

Syntax	non-sub-traffic sub-profile <i>sub-profile-name</i> sla-profile <i>sla-profile-name</i> [subscriber <i>sub-ident-string</i>] no non-sub-traffic
Context	config>service>ies>if>sap>sub-sla-mgmt>single-sub config>service>ies>sub-if>grp-if>sap>sub-sla-mgmt>single-sub
Description	This command configures non-subscriber traffic profiles. The no form of the command removes removes the profiles and disables the feature.

Parameters	sub-profile <i>sub-profile-name</i> — Specifies 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 <i>sla-profile-name</i> — Specifies 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. subscriber <i>sub-ident-string</i> — Specifies 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 VPRN SAP arp-reply-agent to determine the proper handling of received ARP requests from subscribers. For SAPs with arp-reply-agent enabled with the optional <i>sub-ident</i> parameter, the static subscriber host's <i>sub-ident-string</i> 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 <i>sub-ident</i> 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 <i>sub-ident</i> 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

Syntax	[no] profiled-traffic-only
Context	config>service>ies>if>sap>sub-sla-mgmt>single-sub config>service>ies>sub-if>grp-if>sap>sub-sla-mgmt>single-sub
Description	This command enables profiled traffic only for this SAP. The no form of the command disables the command.

Multicast Commands

fast-leave

Syntax	[no] fast-leave
Context	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
Description	This command enables fast leave. When IGMP fast leave processing is enabled, the SR-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'). 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.
Default	no fast-leave

from-vpls

Syntax	from-vpls <i>vpls-id</i> no from-vpls
Context	config>service>vpls>sap>snooping>mvr config>service>vpls>sap>mld-snooping>mvr
Description	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.
Default	no from-vpls
Parameters	<i>vpls-id</i> — Specifies the MVR VPLS from which multicast channels should be copied into this SAP.

group

Syntax	[no] group <i>grp-address</i>
Context	config>service>vpls>sap>snooping>static config>service>vpls>spoke-sdp>snooping>static config>service>vpls>mesh-sdp>snooping>static
Description	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.
Default	none
Parameters	<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

Syntax	group-policy <i>policy-name</i> no group-policy
Context	config>service>vpls>snooping>mvr config>service>vpls>mld-snooping>mvr
Description	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 no form of the command removes the MVR policy association from the VPLS.
Default	no import policy is specified.
Parameters	<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 section Enabling IGMP group membership report filtering in the 7750 SR OS Router Configuration Guide.

igmp-snooping

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

mld-snooping

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

import

Syntax	import <i>policy-name</i> no import
Context	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
Description	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 no form of the command removes the policy association from the SAP or SDP.
Default	no import (No import policy is specified)
Parameters	<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

Syntax	last-member-query-interval <i>tenths-of-seconds</i> no last-member-query-interval
Context	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
Description	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.

Triple Play Service Configuration Commands

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

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

max-num-groups

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

mcac

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

policy

Syntax	policy <i>policy-name</i> no policy

Context	config>service>pw-template>igmp-snooping>mcac config>service>vpls>mesh-sdp>snooping>mcac
Description	This command configures the multicast CAC policy name.
Parameters	<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

Syntax	unconstrained-bw <i>bandwidth</i> mandatory-bw <i>mandatory-bw</i> no unconstrained-bw
Context	config>service>pw-template>igmp-snooping>mcac config>service>vpls>mesh-sdp>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	<i>bandwidth</i> — The bandwidth assigned for interface's MCAC policy traffic, in kilo-bits per second (kbps). Values 0 — 2147483647 mandatory-bw <i>mandatory-bw</i> — Specifies the bandwidth pre-reserved for all the mandatory channels on a given interface in kilo-bits per second (kbps). 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. If the value of <i>mandatory-bw</i> is equal to the value of <i>bandwidth</i> , 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 <i>mandatory-bw</i> should always be less than or equal to that of <i>bandwidth</i> . An attempt to set the value of <i>mandatory-bw</i> greater than that of <i>bandwidth</i> , will result in inconsistent value error. Values 0 — 2147483647

mrouter-port

Syntax	[no] mrouter-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 or SDP as an mrouter-port will have a double effect. First, all multicast traffic received on another SAP or SDP will be copied to this SAP or SDP. Second, 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.

This parameter can only be enabled on a SAP or spoke SDP, not on a mesh SDP.

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	<i>seconds</i> — The time interval, in seconds, that the router transmits general host-query messages.
Values	2 — 1024

query-src-ip

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

query-response-interval

Syntax	query-response-interval <i>seconds</i>
Context	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
Description	If the send-queries 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.
Default	10
Parameters	<i>seconds</i> — Specifies the length of time to wait to receive a response to the host-query message from the host.
Values	1 — 1023

report-src-ip

Syntax	report-src-ip <i>ip-address</i> no report-src-ip
Context	config>service>vpls>igmp-snooping
Description	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

Triple Play Service Configuration Commands

interoperability with some multicast routers, the source IP address of IGMP group reports can be configured using this command.

Default 0.0.0.0

Parameters *ip-address* — The source IP source address in transmitted IGMP reports.

robust-count

Syntax **robust-count** *robust-count*
no robust-count

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>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 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.

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>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

Description 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 **send-queries** command parameter can not be set.

Default no send-queries

source

[no] source *ip-address*

Context	config>service>vpls>sap>snooping>static>group config>service>vpls>spoke-sdp>snooping>static>group config>service>vpls>mesh-sdp>snooping>static>group
Description	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.
Default	none
Parameters	<i>ip-address</i> — Specifies the IPv4 unicast address.

starg

Syntax	[no] starg
Context	config>service>vpls>sap>snooping>static>group config>service>vpls>spoke-sdp>snooping>static>group config>service>vpls>mesh-sdp>snooping>static>group
Description	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 no form of the command to remove the starg entry from the configuration.
Default	no starg

static

Syntax	static
Context	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
Description	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.
Default	none

to-sap

Syntax	to-sap <i>sap-id</i> no to-sap
Context	config>service>vpls>sap>snooping>mvr
Description	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. This command configures the SAP to which the multicast data needs to be copied.
Default	no to-sap
Parameters	<i>sap-id</i> — Specifies the SAP to which multicast channels should be copied. See Common Service Commands on page 1020 for <i>sap-id</i> command syntax.

DHCP and Anti-Spoofing Commands

anti-spoof

Syntax	anti-spoof {ip mac ip-mac} no anti-spoof
Context	config>service>vpls>sap config>service>ies>sap
Description	This command enables anti-spoof filtering and optionally changes the anti-spoof matching type for the SAP. 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 (ip, mac, ip-mac) defines the anti-spoof filter type enforced by the SAP when anti-spoof filtering is enabled. The no form of the command disables anti-spoof filtering on the SAP.
Default	no anti-spoof
Parameters	ip — 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 ip command will fail. mac — 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 anti-spoof mac command will fail. ip-mac — 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 ip-mac command will fail.

arp-host

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

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	<i>max-num-hosts</i> — specifies the maximum number of ARP hosts allowed on this SAP.
Values	1 — 32767

min-auth-interval

Syntax	min-auth-interval <i>min-auth-interval</i> 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	<i>min-auth-interval</i> — 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.

calling-station-id

Syntax	[no] calling-station-id {mac remote-id sap-id sap-string}
Context	<pre>config>service>ies>if>sap config>service>ies>sub-if>grp-if config>service>vpls>sap config>service>vprn>interface config>service>vprn>sub-if>grp-if config>subscr-mgmt>auth-plcy>include-radius-attribute</pre>
Description	This command enables the inclusion of the calling-station-id attribute in RADIUS authentication requests and RADIUS accounting messages.
Default	no calling-station-id
Parameters	mac — Specifies that the mac-address will be sent. remote-id — Specifies that the remote-id will be sent. sap-id — Specifies that the sap-id will be sent. sap-string — 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.

host

Syntax	host {[ip <i>ip-address</i>] [mac <i>mac-address</i>]} [subscriber-sap-id subscriber <i>sub-ident-string</i>] [sub-profile <i>sub-profile-name</i>] [sla-profile <i>sla-profile-name</i>] [ancp-string <i>ancp-string</i>] [inter-dest-id <i><intermediate-destination-id></i>] no host {[ip <i>ip-address</i>] [mac <i>ieee-address</i>]} no host all
Context	config>service>vpls>sap config>service>ies>if>sap config>service>vprn>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 <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.
Default	none
Parameters	ip <i>ip-address</i> — Specify this optional parameter when defining a static host. The IP address must be specified for anti-spoof ip and anti-spoof ip-mac. Only one static host may be configured on the SAP with a given IP address. mac <i>mac-address</i> — Specify this optional parameter when defining a static host. The MAC address must be specified for anti-spoof mac and anti-spoof ip-mac. 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 or VPRN) 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 <i>mac-address</i>. 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.

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

Syntax	egress
Context	config>service>vpls>sap config>service>vpls>mesh-sdp config>service>ies>if>sap
Description	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

Syntax	ingress
Context	config>service>vpls>sap>egress config>service>vpls>mesh-sdp config>service>vpls>spoke-sdp config>service>ies>if>sap
Description	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

Syntax	filter ip <i>ip-filter-id</i> filter ipv6 <i>ipv6-filter-id</i> filter mac <i>mac-filter-id</i> no filter [ip <i>ip-filter-id</i>] [mac <i>mac-filter-id</i>] [ipv6 <i>ipv6-filter-id</i>]
Context	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
Description	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**.

Special Cases	VPLS — Both MAC and IP filters are supported on a VPLS service SAP.
Parameters	ip <i>ip-filter-id</i> — Specifies IP filter policy. The filter ID must already exist within the created IP filters. Values 1 — 65535 ipv6 <i>ipv6-filter-id</i> — Specifies the IPv6 filter policy. The filter ID must already exist within the created IPv6 filters. Values 1 — 65535 mac <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. Values 1 — 65535

agg-rate-limit

Syntax	agg-rate-limit <i>agg-rate</i> [queue-frame-based-accounting] no agg-rate-limit
Context	config>service>ies>if>sap>egress config>service>vpls>sap>egress
Description	This command defines a maximum total rate for all egress queues on a service SAP or multi-service site. The agg-rate-limit command is mutually exclusive with the egress scheduler policy. When an egress scheduler policy is defined, the agg-rate-limit command will fail. If the agg-rate-limit command is specified, at attempt to bind a scheduler-policy to the SAP or multi-service site will fail. A multi-service site must have a port scope defined that ensures all queues associated with the site are on the same port or channel. If the scope is not set to a port, the agg-rate-limit command will fail. Once an agg-rate-limit has been assigned to a multi-service site, the scope cannot be changed to card level. A port scheduler policy must be applied on the egress port or channel the SAP or multi-service site are bound to in order for the defined 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.

If a port scheduler is not defined on the egress port, the queues are allowed to operate based on their own bandwidth parameters.

The **no** form of the command removes the aggregate rate limit from the SAP or multi-service site.

Parameters *agg-rate* — defines the rate, in kilobits-per-second, that the maximum aggregate rate that the queues on the SAP or multi-service-site can operate.

Values 1 — 40000000, max

queue-frame-based-accounting — This keyword enables frame based accounting on all queues associated with the SAP or Multi-Service Site. If frame based accounting is required when an aggregate limit is not necessary, the max keyword should precede the queue-frame-based-accounting keyword. If frame based accounting must be disabled, execute agg-rate-limit without the queue-frame-based-accounting keyword present.

Default Frame based accounting is disabled by default

filter

Syntax

```
filter ip ip-filter-id
filter ipv6 ipv6-filter-id
no filter [ip ip-filter-id] [ipv6 ipv6-filter-id]
```

Context

```
config>service>ies>if>sap>egress
config>service>ies>if>sap>ingress
config>service>ies>sub-if>grp-if>sap>egress
config>service>ies>sub-if>grp-if>sap>ingress
config>service>vprn>sub-if>grp-if>sap>egress
config>service>vprn>sub-if>grp-if>sap>ingress
```

Description

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.

The **filter** 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 **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 the 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. 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**.

Special Cases

IES — Only IP filters are supported on an IES IP interface, and the filters only apply to routed traffic.

Parameters	ip <i>ip-filter-id</i> — Specifies IP filter policy. The filter ID must already exist within the created IP filters.
Values	1 — 65535
	ipv6 <i>ipv6-filter-id</i> — Specifies the IPv6 filter policy. The filter ID must already exist within the created IPv6 filters.
Values	1 — 65535

qinq-mark-top-only

Syntax	[no] qinq-mark-top-only
Context	config>service>vprn>if>sap>egress config>service>vprn>sub-if>grp-if>sap>egress
Description	When enabled (the encapsulation type of the access port where this SAP is defined as qinq), the qinq-mark-top-only command specifies which P-bits to mark during packet egress. When disabled, both set of P-bits are marked. When the enabled, only the P-bits in the top Q-tag are marked.
Default	no qinq-mark-top-only

multicast-group

Syntax	multicast-group <i>group-name</i> no multicast-group
Context	config>service>vpls>sap>egress
Description	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: • If an egress-filter is specified on the egress multicast group, the SAP must have the same egress filter applied. • If an egress-filter is not defined on the egress multicast group, the SAP cannot have an egress filter applied. • 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. • 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. • The access port the SAP is created on cannot currently be an egress mirror source. Once a SAP is a member of an egress multicast group, the following rules apply: • The egress filter defined on the SAP cannot be removed or modified. Egress filtering is managed at the egress multicast group for member SAPs.

- 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.
- 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.

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.

Default	no multicast-group
Parameters	<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.
Values	Any valid egress multicast group name.
Default	None, an egress multicast group name must be explicitly specified.

qos

Syntax	qos <i>policy-id</i> no qos
Context	config>service>vpls>sap>egress config>service>vpls>sap>ingress config>service>ies>if>sap>egress config>service>ies>if>sap>ingress config>service>vprn>if>sap>egress config>service>vprn>sub-if>grp-if>sap>egress
Description	This command associates a Quality of Service (QoS) policy with an ingress or egress Service Access Point (SAP) or IP interface. 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.

The **qos** command is used to associate both ingress and egress QoS policies. The **qos** 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.

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.

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.

The **no** form of this command removes the QoS policy association from the SAP or IP interface, and the QoS policy reverts to the default.

policy-id — The ingress/egress policy ID to associate with SAP or IP interface on ingress/egress. The policy ID must already exist.

Values 1 — 65535

shared-queuing — 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.

queue-override

Syntax	[no] queue-override
Context	config>service>ies>if>sap>egress config>service>ies>sub-if>grp-if>sap>egress
Description	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

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

adaptation-rule

Syntax	adaptation-rule [pir { max min closest }] [cir { max min closest }] no adaptation-rule
Context	config>service>ies>if>sap>egress>queue-override>queue config>service>ies>sub-if>grp-if>sap>egress>queue-override>queue
Description	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. The no 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 adaptation-rule is removed, the default constraints for rate and cir apply.
Default	no adaptation-rule
Parameters	pir — The pir parameter defines the constraints enforced when adapting the PIR rate defined within the queue queue-id rate command. The pir parameter requires a qualifier that defines the constraint used when deriving the operational PIR for the queue. When the rate command is not specified, the default applies. max — The max (maximum) option is mutually exclusive with the min and closest options. When max is defined, the operational PIR for the queue will be equal to or less than the administrative rate specified using the rate command. min — The min (minimum) option is mutually exclusive with the max and closest options. When min is defined, the operational PIR for the queue will be equal to or greater than the administrative rate specified using the rate command. 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

Syntax	avg-frame-overhead <i>percent</i> no avg-frame-overhead
Context	config>service>ies>if>sap>egress>queue-override config>service>ies>sub-if>grp-if>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 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 10000×0.1 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 50×20 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 500×1.1 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 7500×1.1 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.

Default	0
Parameters	<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.
Values	0 — 100

cbs

Syntax	cbs <i>size-in-kbytes</i> no cbs
Context	config>service>ies>if>sap>egress>queue-override>queue config>service>ies>sub-if>grp-if>sap>egress>queue-override>queue
Description	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 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.

Default	no cbs
Parameters	<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).
Values	0 — 131072 or default

high-prio-only

Syntax	high-prio-only <i>percent</i> no high-prio-only
Context	config>service>ies>if>sap>egress>queue-override>queue config>service>ies>sub-if>grp-if>sap>egress>queue-override>queue
Description	This command can be used to override specific attributes of the specified queue's high-prio-only parameters. The high-prio-only command configures the percentage of buffer space for the queue, used exclusively by high priority packets. 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 restores the default high priority reserved size.
Parameters	<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. Values 0 — 100 default

mbs

Syntax	mbs { <i>size-in-kbytes</i> default } no mbs
Context	config>service>ies>if>sap>egress>queue-override>queue config>service>ies>sub-if>grp-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	<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.
Values	0 — 131072 or default

rate

Syntax	rate <i>pir-rate</i> [<i>cir cir-rate</i>] no rate
Context	config>service>ies>if>sap>egress>queue-override>queue config>service>ies>sub-if>grp-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 <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. 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 <i>queue-id</i>. The no 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 (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	<i>pir-rate</i> — 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 *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. The **sum** keyword specifies that the CIR be used as the summed CIR values of the children schedulers or queues.

Values 0 — 100000000, max, sum

Default 0

scheduler-policy

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

Context config>service>vpls>sap>ingress
config>service>vpls>sap>egress
config>service>ies>if>sap>egress
config>service>ies>if>sap>ingress
config>service>vprn>if>sap>egress
config>service>vprn>sub-if>grp-if>sap>egress
config>service>vprn>if>sap>ingress
config>service>vprn>sub-if>grp-if>sap>ingress

Description 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 **config>qos>scheduler-policy** *scheduler-policy-name* 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

Syntax	[no] block-on-mesh-failure
Context	config>service>vpls>spoke-sdp
Description	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).
Default	disabled

max-nbr-mac-addr

Syntax	max-nbr-mac-addr <i>table-size</i> no max-nbr-mac-addr
Context	config>service>vpls>sap config>service>vpls>spoke-sdp
Description	This command specifies the maximum number of FDB entries for both learned and static MAC addresses for this SAP or spoke SDP. When the configured limit has been reached, and discard-unknown-source has been enabled for this SAP or spoke SDP (see the discard-unknown-source command), packets with unknown source MAC addresses will be discarded. The no form of the command restores the global MAC learning limitations for the SAP or spoke SDP.
Default	no max-nbr-mac-addr
Parameters	<i>table-size</i> — The maximum number of MAC entries in the FDB from this SAP. Values 1 — 511999

multi-service-site

Syntax	multi-service-site <i>customer-site-name</i> no multi-service-site
Context	config>service>vpls>sap config>service>ies>sap config>service>ies>subscriber-interface>grp-if>sap config>service>vprn>if>sap config>service>vprn>sub-if>grp-if>sap
Description	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.

This command is mutually exclusive with the SAP ingress and egress **scheduler-policy** commands. If a **scheduler-policy** has been applied to either the ingress or egress nodes on the SAP, the **multi-service-site** command will fail without executing. The locally applied scheduler policies must be removed prior to executing the **multi-service-site** command.

The **no** 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.

Default None

customer-site-name — 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.

Values Any valid customer-site-name created within the context of the customer-id.

static-host

Syntax **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*

Context config>service>ies>if>sap
config>service>vpls>sap
config>service>ies>sub-if>grp-if>sap
config>service>vprn>if>sap
config>service>vprn>sub-if>grp-if>sap

Description This command configures a static host on this SAP.

Syntax **ip** *ip-address* — Specifies the IPv4 unicast address.

mac *ieee-address* — Specify this optional parameter when defining a static host. Every static host definition must have at least one address defined, IP or MAC.

force — Specifies the forced removal of the static host addresses.

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

Syntax **ancp-string** *ancp-string*
no ancp-string

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

Triple Play Service Configuration Commands

```
config>service>ies>sub-if>grp-if>sap>static-host
config>service>vprn>if>sap
config>service>vprn>sub-if>grp-if>sap
```

- Description** This command specifies the ANCP string associated to this SAP host.
- Parameters** *ancp-string* — Specifies the ANCP string up to 63 characters in length.

app-profile

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

inter-dest-id

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

managed-routes

- Syntax** **managed-routes**
- Context** config>service>vprn>sub-if>grp-if>sap>static-host>managed-routes
- Description** This command configures managed routes.

route

Syntax	route <i>{ip-prefix/length ip-prefix netmask}</i> [create] no route <i>{ip-prefix/length ip-prefix netmask}</i>
Context	config>service>vprn>sub-if>grp-if>sap>static-host>managed-routes
Description	This command assigns managed-route to a given subscriber-host. As a consequence, a static-route pointing subscriber-host ip address as a next hop will be installed in FIB. Up to 16 managed routes per subscriber-host can be configured. The no form of the command removes the respective route. Per default, there are no managed-routes configured.

sla-profile

Syntax	sla-profile <i>sla-profile-name</i> no sla-profile
Context	config>service>ies>if>sap>static-host config>service>vpls>sap>static-host config>service>ies>sub-if>grp-if>sap>static-host config>service>vprn>if>sap>static-host config>service>vprn>sub-if>grp-if>sap>static-host
Description	This command specifies 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.
Parameters	<i>sla-profile-name</i> — Specifies the SLA profile name.

sub-profile

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

subscriber

Syntax	subscriber <i>sub-ident</i> no subscriber
Context	config>service>ies>if>sap>static-host config>service>vpls>sap>static-host config>service>ies>sub-if>grp-if>sap>static-host config>service>vprn>if>sap>static-host config>service>vprn>sub-if>grp-if>sap>static-host
Description	This command specifies an existing subscriber identification profile to be associated with the static subscriber host.
Parameters	<i>sub-ident</i> — 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 config>service>ies>sub-if>grp-if>sap>static-host config>service>vprn>if>sap>static-host config>service>vprn>sub-if>grp-if>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 <i>ieee-mac-address</i>
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 7750 SR 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 7750 SR router. If two *sdp-id* bindings terminate on the same 7750 SR, 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

Syntax	spoke-sdp <i>sdp-id[:vc-id]</i> [vc-type { ether vlan }] [split-horizon-group <i>group-name</i>] no spoke-sdp <i>sdp-id[:vc-id]</i>
Context	config>service>vpls
Description	This command binds a service to an existing Service Distribution Point (SDP). 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. 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 an SDP with a VPLS service. If the sdp sdp-id 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. SDPs must be explicitly associated and bound to a service. If an SDP is not bound to a service, no far-end 7750 SR 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 <i>sdp-id</i> is bound to a service.
Special Cases	VPLS — 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.
Parameters	<i>sdp-id</i> — The SDP identifier. Values 1 — 17407 <i>vc-id</i> — The virtual circuit identifier. 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.

Values ether, vlan

ether — Defines the VC type as Ethernet. The **ether**, **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 **ether**, **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.

spoke-sdp

Syntax	spoke-sdp <i>sdp-id[:vc-id]</i> no spoke-sdp <i>sdp-id[:vc-id]</i>
Context	config>service>ies>interface
Description	This command binds a service to an existing Service Distribution Point (SDP). 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. 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 an SDP with an IES service. If the sdp sdp-id 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. SDPs must be explicitly associated and bound to a service. If an SDP is not bound to a service, no far-end 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. The spoke SDP must be shut down first before it can be deleted from the configuration.
Default	No <i>sdp-id</i> is bound to a service.
Special Cases	IES — At most, only one <i>sdp-id</i> can be bound to an IES service.

Parameters	<i>sdp-id</i> — The SDP identifier. Allowed values are integers in the range of 1 and 17407 for existing SDPs. <i>vc-id</i> — The virtual circuit identifier. Values 1 — 4294967295

egress

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

ingress

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

vc-label

Syntax	[no] vc-label egress-vc-label
Context	<pre>config>service>vpls>mesh-sdp>egress config>service>vpls>spoke-sdp>egress config>service>ies>if>spoke-sdp>egress</pre>
Description	This command configures the egress VC label.
Parameters	<i>vc-label</i> — A VC egress value that indicates a specific connection. Values 16 — 1048575

vc-label

Syntax	[no] vc-label ingress-vc-label
Context	<pre>config>service>vpls>mesh-sdp>ingress config>service>vpls>spoke-sdp>ingress config>service>ies>if>spoke-sdp>ingress</pre>

Triple Play Service Configuration Commands

Description	This command configures the ingress VC label.
Parameters	<i>vc-label</i> — A VC ingress value that indicates a specific connection.
Values	2048 — 18431

vlan-vc-tag

Syntax	vlan-vc-tag <i>0..4094</i> no vlan-vc-tag [<i>0..4094</i>]
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	<i>0..4094</i> — Specifies a valid VLAN identifier to bind an 802.1Q VLAN tag ID.

avg-frame-overhead

Syntax	avg-frame-overhead <i>percent</i> 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 10000×0.1 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 50×20 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 500×1.1 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 7500×1.1 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.

Default	0
Parameters	<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.
Values	0 — 100

queue-override

Syntax	[no] queue-override
Context	config>service>vpls>sap>egress config>service>vpls>sap>ingress
Description	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

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

adaptation-rule

Syntax	adaptation-rule [pir { max min closest }] [cir { max min closest }] no adaptation-rule
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 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. The no 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 adaptation-rule is removed, the default constraints for rate and cir apply.
Default	no adaptation-rule
Parameters	pir — The pir parameter defines the constraints enforced when adapting the PIR rate defined within the queue queue-id rate command. The pir parameter requires a qualifier that defines the constraint used when deriving the operational PIR for the queue. When the rate command is not specified, the default applies. max — The max (maximum) option is mutually exclusive with the min and closest options. When max is defined, the operational PIR for the queue will be equal to or less than the administrative rate specified using the rate command. min — The min (minimum) option is mutually exclusive with the max and closest options. When min is defined, the operational PIR for the queue will be equal to or greater than the administrative rate specified using the rate command. 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.

cbs

Syntax	cbs <i>size-in-kbytes</i> no cbs
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 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 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.

Default	no cbs
Parameters	<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).
Values	0 — 131072 or default

high-prio-only

Syntax	high-prio-only <i>percent</i> no high-prio-only
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 high-prio-only parameters. The high-prio-only command configures the percentage of buffer space for the queue, used exclusively by high priority packets. 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 restores the default high priority reserved size.
Parameters	<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. Values 0 — 100 default

mbs

Syntax	mbs { <i>size-in-kbytes</i> default } no mbs
Context	config>service>vpls>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	<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. Values 0 — 131072 or default

mbs

Syntax	mbs { <i>size-in-kbytes</i> default } no mbs
Context	config>service>vpls>sap>ingress>queue-override>queue
Description	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. 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. If the CBS value is larger than the MBS value, an error will occur, preventing the MBS change. 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-

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

Syntax	rate <i>pir-rate</i> [<i>cir cir-rate</i>] no rate
Context	config>service>vpls>sap>egress>sched-override>scheduler
Description	This command can be used to override specific attributes of the specified scheduler rate. The rate 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. 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.

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

Syntax	match-qinq-dot1p {top bottom} no match-qinq-dot1p
Context	config>service>vpls>sap>ingress config>service>ies>if>sap>ingress config>service>vprn>sub-if>grp-if>sap>ingress
Description	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

Syntax [no] discard-unknown-source

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

Description 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 [max-nbr-mac-addr on page 130](#)) 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**

ima — Specifies Inverse Multiplexing over ATM. An IMA Group is a collection of physical links bundled together and assigned to an ATM Port.

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 TDM	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 TDM	BCP-Dot1q	0 - 4094	The SAP is identified by the 802.1Q tag on the channel.
SONET/SDH TDM	Frame Relay	16 — 991	The SAP is identified by the data link connection identifier (DLCI).
SONET/SDH ATM	ATM	vpi (NNI) 0 — 4095 vpi (UNI) 0 — 255 vci 1, 2, 5 — 65535	The SAP is identified by the PVC identifier (vpi/vci).

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

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

l2pt-termination

Syntax	[no] l2pt-termination
Context	config>service>vpls>sap config>service>vpls>spoke-sdp
Description	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.
Default	no l2pt-termination

def-mesh-vc-id

Syntax	[no] def-mesh-vc-id <i>vc-id</i>
Context	config>service>vpls
Description	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 SRs 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 no form of this command disables the VC-ID.
Default	none

Values 1 — 4294967295

mac-move

Syntax	[no] mac-move
Context	config>service>vpls
Description	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. When enabled in a VPLS, mac-move 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 shutdown/no shutdown 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 config>service>vpls>sap>limit-mac-move or config>service>vpls>spoke-sdp>limit-mac-move contexts, see page 93. This means that when the re-learn rate has exceeded the limit, another (blockable) SAP will be disabled instead. The mac-move 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. 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. 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. The no form of this command disables MAC move.
Default	not enabled

mac-protect

Syntax	mac-protect
Context	config>service>vpls
Description	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.
Default	disabled

mac

Syntax	[no] mac <i>ieee-address</i>
Context	config>service>vpls>mac-protect
Description	This command specifies the 48-bit IEEE 802.3 MAC address.
Parameters	<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

Syntax	mac-subnet-length <i>subnet-length</i> no mac-subnet-length
Context	config>service>vpls
Description	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. The no form of this command switches back to full MAC lookup.
Parameters	<i>subnet-length</i> — Specifies the number of bits to be considered when performing MAC learning or MAC switching.
Values	24 — 48

move-frequency

Syntax	move-frequency <i>frequency</i> no move-frequency
Context	config>service>vpls>mac-move
Description	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. 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. The no form of the command reverts to the default value.
Default	2 (when mac-move is enabled)
Parameters	<i>frequency</i> — Specifies the rate, in 5-second intervals for the maximum number of relearns.

Values 1 — 100

retry-timeout

Syntax	retry-timeout <i>timeout</i> no retry-timeout
Context	config>service>vpls>mac-move
Description	This object indicates the time in seconds to wait before a SAP that has been disabled after exceeding the maximum relearn rate is reenabled. 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. The no form of the command reverts to the default value.
Default	10 (when mac-move is enabled)
Parameters	<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.
	Values 0 — 120

mfib-table-high-wmark

Syntax	[no] mfib-table-high-wmark <i>high-water-mark</i>
Context	config>service>vpls
Description	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.
Parameters	<i>high-water-mark</i> — Specifies the multicast FIB high watermark as a percentage.
	Values 1 — 100
	Default 95%

mfib-table-low-wmark

Syntax	[no] mfib-table-low-wmark <i>low-water-mark</i>
Context	config>service>vpls
Description	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.
Parameters	<i>low-water-mark</i> — Specifies the multicast FIB low watermark as a percentage.

Values 1 — 100

Default 90%

mfib-table-size

Syntax	mfib-table-size <i>size</i> no mfib-table-size
Context	config>service>vpls
Description	This command specifies the maximum number of (s,g) entries in the multicast forwarding database (MFIB) for this VPLS instance. 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. 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. The no form of this command removes the configured maximum MFIB table size.
Default	none
Parameters	<i>size</i> — The maximum number of (s,g) entries allowed in the Multicast FIB. Values 1 — 16383

send-flush-on-failure

Syntax	[no] send-flush-on-failure
Context	config>service>vpls
Description	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. This feature cannot be enabled on management VPLS.
Default	no send-flush-on-failure

restrict-protected-src

Syntax	[no] restrict-protected-src
Context	config>service>vpls>split-horizon-group config>service>vpls>sap

Triple Play Service Configuration Commands

Description	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.
Default	no restrict-protected-src

restrict-unprotected-dst

Syntax	[no] restrict-unprotected-dst
Context	config>service>vpls>split-horizon-group config>service>vpls>sap
Description	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.
Default	no restrict-unprotected-dst

DHCP Commands

dhcp

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

gi-address

Syntax	gi-address <i>ip-address</i> [<i>src-ip-addr</i>] no gi-address
Context	<pre>config>service>ies>if>dhcp config>service>vprn>interface>dhcp config>service>ies>subscriber-interface>grp-if>dhcp config>service>ies>subscriber-interface>dhcp</pre>
Description	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. 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.
Default	no gi-address
Parameters	<i>ip-address</i> — Specifies the host IP address to be used for DHCP relay packets. <i>src-ip-address</i> — Specifies that this GI address is to be the source IP address for DHCP relay packets.

lease-populate

Syntax	lease-populate [<i>number-of-entries</i>] no lease-populate
Context	config>service>vpls>sap>dhcp

```
config>service>ies>interface>dhcp
config>service>vpn>interface>dhcp
config>service>ies>subscriber-interface>grp-if>dhcp
```

Description	This command enables and disables 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 SAP. The optional number-of-entries parameter is used to define the number lease state table entries allowed for this SAP 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 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. • Generate dynamic ARP replies if arp-reply-agent is enabled.
Default	no lease-populate

match-circuit-id

Syntax	[no] match-circuit-id
Context	config>service>ies>sub-if>grp-if>dhcp
Description	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. When we 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. When the match-circuit-id command is enabled we use this as part of the key to guarantee correctness in our lookup. This is really only needed when we are dealing with an IP aware DSLAM that proxies the client HW mac address.
Default	no match-circuit-id

option

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

action

Syntax	action {replace drop keep} no action
Context	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
Description	This command configures the Relay Agent Information Option (Option 82) processing. The no form of this command returns the system to the default value.
Default	The default is to keep the existing information intact.
Parameters	replace — 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). drop — The DHCP packet is dropped if an Option 82 field is present, and a counter is incremented. keep — 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

Syntax	circuit-id [ascii-tuple vlan-ascii-tuple] no circuit-id
Context	config>service>vpls>sap>dhcp>option
Description	When enabled, the router sends an ASCII-encoded tuple in the circuit-id 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 **circuit-id** sub-option of the DHCP packet will be left empty.

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

Default	circuit-id
Parameters	ascii-tuple — Specifies that the ASCII-encoded concatenated tuple consisting of the access-node-identifier, service-id, and interface-name is used. 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 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

Syntax	circuit-id [ascii-tuple ifindex sap-id vlan-ascii-tuple] no circuit-id
Context	config>service>ies>if>dhcp>option config>service>vprn>if>dhcp>option config>service>vprn>subscriber-interface>grp-if>dhcp>option
Description	When enabled, the router sends an ASCII-encoded tuple in the circuit-id 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 circuit-id sub-option of the DHCP packet will be left empty. The no form of this command returns the system to the default.
Default	circuit-id ascii-tuple — Specifies that the ASCII-encoded concatenated tuple will be used which consists of the access-node-identifier, service-id, and interface-name, separated by “ ”. ifindex — Specifies that the interface index will be used. (The If Index of a router interface can be displayed using the command <code>show>router>interface>detail</code>) sap-id — Specifies that the SAP identifier will be used. 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

Syntax	remote-id [mac string <i>string</i>] no remote-id
Context	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
```

```
config>subscr-mgmt>msap-policy>vpls-only>dhcp>option
```

Description This command specifies what information goes into the remote-id sub-option in the DHCP Relay packet.

If disabled, the **remote-id** sub-option of the DHCP packet will be left empty.

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

Default remote-id

Parameters **mac** — This keyword specifies the MAC address of the remote end is encoded in the sub-option.
string string — Specifies the remote-id.

snoop

Syntax [no] snoop

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

Description 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.

Use the **no** form of the command to disable DHCP snooping on the specified VPLS SAP or SDP binding.

Default no snoop

server

Syntax **server server1** [server2...(up to 8 max)]

Context config>service>ies>if>dhcp
 config>service>vprn>if>dhcp
 config>service>ies>subscriber-interface>grp-if>dhcp

Description 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.

There can be a maximum of 8 DHCP servers configured.

Default no server

Parameters *server* — Specify the DHCP server IP address.

trusted

Syntax	[no] trusted
Context	config>service>ies>if>dhcp config>service>vpn>if>dhcp config>service>ies>subscriber-interface>grp-if>dhcp
Description	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. 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 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

Syntax	egress-multicast-group <i>egress-multicast-group-name</i> no egress-multicast-group <i>group-name</i>
Context	config>service
Description	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. Only SAPs defined on Ethernet access ports are allowed into an egress-multicast-group. 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. 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. 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. 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. 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.

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.

Parameters	<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.
Default	None, each egress multicast group must be explicitly configured.
Values	Up to 32 egress multicast groups may be created on the system.

description

Syntax	description <i>description-string</i> no description
Context	config>service>egress-multicast-group
Description	This command defines an ASCII string associated with egress-multicast-group-name. The no form of the command removes an existing description string from egress-multicast-group.
Default	none
Parameters	<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.
Values	An ASCII string up to 80 characters in length.

dest-chain-limit

Syntax	dest-chain-limit <i>destinations per pass</i> no dest-chain-limit
Context	config>service>egress-multicast-group
Description	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. The no form of the command restores the default value of 10 to the dest-chain-limit parameter for the egress-multicast-group.
Default	no dest-chain-limit
Parameters	<i>destinations per pass</i> — This parameter must be specified when executing the dest-chain-limit 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. 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. 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. 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. 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 dest-chain-limit and is provided for debugging purposes only. Setting the value to 1 is persistent between reboots of the system.
Values	0 — 255
Default	10

sap-common-requirements

Syntax	sap-common-requirements
Context	config>service>egress-multicast-group
Description	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

Syntax	egress-filter [ip <i>ip-filter-id</i>] egress-filter [ipv6 <i>ipv6-filter-id</i>] egress-filter [mac <i>mac-filter-id</i>] no egress-filter [ip <i>ip-filter-id</i>] [ipv6 <i>ipv6-filter-id</i>] [mac <i>mac-filter-id</i>]
Context	config>service>egress-multicast-group>sap-common-requirements
Description	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. Once the SAP is made a member of the egress-multicast-group, the SAP's egress filter cannot be changed on the SAP. Changing the egress-filter parameters within the sap-common-requirements 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. 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. The no form of the command removes the egress-filter from each member SAP. The no egress-filter command specifies that an egress filter (IP or MAC)(IP, IPv6 or MAC) is not applied to a new member SAP within the egress-multicast-group.
Default	no filter. The egress filter ID must be defined with the associated ip or mac 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.
Parameters	ip <i>ip-filter-id</i> — Specifies IP filter policy. The filter ID must already exist within the created IP filters. Values 1 — 65535 ipv6 <i>ipv6-filter-id</i> — Specifies the IPv6 filter policy. The filter ID must already exist within the created IPv6 filters. Values 1 — 65535

mac *mac-filter-id* — 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.

Values 1 — 65535

encap-type

Syntax	encap-type {dot1q null} no encap-type
Context	config>service>egress-multicast-group>sap-common-requirements
Description	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 config>port>ethernet>access>encap-type 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. If at least one SAP is currently a member of the efficient-multicast-group, the encap-type cannot be changed within the sap-common-requirements node. If the efficient-multicast-group does not contain any member SAPs, the encap-type may be changed at anytime. 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. The no 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.
Default	dot1q — For an egress-multicast-group. null — If member SAPs are on a null encapsulated access port.
Parameters	null — The null keyword is mutually exclusive with the dot1q 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. dot1q — The dot1q keyword is mutually exclusive with the null 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.

dot1q-etype

Syntax	dot1q-etype [0x0600..0xffff] no dot1q-etype
Context	config>service>egress-multicast-group>sap-common-requirements
Description	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 config>port>ethernet>access>dot1q-

etype 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.

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.

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.

Default The default dot1q-etype is 0x8100 for an egress-multicast-group.

Parameters *ethertype* — 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.

Values [0x0600 — 0xffff]: [1536 — 65535] in decimal or hex

Default 0x8100

Interface Commands

interface

Syntax	interface <i>ip-int-name</i> no interface <i>ip-int-name</i>
Context	config>service>ies
Description	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. The interface command, under the context of services, is used to create and maintain IP routing interfaces within IES service IDs. The interface 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. Interface names are case sensitive and must be unique within the group of defined IP interfaces defined for config router interface and config service ies interface (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. The available IP address space for local subnets and routes is controlled with the config router service-prefix command. The service-prefix 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 config router and config service domains. 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. The no form of this command removes the IP interface and all the associated configuration. The interface must be administratively shutdown before issuing the no interface command. For IES services, the IP interface must be shutdown before the SAP on that interface may be removed. IES services do not have the shutdown command in the SAP CLI context. IES service SAPs rely on the interface status to enable and disable them.
Parameters	<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 config router interface and config service ies interface 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 <i>ip-int-name</i> already exists within the service ID, the context will be changed to maintain that IP interface. If <i>ip-int-name</i> 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 <i>ip-address/mask-length</i> [broadcast { all-ones host-ones }] no address
Context	config>service>ies>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 7750 SR. 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 “*P*” 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

Syntax	[no] allow-directed-broadcast
Context	config>service>ies>interface config>service>vprn>interface
Description	This command enables the forwarding of directed broadcasts out of the IP interface. A directed broadcast is a packet received on a local router interface destined for the subnet broadcast address on another IP interface. The allow-directed-broadcasts command on an IP interface enables or disables the transmission of packets destined to the subnet broadcast address of the egress IP interface. 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. When disabled, directed broadcast packets discarded at this egress IP interface will be counted in the normal discard counters for the egress SAP. By default, directed broadcasts are not allowed and will be discarded at this egress IP interface. The no form of this command disables the forwarding of directed broadcasts out of the IP interface.
Default	no allow-directed-broadcasts — Directed broadcasts are dropped

loopback

Syntax	[no] loopback
Context	config>service>ies>interface
Description	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. 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.
Default	None

mac

Syntax	mac ieee-address no mac
Context	config>service>ies>interface config>service>ies>sub-if>grp-if
Description	This command assigns a specific MAC address to an IES IP interface. The no form of the command returns the MAC address of the IP interface to the default value.

Default	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	<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.

sap

Syntax	[no] sap sap-id [create]
Context	config>service>ies>interface config>service>ies>sub-if>grp-if config>service>vprn>interface config>service>vprn>sub-if>grp-if
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 7750. 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. 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 — An IES SAP can be defined with Ethernet ports, SONET/SDH or TDM channels. A SAP is defined within the context of an IP routed interface. Each IP interface is limited to a single SAP definition. Group interfaces allow more than one SAP. Attempts to create a second SAP on an IP interface will fail and generate an error; the original SAP will not be affected.

secondary

Syntax	secondary { <i>ip-address/mask</i> <i>ip-address netmask</i> } [broadcast all-ones host-ones] [igp-inhibit] no secondary <i>ip-address</i>
Context	config>service>ies>interface This command assigns a secondary IP address/IP subnet/broadcast address format to the interface.
Default	none
Parameters	<i>ip-address</i> — The IP address of the IP interface. The <i>ip-address</i> 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). <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. <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. 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: <i>host-ones</i>) 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 <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. 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. igp-inhibit — The optional igp-inhibit 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

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

Syntax	tos-marking-state {trusted untrusted} no tos-marking-state
Context	config>service>ies>interface config>service>ies>sub-if>grp-if
Description	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. 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. 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. 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 save config command will not store the default tos-marking-state trusted state for network IP interfaces unless the detail parameter is also specified. The no 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.
Default	trusted
Parameters	trusted — 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. untrusted — 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.

srrp

Syntax	[no] srrp srrp-id
Context	config>service>ies>sub-if>grp-if config>service>vprn>sub-if>grp-if

Description	This command creates a Subscriber Router Redundancy Protocol (SRRP) instance on a group IP interface. An SRRP instance manages all subscriber subnets within the group interfaces subscriber IP interface or other subscriber IP interfaces that are associated through a wholesale/retail relationship. Only one unique SRRP instance can be configured per group interface. The no form of the command removes an SRRP instance from a group IP interface. Once removed, the group interface ignores ARP requests for the SRRP gateway IP addresses that may exist on subscriber subnets associated with the group IP interface. Then the group interface stops routing using the redundant IP interface associated with the group IP interface and will stop routing with the SRRP gateway MAC address. Ingress packets destined to the SRRP gateway MAC will also be silently discarded. This is the same behavior as a group IP interface that is disabled (shutdown).
Default	no srrp
Parameters	<i>srrp-id</i> — Specifies a 32 bit instance ID that must be unique to the system. The instance ID must also match the instance ID used by the remote router that is participating in the same SRRP context. SRRP is intended to perform a function similar to VRRP where adjacent IP hosts within local subnets use a default gateway to access IP hosts on other subnets. Values 1 — 4294967295

bfd-enable

Syntax	[no] bfd-enable [<i>service-id</i>] interface <i>interface-name</i> dst-ip <i>ip-address</i>
Context	<pre>config>service>ies>sub-if>grp-if>srrp config>service>vprn>sub-if>grp-if>srrp</pre>
Description	This commands assigns a bi-directional forwarding (BFD) session providing heart-beat mechanism for the given VRRP/SRRP instance. There can be only one BFD session assigned to any given VRRP/SRRP instance, but there can be multiple SRRP/VRRP sessions using the same BFD session. BFD control the state of the associated interface. By enabling BFD on a given protocol interface, the state of the protocol interface is tied to the state of the BFD session between the local node and the remote node. The parameters used for the BFD are set via the BFD command under the IP interface. The no form of this command removes BFD from the configuration.
Default	none
Parameters	<i>service-id</i> — Specifies the service ID of the interface running BFD. Values service-id: 1 — 214748364 svc-name: A string up to 64 characters in length. interface <i>interface-name</i> — Specifies the name of the interface running BFD. dst-ip <i>ip-address</i> — Specifies the destination address of the interface running BFD.

gw-mac

Syntax	gw-mac <i>mac-address</i> no gw-mac
Context	config>service>ies>sub-if>grp-if>srrp config>service>vprn>sub-if>grp-if>srrp
Description	This command overrides the default SRRP gateway MAC address used by the SRRP instance. Unless specified, the system uses the same base MAC address for all SRRP instances with the last octet overridden by the lower 8 bits of the SRRP instance ID. . The same SRRP gateway MAC address should be in-use by both the local and remote routers participating in the same SRRP context. One reason to change the default SRRP gateway MAC address is if two SRRP instances sharing the same broadcast domain are using the same SRRP gateway MAC. The system will use the SRRP instance ID to separate the SRRP messages (by ignoring the messages that does not match the local instance ID), but a unique SRRP gateway MAC is essential to separate the routed packets for each gateway IP address. The no form of the command removes the explicit SRRP gateway MAC address from the SRRP instance. The SRRP gateway MAC address can only be changed or removed when the SRRP instance is shutdown.
Parameters	<i>mac-address</i> — Specifies a MAC address that is used to override the default SRRP base MAC address Values Any MAC address except all zeros, broadcast or multicast addresses. The offset is expressed in normal Ethernet MAC address notation. The defined gw-mac cannot be 00:00:00:00:00:00, ff:ff:ff:ff:ff:ff or any multicast address. If not specified, the system uses the default SRRP gateway MAC address with the last octet set to the 8 least significant bits of the SRRP instance ID.

keep-alive-interval

Syntax	keep-alive-interval <i>interval</i> no keep-alive-interval
Context	config>service>ies>sub-if>grp-if>srrp config>service>vprn>sub-if>grp-if>srrp
Description	This command defines the interval between SRRP advertisement messages sent when operating in the master state. The interval is also the basis for setting the master-down timer used to determine when the master is no longer sending. The system uses three times the keep-alive interval to set the timer. Every time an SRRP advertisement is seen that is better then the local priority, the timer is reset. If the timer expires, the SRRP instance assumes that a master does not exist and initiates the attempt to become master. When in backup state, the SRRP instance takes the keep-alive interval of the master as represented in the masters SRRP advertisement message. Once in master state, the SRRP instance uses its own configured keep-alive interval. The keep-alive-interval may be changed at anytime, but will have no effect until the SRRP instance is in the master state.

The **no** form of the command restores the default interval.

Parameters	<i>interval</i> — Specifies the interval between SRRP advertisement messages sent when operating in the master state.
Values	1 — 100 hundreds of milli-seconds
Default	10

message-path

Syntax	message-path <i>sap-id</i> no message-path
Context	config>service>ies>sub-if>grp-if>srrp config>service>vprn>sub-if>grp-if>srrp
Description	This command defines a specific SAP for SRRP in-band messaging. A message-path SAP must be defined prior to activating the SRRP instance. The defined SAP must exist on the SRRP instances group IP interface for the command to succeed and cannot currently be associated with any dynamic or static subscriber hosts. Once a group IP interface SAP has been defined as the transmission path for SRRP Advertisement messages, it cannot be administratively shutdown, will not support static or dynamic subscriber hosts and cannot be removed from the group IP interface. The SRRP instance message-path command may be executed at anytime on the SRRP instance. Changing the message SAP will fail if a dynamic or static subscriber host is associated with the new SAP. Once successfully changed, the SRRP instance will immediately disable anti-spoof on the SAP and start sending SRRP Advertisement messages if the SRRP instance is activated. Changing the current SRRP message SAP on an active pair of routers should be done in the following manner: 1. Shutdown the backup SRRP instance. 2. Change the message SAP on the shutdown node. 3. Change the message SAP on the active master node. 4. Re-activate the shutdown SRRP instance. Shutting down the backup SRRP instance prevents the SRRP instances from becoming master due to temporarily using differing message path SAPs. If an MCS peering is operational between the redundant nodes and the SRRP instance has been associated with the peering, the designated message path SAP will be sent from each member. The no form of the command can only be executed when the SRRP instance is shutdown. Executing no message-path allows the existing SAP to be used for subscriber management functions. A new message-path SAP must be defined prior to activating the SRRP instance.
Parameters	<i>sap-id</i> — Specifies the physical port identifier portion of the SAP definition. See Common Service Commands on page 1020 for <i>sap-id</i> command syntax.

policy

Syntax	[no] policy <i>vrrp-policy-id</i>
Context	config>service>ies>sub-if>grp-if>srrp config>service>vprn>sub-if>grp-if>srrp
Description	This command associates one or more VRRP policies with the SRRP instance. A VRRP policy is a collection of connectivity and verification tests used to manipulate the in-use priorities of VRRP and SRRP instances. A VRRP policy can test the link state of ports, ping IP hosts, discover the existence of routes in the routing table or the ability to reach L2 hosts. When one or more of these tests fail, the VRRP policy has the option of decrementing or setting an explicit value for the in-use priority of an SRRP instance. More than one VRRP policy may be associated with an SRRP instance. When more than one VRRP policy is associated with an SRRP instance the delta decrement of the in-use priority is cumulative unless one or more test fail that have explicit priority values. When one or more explicit tests fail, the lowest priority value event takes effect for the SRRP instance. When the highest delta-in-use-limit is used to manage the lowest delta derived in-use priority for the SRRP instance. VRRP policy associations may be added and removed at anytime. A maximum of two VRRP policies can be associated with a single SRRP instance. The no form of the command removes the association with <i>vrrp-policy-id</i> from the SRRP instance.
Parameters	<i>vrrp-policy-id</i> — Specifies one or more VRRP policies with the SRRP instance.
Values	1 — 9999

priority

Syntax	priority <i>priority</i> no priority
Context	config>service>ies>sub-if>grp-if>srrp config>service>vprn>sub-if>grp-if>srrp
Description	This command overrides the default base priority for the SRRP instance. The SRRP instance priority is advertised by the SRRP instance to its neighbor router and is compared to the priority received from the neighbor router. The router with the best (highest) priority enters the master state while the other router enters the backup state. If the priority of each router is the same, the router with the lowest source IP address in the SRRP advertisement message assumes the master state. The base priority of an SRRP instance can be managed by VRRP policies. A VRRP policy defines a set of connectivity or verification tests which, when they fail, may lower an SRRP instances base priority (creating an in-use priority for the instance). Every time an SRRP instances in-use priority changes when in master state, it sends an SRRP advertisement message with the new priority. If the dynamic priority drops to zero or receives an SRRP Advertisement message with a better priority, the SRRP instance transitions to the <i>becoming backup</i> state. When the priority command is not specified, or the no priority command is executed, the system uses a default base priority of 100. The priority command may be executed at anytime.

The **no** form of the command restores the default base priority to the SRRP instance. If a VRRP policy is associated with the SRRP instance, it will use the default base priority as the basis for any modifications to the SRRP instances in-use priority.

Parameters *priority* — Specifies a base priority for the SRRP instance to override the default.

Values 1 — 254

Default 100

arp-populate

Syntax **[no] arp-populate**

Context config>service>ies>interface
config>service>ies>sub-if>grp-if
config>service>vprn>interface

Description 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 **lease-populate**), and optionally with static entries entered with the **host** command.

Enabling the **arp-populate** command will remove any dynamic ARP entries learned on this interface from the ARP cache.

The **arp-populate** command will fail if an existing static ARP entry exists for this interface.

The **arp-populate** command will fail if an existing static subscriber host on the SAP does not have both MAC and IP addresses specified.

Once **arp-populate** is enabled, creating a static subscriber host on the SAP without both an IP address and MAC address will fail.

When **arp-populate** 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 **arp-populate** enabled. The **arp-populate** command can only be enabled on IES and VPRN interfaces supporting Ethernet encapsulation.

Use the **no** 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.

Default not enabled

arp-timeout

Syntax **arp-timeout seconds**
no arp-timeout

Context config>service>ies>interface
config>service>vprn>interface
config>service>ies>sub-if>grp-if

Description	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 arp-timeout is set to a value of zero seconds, ARP aging is disabled. 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.
Default	14400 seconds
Parameters	<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. Values 0 — 65535

host-connectivity-verify

Syntax	host-connectivity-verify [source { vrrp interface }] [interval <i>interval</i>] [action { remove alarm }]
Context	config>service>ies>if
Description	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.
Default	no host-connectivity-verify
Parameters	source {vrrp interface} — Specifies the source to be used for generation of subscriber host connectivity verification packets. The vrrp keyword specifies that the VRRP state should be used to select proper IP and MAC (active uses VRID, back-up uses interface addresses). The interface 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. interval <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. Values 1 — 6000 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

Syntax	icmp
Context	config>service>ies>interface config>service>ies>sub-if>grp-if config>service>vprn>interface
Description	This command enables the contex to configure Internet Control Message Protocol (ICMP) parameters on a service.

private-retail-subnets

Syntax	[no] private-retail-subnets
Context	config>service>vprn>sub-if
Description	This command controls the export of subnets to the forwarding service. When this attribute is configured, subnets defined on this retail subscriber interface will no longer be exported to the associated wholesale VPRN and will remain private to the retail VPRN. This is useful in a PPPoE business service context as it allows retail services to use overlapping IP address spaces even if these services are associated with the same wholesale service. PPPoE sessions are actually terminated in the retail service although their traffic transits on a SAP belonging to the wholesale service. This configuration is incompatible, however, with IPoE host management (DHCP, static-host and ARP-host) as these host types require that the retail subnets are exported to the wholesale VPRN. Thus, if PPPoE sessions need to coexist with IPoE hosts, this attribute should not be configured on this retail interface. This command will fail if the subscriber interface is not associated with a wholesale service. If the retail VPRN is of the type hub, this attribute is mandatory. Then, it will be enabled by default and it will not be possible to deconfigure it.

local-proxy-arp

Syntax	[no] local-proxy-arp
Context	config>service>ies>interface config>service>ies>sub-if>grp-if config>service>vprn>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.

Triple Play Service Configuration Commands

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
config>service>ies>sub-if>grp-if
config>service>vprn>interface

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
config>service>ies>sub-if>grp-if
config>service>vprn>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

Syntax **redirects** [*number seconds*]
no redirects

Context config>service>ies>if>icmp
config>service>ies>sub-if>grp-if

```
config>service>vprn>interface
```

Description	This command configures the rate for Internet Control Message Protocol (ICMP) redirect messages issued on the router interface. 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. The redirects 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. 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>) The no form of this command disables the generation of icmp redirects on the router interface.
Default	redirects 100 10 — Maximum of 100 redirect messages in 10 seconds.
Parameters	<i>number</i> — The maximum number of ICMP redirect messages to send. This parameter must be specified with the <i>second</i> parameter. Values 10 — 1000 <i>seconds</i> — The time frame in seconds used to limit the <i>number</i> of ICMP redirect messages that can be issued. Values 1 — 60

remote-proxy-arp

Syntax	[no] remote-proxy-arp
Context	<pre>config>service>ies>interface config>service>vprn>interface config>service>ies>sub-if>grp-if</pre>
Description	This command enables remote proxy ARP on the interface. 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.
Default	no remote-proxy-arp

static-arp

Syntax	static-arp ip-address ieee-mac-address no static-arp ip-address ieee-mac-address
Context	<pre>config>service>ies>interface config>service>vprn>interface</pre>

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	<i>ip-address</i> — Specifies the IP address for the static ARP in IP address dotted decimal notation. <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.

ttl-expired

Syntax	ttl-expired <i>number seconds</i> 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	<i>number</i> — The maximum number of ICMP TTL expired messages to send, expressed as a decimal integer. This parameter must be specified with the <i>seconds</i> parameter. Values 10 — 1000 <i>seconds</i> — The time frame in seconds used to limit the <i>number</i> of ICMP TTL expired messages that can be issued, expressed as a decimal integer. Values 1 — 60

unreachables

Syntax	unreachables [<i>number seconds</i>] no unreachables
Context	config>service>ies>if>icmp

Description	This command configures the rate for ICMP host and network destination unreachable messages issued on the router interface. The unreachables 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. By default, generation of ICMP destination unreachables messages is enabled at a maximum rate of 100 per 10 second time interval. The no form of this command disables the generation of icmp destination unreachables on the router interface.
Default	unreachables 100 10
Parameters	<i>number</i> — The maximum number of ICMP unreachable messages to send. This parameter must be specified with the <i>seconds</i> parameter. Values 10 — 1000 <i>seconds</i> — The time frame in seconds used to limit the <i>number</i> of ICMP unreachable messages that can be issued. Values 1 — 60

Interface IPv6 Commands

ipv6

Syntax	[no] ipv6
Context	config>service>ies>interface
Description	This command enables the context to configure IPv6 for an interface.

address

Syntax	address <i>ipv6-address/prefix-length</i> [eui-64] no address <i>ipv6-address/prefix-length</i>		
Context	config>service>ies>if>ipv6		
Description	This command assigns an IPv6 address to the IES interface.		
Parameters	<i>ipv6-address/prefix-length</i> — Specify the IPv6 address on the interface.		
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	
	prefix-length	1 — 128	
	eui-64 — When the eui-64 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.		

dhcp6-relay

Syntax	[no] dhcp6-relay
Context	config>service>ies>if>ipv6
Description	This command enables the context to configure DHCPv6 relay parameters for the IES interface. The no form of the command disables DHCPv6 relay.

neighbor-resolution

Syntax	[no] neighbor-resolution
Context	config>service>ies>if>ipv6>dhcp6-relay
Description	This command enables neighbor resolution with DHCPv6 relay. The no form of the command disables neighbor resolution.

option

Syntax	[no] option
Context	config>service>ies>if>ipv6>dhcp6-relay
Description	This command enables the context to configure DHCPv6 relay information options. The no form of the command disables DHCPv6 relay information options.

interface-id

Syntax	interface-id interface-id ascii-tuple interface-id ifindex interface-id sap-id no interface-id
Context	config>service>ies>if>ipv6>dhcp6>option
Description	This command enables the sending of interface ID options in the DHCPv6 relay packet. The no form of the command disables the sending of interface ID options in the DHCPv6 relay packet
Parameters	ascii-tuple — Specifies that the ASCII-encoded concatenated tuple will be used which consists of the access-node-identifier, service-id, and interface-name, separated by “ ”. ifindex — 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) sap-id — Specifies that the SAP identifier will be used.

remote-id

Syntax	[no] remote-id
Context	config>service>ies>if>ipv6>dhcp6>option
Description	This command enables the sending of remote ID option in the DHCPv6 relay packet.

Triple Play Service Configuration Commands

The client DHCP Unique Identifier (DUID) is used as the remote ID.

The **no** form of the command disables the sending of remote ID option in the DHCPv6 relay packet.

server

Syntax	server <i>ipv6z-address</i> [<i>ipv6z-address</i> ...(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	<i>ipv6-address</i> — Specifies the IPv6 addresses of the DHCP servers where the DHCP6 requests will be forwarded. Up to 8 addresses can be specified.								
Values	<table><tr><td>ipv6-address:</td><td>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></table>	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
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								

source-address

Syntax	source-address <i>ipv6-address</i> no source-address								
Context	config>service>ies>if>ipv6>dhcp6								
Description	This command configures the source IPv6 address of the DHCPv6 relay messages.								
Parameters	<i>ipv6-address</i> — Specifies the source IPv6 address of the DHCPv6 relay messages.								
Values	<table><tr><td>ipv6-address:</td><td>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></table>	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
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

Syntax	[no] dhcp6-server
Context	config>service>ies>if>ipv6
Description	This command enables the context to configure DHCPv6 server parameters for the IES interface.

The **no** form of the command disables the DHCP6 server.

max-nbr-of-leases

Syntax	max-nbr-of-leases <i>max-nbr-of-leases</i> no max-nbr-of-leases
Context	config>service>ies>if>ipv6>dhcp6-server
Description	This command configures the maximum number of lease states installed by the DHCP6 server function allowed on this interface. The no form of the command returns the value to the default.
Default	8000
Parameters	<i>max-nbr-of-leases</i> — Specifies the maximum number of lease states installed by the DHCP6 server function allowed on this interface. Values 0 — 8000

prefix-delegation

Syntax	[no] prefix-delegation
Context	config>service>ies>if>ipv6>dhcp6-server
Description	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 no form of the command disables prefix-delegation.

prefix

Syntax	[no] prefix <i>ipv6-address/prefix-length</i>		
Context	config>service>ies>if>ipv6>dhcp6-server>pfx-delegate		
Description	This command specifies the IPv6 prefix that will be delegated by this system.		
Parameters	<i>ipv6-address/prefix-length</i> — Specify the IPv6 address on the interface.		
	Values	ipv6-address/prefix: ipv6-address	x::x::x::x::x::x (eight 16-bit pieces) x::x::x::x:d.d.d.d x [0 — FFFF]H d [0 — 255]D prefix-length 1 — 128

duid

Syntax	duid <i>duid</i> [<i>iaid</i> <i>iaid</i>] no duid
Context	config>service>ies>if>ipv6>dhcp6>pfx-delegate>prefix
Description	This command configures the DHCP Unique Identifier (DUID) of the DHCP client.
Parameters	<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. iaid <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.

preferred-lifetime

Syntax	preferred-lifetime <i>seconds</i> preferred-lifetime infinite no preferred-lifetime
Context	config>service>ies>if>ipv6>dhcp6>pfx-delegate>prefix
Description	This command configures the IPv6 prefix/mask preferred life time. The preferred-lifetime value cannot be bigger than the valid-lifetime value. The no form of the command reverts to the default value.
Default	604800 seconds (7 days)
Parameters	<i>seconds</i> — Specifies the time, in seconds, that this prefix remains preferred. Values 1 — 4294967294 infinite — Specifies that this prefix remains preferred infinitely.

valid-lifetime

Syntax	valid-lifetime <i>seconds</i> valid-lifetime infinite no valid-lifetime
Context	config>service>ies>if>ipv6>dhcp6>pfx-delegate>prefix
Description	This command configures the time, in seconds, that the prefix is valid. The maximum value 4294967295 is considered equal to infinity. The no form of the command reverts to the default value.
Default	2592000 seconds (30 days)
Parameters	<i>seconds</i> — Specifies the time, in seconds, that this prefix remains valid.

Values 1 — 4294967294

infinite — Specifies that this prefix remains valid infinitely.

icmp6

Syntax	icmp6
Context	config>service>ies>if>ipv6
Description	This command configures ICMPv6 parameters for the interface.

packet-too-big

Syntax	packet-too-big [<i>number seconds</i>] no packet-too-big
Context	config>service>ies>if>ipv6>icmp6
Description	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 no form of the command disables the sending of ICMPv6 “packet-too-big” messages.
Default	100 10
Parameters	<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. Values 10 — 1000 Default 100 <i>seconds</i> — Specifies the time frame in seconds that is used to limit the number of “packet-too-big” ICMP messages issued. Values 1 — 60 Default 10

param-problem

Syntax	param-problem [<i>number seconds</i>] no packet-too-big
Context	config>service>ies>if>ipv6>icmp6
Description	This command specifies whether “parameter-problem” ICMP messages should be sent. When enabled, “parameter-problem” ICMP messages are generated by this interface. The no form of the command disables the sending of “parameter-problem” ICMP messages.

Triple Play Service Configuration Commands

Default	100 10
	<i>number</i> — Specifies the number of “parameter-problem” ICMP messages to send in the time frame specified by the <i>seconds</i> parameter.
Values	10 — 1000
Default	100
	<i>seconds</i> — 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 [<i>number seconds</i>] 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
	<i>number</i> — Specifies the number of version 6 redirects are to be issued in the time frame specified by the <i>seconds</i> parameter.
Values	10 — 1000
Default	100
	<i>seconds</i> — 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 [<i>number seconds</i>] 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

Syntax **neighbor** *ipv6-address mac-address*
no neighbor *ipv6-address*

Context config>service>ies>if>ipv6

Description This command configures IPv6-to-MAC address mapping on the IES interface.

Parameters *ipv6-address* — The IPv6 address of the interface for which to display information.

Values 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 [1..128]

mac-address — 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.

proxy-nd-policy

Syntax **proxy-nd-policy** *policy-name [policy-name...(up to 5 max)]*
no proxy-nd-policy

Context config>service>ies>if>ipv6

Description 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.

The **no** form of this command reverts to the default value.

Default no proxy-nd-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.

Show Commands

egress-label

Syntax `egress-label egress-label1 [egress-label2]`

Context `show>service`

Description Display services using the range of egress labels.

If only the mandatory *egress-label1* parameter is specified, only services using the specified label are displayed.

If both *egress-label1* and *egress-label2* parameters are specified, the services using the range of labels X where *egress-label1* <= X <= *egress-label2* are displayed.

Use the **show router ldp bindings** command to display dynamic labels.

Parameters *egress-label1* — The starting egress label value for which to display services using the label range. If only *egress-label1* is specified, services only using *egress-label1* are displayed.

Values 0, 2049 — 131071

egress-label2 — The ending egress label value for which to display services using the label range.

Default The *egress-label1* value.

Values 2049 — 131071

Output **Show Service Egress Command Output** — 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

Syntax	fdb-info
Context	show>service
Description	Displays global FDB usage information.
Output	Show FDB-Info Command Output — 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 MAC's 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 MAC's. The rate is computed as the maximum number of re-learns allowed in a 5 second interval: e.g. 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 MAC's 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

Syntax	fdb-mac <i>ieee-address</i>
Context	show>service
Description	Displays the FDB entry for a given MAC address.
Parameters	<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.
Output	Show FDB-MAC Command Output — 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

Syntax	ingress-label <i>ingress-label1</i> [<i>ingress-label2</i>]
Context	show>service
Description	Display services using the range of ingress labels. If only the mandatory <i>ingress-label1</i> parameter is specified, only services using the specified label are displayed. 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> <= X <= <i>ingress-label2</i> are displayed. Use the show router vprn-service-id ldp bindings command to display dynamic labels.
Parameters	<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. Values 0, 2048 — 131071 <i>ingress-label2</i> — The ending ingress label value for which to display services using the label range. Default The <i>ingress-label1</i> value. Values 2048 — 131071
Output	Show Service Ingress Label — 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

Syntax **sap-using** [**sap** *sap-id*]
sap-using **interface** [*ip-address* | *ip-int-name*]
sap-using [**ingress** | **egress**] **atm-td-profile** *td-profile-id*
sap-using [**ingress** | **egress**] **filter** *filter-id*
sap-using [**ingress** | **egress**] **qos-policy** *qos-policy-id*
sap-using **authentication-policy** *auth-plcy-name*

Context show>service

Description Displays SAP information.
 If no optional parameters are specified, the command displays a summary of all defined SAPs.
 The optional parameters restrict output to only SAPs matching the specified properties.

Parameters **ingress** — Specifies matching an ingress policy.
egress — Specifies matching an egress policy.
qos-policy *qos-policy-id* — The ingress or egress QoS Policy ID for which to display matching SAPs.
 Values 1 — 65535
atm-td-profile *td-profile-id* — Displays SAPs using this traffic description.
filter *filter-id* — The ingress or egress filter policy ID for which to display matching SAPs.
 Values 1 — 65535
authentication *auth-plcy-name* — The session authentication policy for which to display matching SAPs.
sap-id — Specifies the physical port identifier portion of the SAP definition. See [Common Service Commands on page 1020](#) for *sap-id* command syntax.
interface — Specifies matching SAPs with the specified IP interface.
ip-addr — The IP address of the interface for which to display matching SAPs.
 Values 1.0.0.0 — 223.255.255.255
ip-int-name — The IP interface name for which to display matching SAPs.

Output **Show Service SAP** — 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.

Label	Description (Continued)
E.QoS	The SAP egress QoS policy number specified on the egress SAP.
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#
```

sdp

Syntax	sdp [<i>sdp-id</i> far-end <i>ip-addr</i>] [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	<i>sdp-id</i> — Displays only information for the specified SDP ID. Default All SDPs Values 1 — 17407 far-end <i>ip-addr</i> — 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.

Label	Description (Continued)
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.
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.

Label	Description (Continued)
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

Syntax	sdp-using [<i>sdp-id</i> [: <i>vc-id</i>] <i>far-end ip-address</i>]
Context	show>service
Description	Display services using SDP or far-end address options.
Parameters	<i>sdp-id</i> — Displays only services bound to the specified SDP ID. Values 1 — 17407 <i>vc-id</i> — The virtual circuit identifier. Values 1 — 4294967295 <i>far-end ip-address</i> — 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.

Label	Description (Continued)
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] [vprn] [sdp sdp-id]
Context	show>service
Description	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. vprn — Displays matching VPRN 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 summary active-subscribers [subscriber <i>sub-ident-string</i> [sap <i>sap-id</i> sla-profile <i>sla-profile-name</i>]] [detail] active-subscribers hierarchy [subscriber <i>sub-ident-string</i>] host-tracking [subscriber <i>sub-ident-string</i>] [detail statistics summary] host-tracking groups [group <i>group-ip-address</i>] [detail summary]
Context	show>service
Description	This command displays active subscriber information.
Parameters	sap <i>sap-id</i> — Displays SAP information for the specified SAP ID. See Common Service Commands on page 1020 for <i>sap-id</i> command syntax. sla-profile <i>sla-profile-name</i> — Displays information for the specified SLA profile name.

summary — Displays active subscriber information in a brief format.

subscriber *sub-ident-string* — Displays information for the specified subscriber identification string.

hierarchy — Displays the subscriber hierarchy.

detail — Displays detailed output.

Sample Output

```
A:Dut-Q# show service active-subscribers
-----
(1) SLA Profile Instance sap:lag-2:421 - sla:sipslam
-----
IP Address          MAC Address          PPPoE-SID Origin
-----
1.1.1.1             Unknown              N/A              Static
-----
A:Dut-Q#

A:Dut-A# show service active-subscribers summary
=====
Active Subscriber table summary
=====
Total Count          : 6
=====
A:Dut-A#

A:Dut-A# show service active-subscribers hierarchy
=====
Active Subscriber hierarchy
=====
-- alcatel_100 (sub_default)
|
|-- sap:1/2/1:100 - sla:sla_default
| |
| |-- 10.100.1.3 - 00:10:00:00:00:01 (-/D/-)
| |
| |
| |-- sap:1/2/1:101 - sla:sla_default
| |
| |-- 10.100.1.4 - 00:10:00:00:00:02 (-/D/-)
| |
| |
| |-- sap:1/2/1:102 - sla:sla_default
| |
| |-- 10.100.1.5 - 00:10:00:00:00:03 (-/D/-)
| |
| |
-- alcatel_110 (sub_default)
|
|-- sap:1/2/1:110 - sla:sla_default
| |
| |-- 10.110.1.3 - 00:10:10:00:00:01 (-/D/-)
| |
| |
```

Triple Play Service Configuration Commands

```
|
|-- sap:1/2/1:111 - sla:sla_default
|   |
|   |-- 10.110.1.4 - 00:10:10:00:00:02 (-/D/-)
|   |
|
|-- sap:1/2/1:112 - sla:sla_default
|   |
|   |-- 10.110.1.5 - 00:10:10:00:00:03 (-/D/-)
|   |

-- alcatel_120 (sub_default)
|
|-- sap:1/2/1:120 - sla:sla_default
|   |
|   |-- 10.120.1.3 - 00:10:20:00:00:01 (-/D/-)
|   |
|
|-- sap:1/2/1:121 - sla:sla_default
|   |
|   |-- 10.120.1.4 - 00:10:20:00:00:02 (-/D/-)
|   |
|
|-- sap:1/2/1:122 - sla:sla_prof120_VOIP
|   |
|   |-- 10.120.1.5 - 00:10:20:00:00:03 (-/D/-)
|   |

-- alcatel_130 (sub_default)
|
|-- sap:1/2/1:130 - sla:sla_default
|   |
|   |-- 10.130.1.3 - 00:10:30:00:00:01 (-/D/-)
|   |

-- alcatel_140 (sub_default)
|
|-- sap:1/2/1:140 - sla:sla_default
|   |
|   |-- 10.140.1.3 - 00:10:40:00:00:01 (-/D/-)
|   |

-- alcatel_80 (sub_default)
|
|-- sap:1/2/1:80 - sla:sla_default
|   |
|   |-- 10.80.1.3 - 00:80:00:00:00:01 (-/D/-)
|   |
|   |-- 10.80.1.4 - 00:80:00:00:00:02 (-/D/-)
|   |
|   |-- 10.80.1.5 - 00:80:00:00:00:03 (-/D/-)
|   |

-- alcatel_81 (sub_prof81)
|
|-- sap:1/2/1:80 - sla:sla_prof81_VOIP
|   |
|   |-- 10.80.1.6 - 00:80:00:00:00:04 (-/D/-)
|   |

-- alcatel_90 (sub_default)
```

```

|
|-- sap:1/2/1:90 - sla:sla_default
| |
| |-- 10.90.1.3 - 00:90:00:00:00:01 (-/D/-)
| |
|

-- client_PC1 (sub_profPC1)
|
|-- sap:1/2/2:4000 - sla:sla_profPC1
| |
| |-- 0.0.0.0 - 00:00:00:00:00:00 (-/-/N)
| |
| |-- 10.24.1.253 - 00:13:21:67:a4:cd (-/D/-)
| |
|
|-- sap:lag-1 - sla:sla_profPC1
| |
| |-- 1.2.3.4 - 00:05:04:03:02:01 (S/-/-)
| |

-- static (sub_default)
|
|-- sap:1/2/1:80 - sla:sla_default
| |
| |-- 10.80.123.123 - 00:00:12:34:56:78 (S/-/-)
| |

=====
A:Dut-A#

A:Dut-A# show service active-subscribers subscriber alcatel_100 hierarchy
=====
Active Subscriber hierarchy
=====
-- alcatel_100 (sub_prof100)
|
|-- sap:1/2/1:101 - sla:sla_prof100_VOIP
| |
| |-- 10.100.1.4 - 00:10:00:00:00:02 (-/D/-)
| |
|
|-- sap:1/2/1:102 - sla:sla_default
| |
| |-- 10.100.1.5 - 00:10:00:00:00:03 (-/D/-)
| |

=====
A:Dut-A#

A:Dut-A# show service active-subscribers subscriber alcatel_100
=====
Active Subscribers
-----
Subscriber alcatel_100 (sub_default)
-----
(1) SLA Profile Instance sap:1/2/1:100 - sla:sla_default
-----
IP Address      MAC Address      Origin(*)
-----
10.100.1.3      00:10:00:00:00:01 -/D/-

```

Triple Play Service Configuration Commands

```
-----  
(2) SLA Profile Instance sap:1/2/1:101 - sla:sla_default  
-----
```

IP Address	MAC Address	Origin(*)

10.100.1.4	00:10:00:00:00:02	-/D/-

```
-----  
(3) SLA Profile Instance sap:1/2/1:102 - sla:sla_default  
-----
```

IP Address	MAC Address	Origin(*)

10.100.1.5	00:10:00:00:00:03	-/D/-

```
=====
```

(*) S=Static Host, D=DHCP Lease, N=Non-Sub-Traffic

```
=====
```

A:Dut-A#

A:Dut-A# show service active-subscribers subscriber alcatel_100 sap 1/2/1:100 sla-profile sla_default

```
=====
```

Active Subscribers

Subscriber alcatel_100 (sub_default)

```
-----  
(1) SLA Profile Instance sap:1/2/1:100 - sla:sla_default  
-----
```

IP Address	MAC Address	Origin(*)

10.100.1.3	00:10:00:00:00:01	-/D/-

```
=====
```

(*) S=Static Host, D=DHCP Lease, N=Non-Sub-Traffic

```
=====
```

A:Dut-A#

A:Dut-A# show service active-subscribers subscriber alcatel_100 sap 1/2/1:100 sla-profile sla_default detail

```
=====
```

Active Subscribers

Subscriber alcatel_100 (sub_default)

```
-----
```

I. Sched. Policy : service_all

E. Sched. Policy : service_all

Acct. Policy : N/A

Collect Stats : Disabled

```
-----
```

(1) SLA Profile Instance

- sap:1/2/1:100 (VPLS 100)
- sla:sla_default

```
-----
```

Host Limit	: No Limit

Ingress Qos-Policy	: 1000

Egress Qos-Policy	: 1000

Ingress Queuing Type	: Service-queuing

Ingress Filter-Id	: N/A

Egress Filter-Id	: N/A

```
-----
```

IP Address	MAC Address	Origin(*)

10.100.1.3	00:10:00:00:00:01	-/D/-

```
-----
```

SLA Profile Instance statistics

	Packets	Octets
Off. HiPrio	: 0	0
Off. LowPrio	: 42361	8639977
Off. Uncolor	: 0	0
Queueing Stats (Ingress QoS Policy 1000)		
Dro. HiPrio	: 0	0
Dro. LowPrio	: 6783	1392451
For. InProf	: 0	0
For. OutProf	: 35392	7211148
Queueing Stats (Egress QoS Policy 1000)		
Dro. InProf	: 6599	1347340
Dro. OutProf	: 0	0
For. InProf	: 34364	7011246
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	: 0	0
Off. Uncolor	: 0	0
Dro. HiPrio	: 0	0
Dro. LowPrio	: 0	0
For. InProf	: 0	0
For. OutProf	: 0	0
Ingress Queue 3 (Unicast) (Priority)		
Off. HiPrio	: 0	0
Off. LowPrio	: 42361	8639977
Off. Uncolor	: 0	0
Dro. HiPrio	: 0	0
Dro. LowPrio	: 6783	1392451
For. InProf	: 0	0
For. OutProf	: 35392	7211148
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	: 0	0
Dro. OutProf	: 0	0
For. InProf	: 0	0

Triple Play Service Configuration Commands

```
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           : 6599                   1347340
Dro. OutProf          : 0                      0
For. InProf           : 34364                   7011246
For. OutProf          : 0                      0
=====
(*) S=Static Host, D=DHCP Lease, N=Non-Sub-Traffic
A:Dut-A#
```

id

Syntax	id <i>service-id</i>
Context	show>service
Description	Enables the context to display information for a particular service-id.
Parameters	<i>service-id</i> — The unique service identification number that identifies the service in the service domain. all — Displays detailed information about the service. arp — Displays ARP entries for the service. arp-host — Displays ARP host related information. base — Display sbasic service information. fdb — Displays FDB entries. 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

Syntax	all
Context	show>service>id
Description	Displays detailed information for all aspects of the service.

Output **Show All Service-ID Output** — 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.
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.

Label	Description (Continued)
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
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.

Label	Description (Continued)
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-policy	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-policy	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 Priority	The number of high priority packets, as determined by the SAP ingress QoS policy.
Offered Low Priority	The number of low priority packets, as determined by the SAP ingress QoS policy.
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.
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.

Label	Description (Continued)
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
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                     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          : Blockable
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

```

Triple Play Service Configuration Commands

```
RST BPDUs rcvd      : 0                      RST BPDUs tx       : 0
-----
Sdp Id 5:750  -(128.251.10.49)
-----
SDP Id              : 5:750                      Type           : Mesh
VC Type             : Ether                      VC Tag            : n/a
Admin Path MTU      : 0                          Oper Path MTU     : 0
Far End             : 128.251.10.49              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                          I. Dro. Pkts.     : 0
E. Fwd. 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                      Type           : Mesh
VC Type             : Ether                      VC Tag            : n/a
Admin Path MTU      : 0                          Oper Path MTU     : 0
Far End             : 10.10.10.106              Delivery          : MPLS
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                          I. Dro. Pkts.     : 0
E. Fwd. Pkts.       : 0                          E. Fwd. Octets    : 0

Associated LSP LIST :
Lsp Name            : to-49
Admin State          : Down                      Oper State         : Down
Time Since Last Tr* : 02h01m08s
-----
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
Oper State      : Down
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
Total MAC Addr   : 0
Learned MAC Addr : 0
Static MAC Addr  : 0
Admin MTU        : 1518
Oper MTU         : 1518
Ingress qos-policy : 1
Egress qos-policy : 1
Ingress Filter-Id : n/a
Egress Filter-Id  : n/a
Mac Learning     : Enabled
Discard Unkwn Srce: Disabled
Mac Aging        : Enabled

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
Rstp Oper State  : Down
Port Role       : N/A
Port State      : Discarding
Port Number     : 2048
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

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
Off. LoPrio      : 0
Dro. HiPrio      : 0
Dro. LoPrio      : 0
For. InProf      : 0
For. OutProf     : 0

Ingress Queue 11 (Multipoint) (Priority)
Off. HiPrio      : 0
Off. LoPrio      : 0
Dro. HiPrio      : 0
Dro. LoPrio      : 0
For. InProf      : 0
For. OutProf     : 0

Egress Queue 1
For. InProf      : 0
For. OutProf     : 0
Dro. InProf      : 0
Dro. OutProf     : 0
-----
VPLS Rapid Spanning Tree Information
-----
VPLS oper state      : Down
Rstp Admin State     : Up
Mode                  : Rstp
Core Connectivity    : Down
Rstp Oper State      : Down
Vcp Active Prot.     : N/A

Bridge Id            : 80:01.14:30:ff:00:00:01
Bridge Priority       : 32768
Topology Change      : Inactive
Last Top. Change     : 0d 00:00:00
Top. Change Count    : 0
Bridge Instance Id   : 1
Tx Hold Count        : 6
Bridge Hello Time    : 2
Bridge Max Age       : 20
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 1020](#) 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. 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
                  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
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.

Label	Description (Continued)
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.
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 1020](#) 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

- Syntax** **host** [**sap** *sap-id*]
- Context** show>service>id
- Description** This command displays static hosts configured on this service.

Triple Play Service Configuration Commands

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition. See [Common Service Commands on page 1020](#) for *sap-id* 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

Syntax **labels**

Context show>service>id

Description Displays the labels being used by the service.

Output **Show Service-ID Labels** — 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

Syntax **mfib brief**
mfib [group *grp-address*]
mfib statistics [group *grp-address*]

Context show>service>id>

Description This command displays the multicast FIB on the VPLS service.

Parameters **group *grp-ip-address*** — Displays the multicast FIB for a specific multicast group address.
statistics — Displays statistics on the multicast FIB.

Output **Show Output** — 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.

Sample Output

```

A:ALA-1>show>service>id # mfib
=====
IGMP Snooping MFIB for service 10
=====
Source Address  Group Address  Sap/Sdp Id      Fwd/Blk
-----
*              225.0.0.1     sap:2/1/5:1     Fwd

```

Triple Play Service Configuration Commands

```
*                225.0.0.7                sap:2/1/5:7                Fwd
-----
Number of entries: 2

*A:rbac_C# show service id 1 mfib statistics
=====
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:rbac_C#
```

mld-snooping

Syntax	mld-snooping
Context	show>service>id
Description	This command displays MLD snooping information.

all

Syntax	all
Context	show>service>id>mld-snooping
Description	This command displays detailed information about MLD snooping.

Sample Output

```
*A:rbac_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
-----
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

```

Triple Play Service Configuration Commands

```

-----
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

Syntax **base**

Context show>service>id>mld-snooping

Description This command displays basic MLD snooping information.

mrouters

Syntax **mrouters [detail]**

Context show>service>id>mld-snooping

Description This command displays all multicast routers.

Sample Output

```

*A:rbac_C# show service id 1 mld-snooping mrouters
=====
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 mrouters: 1
=====
*A:rbac_C#

```

```

*A:rbac_C# show service id 1 mld-snooping mrouters 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 mrouters: 1
=====

```

Triple Play Service Configuration Commands

mvr

Syntax	mvr
Context	show>service>id>mld-snooping
Description	This command displays multicast VPLS registration information.

port-db

Syntax	port-db sap <i>sap-id</i> port-db sap <i>sap-id</i> detail port-db sap <i>sap-id</i> group <i>grp-ipv6-address</i> port-db sdp <i>sdp-id:vc-id</i> [detail] port-db sdp <i>sdp-id:vc-id</i> group <i>grp-ipv6-address</i>
Context	show>service>id>mld-snooping
Description	This command displays MLD snooping information related to a specific SAP.

Sample Output

```
*A:rbac_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:rbac_C#

*A:rbac_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:rbac_C#

```

proxy-db

Syntax	proxy-db [detail] proxy-db group grp-ipv6-address
Context	show>service>id>mld-snooping
Description	This command displays proxy-reporting database entries.

Sample Output

```

*A:rbac_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:rbac_C#

*A:rbac_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

```

Triple Play Service Configuration Commands

```
-----  
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

Syntax	querier
Context	show>service>id>mld-snooping
Description	This command displays information about the current querier.

Sample Output

```
*A:rbae_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:rbae_C#
```

static

Syntax	static [sap sap-id sdp sdp-id:vc-id]
Context	show>service>id>mld-snooping
Description	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

Syntax	statistics[sap <i>sap-id</i> sdp <i>sdp-id:vc-id</i>]
Context	show>service>id>mld-snooping
Description	This command displays MLD snooping statistics.

Sample Output

```
*A:rbac_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
```

Triple Play Service Configuration Commands

```
MCS Failures           : 0

MVR From VPLS Cfg Drops : 0
MVR To SAP Cfg Drops    : 0
=====
*A:rbae_C#
```

mstp-configuration

Syntax	mstp-configuration
Context	show>service>id
Description	This command displays the MSTP specific configuration data. This command is only valid on a management VPLS.

sap

Syntax	sap sap-id [detail]
Context	show>service>id
Description	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.
Parameters	<i>sap-id</i> — The ID that displays SAP information. See Common Service Commands on page 1020 for <i>sap-id</i> command syntax. detail — Displays detailed information for the SAP.
Output	Show Service-ID SAP — 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 Ether type 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

Label	Description (Continued)
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.
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.

Label	Description (Continued)	
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	Total MAC Addr : 0
Learned MAC Addr	: 0	Static MAC Addr : 0
Admin MTU	: 1518	Oper MTU : 1518
Ingress qos-policy	: 1	Egress qos-policy : 1
Ingress Filter-Id	: n/a	Egress Filter-Id : n/a
Mac Learning	: Enabled	Discard Unkwn Srce: Disabled
Mac Aging	: Enabled	
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	Encap : q-tag
Dot1Q Ethertype	: 0x8100	QinQ Ethertype : 0x8100
Admin State	: Up	Oper State : Down
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	Total MAC Addr : 0
Learned MAC Addr	: 0	Static MAC Addr : 0
Admin MTU	: 1518	Oper MTU : 1518
Ingress qos-policy	: 100	Egress qos-policy : 1
Shared Q plcy	: default	Multipoint shared : Enabled
Ingr IP Fltr-Id	: n/a	Egr IP Fltr-Id : 10
Ingr Mac Fltr-Id	: n/a	Egr Mac Fltr-Id : n/a
Ingr IPv6 Fltr-Id	: n/a	Egr IPv6 Fltr-Id : n/a
tod-suite	: None	qinq-pbit-marking : both
ARP Reply Agent	: Enabled	Host Conn Verify : Enabled
Mac Learning	: Enabled	Discard Unkwn Srce: Disabled
Mac Aging	: Enabled	Mac Pinning : Disabled
L2PT Termination	: 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	Nbr Static Hosts : 1
MCAC Policy Name	:	MCAC Const Adm St : Enable
MCAC Max Unconst BW:	no limit	MCAC Max Mand BW : no limit
MCAC In use Mand BW:	0	MCAC Avail Mand BW: unlimited
MCAC In use Opnl BW:	0	MCAC Avail Opnl BW: unlimited

Auth Policy : none
 Egr MCast Grp :

Rstp Service Access Point specifics

Mac Move	: Blockable		
Rstp Admin State	: Up	Rstp Oper State	: Down
Core Connectivity	: Down		
Port Role	: N/A	Port State	: Discarding
Port Number	: 2048	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		

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
=====
*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.)

Label	Description (Continued)
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.
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
```

Triple Play Service Configuration Commands

```
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

Syntax	gsmp
Context	show>service>id
Description	This enables the command to display GSMP information.

neighbors

Syntax	neighbors group [<i>name</i>] [<i>ip-address</i>]
Context	show>service>id>gsmp
Description	This command display GSMP neighbor information.
Parameters	group — 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.
Output	Show Service-ID GSMP Neighbors Group — 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.

Label	Description (Continued)
Sessions	Displays the number of sessions (open TCP connections) for each configured neighbor.

Sample Output

These commands show the configured neighbors per service, regardless of the fact 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
-----
dslaml                     192.168.1.2            Enabled     0
dslaml                     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 dslaml
=====
GSMP neighbors
=====
Group                      Neighbor                AdminState  Sessions
-----
dslaml                     192.168.1.2            Enabled     0
dslaml                     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 dslaml 192.168.1.2
=====
GSMP neighbors
=====
Group                      Neighbor                AdminState  Sessions
-----
dslaml                     192.168.1.2            Enabled     0
=====
A:active>show>service>id>gsmp#
```

sessions

Syntax	sessions [group <i>name</i>] neighbor <i>ip-address</i>] [port <i>port-number</i>] [association] [statistics]
Context	show>service>id>gsmp
Description	This command displays GSMP sessions information.
Parameters	group — A GSMP group defines a set of GSMP neighbors which have the same properties. name — Specifies a GSMP group name is unique only within the scope of the service in which it is defined. ip-address — Specifies the ip-address of the neighbor. port — Specifies the neighbor TCP port number use for this ANCP session. Values 0 — 65535 association — Displays to what object the ANCP-string is associated. statistics — Displays statistics information about an ANCP session known to the system.
Output	Show Service-ID GSMP sessions — 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.

Label	Description (Continued)
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 (VPRN)
=====
```

```
Port    Ngbr-IPAddr    Gsmp-Group
-----
40590   192.168.1.2    dslaml
-----
```

```
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 (VPRN), 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 (VPRN), 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>gsmp
```

host

Syntax	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] host [sap <i>sap-id</i>] [wholesaler <i>service-id</i>] [port <i>port-id</i>] 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 <i>sap-id</i> — Specifies the physical port identifier portion of the SAP definition. See Common Service Commands on page 1020 for <i>sap-id</i> command syntax. <i>intermediate-destination-id</i> — 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.

wholesaler *service-id* — The service ID of the wholesaler.

Values 1 — 2147483647

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

Syntax	wholesalers
Context	show>service>id
Description	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
=====
```

Triple Play Service Configuration Commands

```
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

Syntax	split-horizon-group [<i>group-name</i>]
Context	show>service>id
Description	This command displays service split horizon groups.
Parameters	<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
-----
DSL-group1            Split horizon group for DSL
-----
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
-----
DSL-group1            Split horizon group for DSL
-----
Associations
-----
SAP                   1/1/3:1
SDP                   108:1
SDP                   109:1
-----
SAPs Associated : 1      SDPs Associated : 2
=====
A:ALA-1#
```

static-host

Syntax	static-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] static-host [sap <i>sap-id</i>] [wholesaler <i>service-id</i>] [port <i>port-id</i>] no-inter-dest-id [detail] static-host summary
Context	show>service>id
Description	This command displays Display static hosts configured on this service.
Parameters	sap <i>sap-id</i> — Specifies the physical port identifier portion of the SAP definition. See Common Service Commands on page 1020 for <i>sap-id</i> command syntax. <i>intermediate-destination-id</i> — 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. wholesaler <i>service-id</i> — The service ID of the wholesaler. Values 1 — 2147483647

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

Syntax	stp [detail]
Context	show>service>id
Description	Displays information for the spanning tree protocol instance for the service.
Parameters	detail — Displays detailed information.

Output **Show Service-ID STP Output** — 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 priority : 1
Last Top. change : 0d 00:00:00

Bridge max age : 20
Topology change : Inactive
Top. change count : 0

Root bridge-id : 00:03:fa:00:00:00

Root path cost : 1
Root hello time : 1
Root priority : 0

Root forward delay: 15
Root max age : 20
Root port : vcp

-----
Spanning Tree Specifics
-----
SAP Identifier : 1/1/7:0
STP Port State : Forwarding
Port Number : 2048
Cost : 10
Designated Port : 34816

RSTP State : Down
BPDU encap : dot1d
Priority : 128
Fast Start : Disabled
Designated Bridge : 02:fa:00:04:54:01
=====
A:ALA-12#

```

authentication

Syntax	authentication
Context	show>service>id
Description	This command enables the context to show session authentication information.

statistics

Syntax	statistics [policy name] [sap sap-id]
Context	show>service>id>authentication
Description	This command displays session authentication statistics for this service.
Parameters	policy name — Specifies an existing authentication policy name. sap-id — Specifies the physical port identifier portion of the SAP definition. See Common Service Commands on page 1020 for <i>sap-id</i> command syntax.

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

- Syntax** **subscriber-hosts** [**sap** *sap-id*] [**ip** *ip-address[/mask]*] [**mac** *ieee-address*] [**sub-profile** *sub-profile-name*] [**sla-profile** *sla-profile-name*] [**detail**]
- Context** show>service>id
- Description** This command displays subscriber host information.
- Parameters** **sap** *sap-id* — Specifies the physical port identifier portion of the SAP definition. See [Common Service Commands on page 1020](#) for *sap-id* command syntax.
- ip** *ip-address[/mask]* — Shows information for the specified IP address and mask.
- mac** *ieee-address* — Displays information only for 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.
- profile** *sub-profile-name* — Displays 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* — Displays this optional parameter to specify an existing SLA profile name.
- detail** — Displays detailed information.
- Output** **Show Service-ID subscriber hosts** — 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 Sub-scriber	The ID of the originating subscriber.

Sample Output

```
A:Dut-R# show service id 100 subscriber-hosts
=====
Subscriber Host table
=====
Sap              Subscriber
IP Address
MAC Address      PPPoE-SID Origin    Fwding State
-----
lag-2:421        sub
1.1.1.1
Unknown          N/A      Static    Not Fwding
-----
Number of subscriber hosts : 1
=====
```


A:Dut-R#

```
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#
```

subscriber-using

Syntax **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*]

Context show>service

Description This command displays selective subscriber information using specific options.

Parameters **service-id** *service-id* — Displays information for the specifies ID that uniquely identifies a service.

sap-id *sap-id* — Displays the physical port identifier portion of the SAP definition. See [Common Service Commands on page 1020](#) for *sap-id* command syntax.

interface *ip-int-name* — Shows DHCP statistics on the specified interface.

port *port-id* — Indicates the SAP or SDP for which this entry contains information.

ip *ip-address[/mask]* — Shows information for the specified IP address and mask.

mac *ieee-address* — Displays information only for 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.

sub-profile *sub-profile-name* — Displays 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* — Displays this optional parameter to specify an existing SLA profile name.

app-profile — Displays the application specified profile. For further information about application profile parameters refer to the 7750 SR-Series OS Integrated Services Adapter Guide.

inter-dest-id *intermediate-destination-id* — Indicates the intermediate destination identifier received from either the DHCP or the RADIUS server or the local user database.

Sample Output

Triple Play Service Configuration Commands

```
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
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

Syntax	redundancy
Context	show
Description	This command enables the context to show multi-chassis redundancy information.

multi-chassis

Syntax	multi-chassis all multi-chassis mc-lag peer <i>ip-address</i> [<i>lag lag-id</i>] multi-chassis mc-lag [<i>peer ip-address</i> [<i>lag lag-id</i>]] statistics multi-chassis sync [<i>peer ip-address</i>] [<i>detail</i>] multi-chassis sync [<i>peer ip-address</i>] statistics
Context	show>redundancy
Description	This command displays multi-chassis redundancy information.
Parameters	all — Displays all multi-chassis information. mc-lag — Displays multi-chassis LAG information.

peer ip-address — Displays the address of the multi-chassis peer.

lag lag-id — Displays the specified LAG ID on this system that forms an multi-chassis LAG configuration with the indicated peer.

statistics — Displays statistics for the multi-chassis peer.

sync — Displays synchronization information.

detail — 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
```

Triple Play Service Configuration Commands

```
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
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

```

Triple Play Service Configuration Commands

```
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
=====
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:pcl#

```

```

A:pcl# 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

```

Triple Play Service Configuration Commands

```
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
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-ring

Syntax	mc-ring peer <i>ip-address</i> statistics mc-ring peer <i>ip-address</i> [<i>ring sync-tag</i> [<i>detail statistics</i>]] mc-ring peer <i>ip-address</i> ring <i>sync-tag</i> <i>ring-node</i> [<i>ring-node-name</i> [<i>detail statistics</i>]] mc-ring global-statistics
Context	show>redundancy>multi-chassis
Description	This command displays multi-chassis ring information.
Parameters	<i>ip-address</i> — Specifies the address of the multi-chassis peer to display. ring <i>sync-tag</i> — Specifies a synchronization tag to be displayed that was used while synchronizing this port with the multi-chassis peer. node <i>ring-node-name</i> — Specifies a ring-node name. global-statistics — Displays global statistics for the multi-chassis ring. detail — Displays detailed peer information for the multi-chassis ring.
Output	Show mc-ring peer ip-address ring Output — 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	noPeer — The peer has no corresponding ring configured. connected — The inband control connection with the peer is operational. broken — The inband control connection with the peer has timed out. conflict — 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. testingRing — The inband control connection with the peer is being set up. Waiting for result. waitingForPeer — Verifying if this ring is configured on the peer. configErr — The ring is administratively up, but a configuration error prevents it from operating properly. halfBroken — The inband control connection indicates that the ring is broken in one direction (towards the peer). 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
```

Triple Play Service Configuration Commands

```
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 signalling messages received from the peer.
Transmitted	Indicates the number of valid MC-Ring signalling 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.

Label	Description (Continued)
Keepalive	Displays the number of valid MC-Ring control packets of type 'keepalive' were received from the peer.
<pre>*A:ALA-48>show>redundancy>multi-chassis# mc-ring peer 192.251.10.104 statistics ===== MC Ring statistics for peer 192.251.10.104 =====</pre>	
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
<pre>===== *A:ALA-48>show>redundancy>multi-chassis#</pre>	

Show mc-ring ring-node Command Output

Label	Description
Oper State	Displays the state of the connection verification (both local and remote). 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 Command Output

Label	Description
Rx	Displays the number of MC-ring signalling packets were received by this system.
Rx Too Short	Displays the number of MC-ring signalling packets were received by this system that were too short.
Rx Wrong Authentication	Displays the number of MC-ring signalling packets were received by this system with invalid authentication.
Rx Invalid TLV	Displays the number of MC-ring signalling packets were received by this system with invalid TLV.
Rx Incomplete	Displays the number of MC-ring signalling packets were received by this system that were incomplete.
Rx Unknown Type	Displays the number of MC-ring signalling packets were received by this system that were of unknown type.
Rx Unknown Peer	Displays the number of MC-ring signalling packets were received by this system that were related to an unknown peer.
Rx Unknown Ring	Displays the number of MC-ring signalling packets were received by this system that were related to an unknown ring.
Rx Unknown Ring Node	Displays the number of MC-ring signalling packets were received by this system that were related to an unknown ring node.
Tx	Displays the number of MC-ring signalling packets were transmitted by this system.
Tx No Buffer	Displays the number of MC-ring signalling packets could not be transmitted by this system due to a lack of packet buffers.
Tx Transmission Failed	Displays the number of MC-ring signalling packets could not be transmitted by this system due to a transmission failure.
Tx Unknown Destination	Displays the number of MC-ring 'unknown destination' signalling 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
```

Triple Play Service Configuration Commands

```
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

Syntax	lease-state [wholesaler <i>service-id</i>] [sap <i>sap-id</i> sdp <i>sdp-id:vc-id</i> interface <i>interface-name</i> ip-address <i>ip-address[/mask]</i> chaddr <i>ieee-address</i> mac <i>ieee-address</i> {[port <i>port-id</i>] [no-inter-dest-id inter-dest-id <i>inter-dest-id</i>]}] [detail]
Context	show>service>id>dhcp
Description	This command displays DHCP lease state information. Note that the wholesaler <i>service-id</i> parameter is applicable only in the VPRN context.
Parameters	wholesaler <i>service-id</i> — The service ID of the wholesaler. When specified in this context, SAP, SDP, interface, IP address and MAC parameters are ignored. Values service-id: 1 — 214748364 svc-name: A string up to 64 characters in length. sap <i>sap-id</i> — Specifies the physical port identifier portion of the SAP definition. See Common Service Commands on page 1020 for <i>sap-id</i> command syntax. sdp <i>sdp-id</i> — The SDP identifier. Values 1 — 17407 <i>vc-id</i> — 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 <i>ip-address[/mask]</i> — Shows information for the specified IP address and mask. port <i>port-id</i> — 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 <i>ip-int-name</i> — 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 LifeTime	Lease Origin	MC 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 LifeTime	Lease Origin	MC 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

Syntax	statistics [[sap sap-id]][sdp sdp-id:vc-id]][interface interface-name]
Context	show>service>id>dhcp
Description	This command displays DHCP relay statistics.
Parameters	interface ip-int-name — Display DHCP statistics on the specified interface. interface interface-name — Displays DHCP statistics for the specified interface name. sap sap-id — Displays DHCP statistics for the specified SAP. See Common Service Commands on page 1020 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

- Syntax** `summary [interface interface-name | saps]`
- Context** `show>service>id>dhcp`
- Description** This command displays DHCP configuration summary information.
- Parameters** `interface interface-name` — Displays summary information for the specified existing interface.
`sap` — 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#
```

```
VRPN:
```

```
*A:ALA-49>show>service# id 1 dhcp summary
```

```
=====
DHCP Summary, service 1
```

Interface Name SapId/Sdp	Arp Populate	Used/ Provided	Info Option	Admin State
SpokeSDP	No	0/0	Keep	Down
sdp:spoke-3:4		0/0		
test	No	0/0	Keep	Down
sap:9/1/4:50/5		0/0		
to-cel	No	0/0	Keep	Up
sap:1/1/10:1		0/0		

```
-----
Interfaces: 3
=====
```

```
*A:ALA-49>show>service#
```

```
*A:ALA-48# show service id 88 dhcp summary saps
```

```
=====
DHCP Summary, service 88
```

Interface Name SapId/Sdp	Arp Populate	Used/ Provided	Info Option	Admin 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 SapId/Sdp	Arp Populate	Used/ Provided	Info Option	Admin State
SpokeTerm	No	0/0	Keep	Up
sdp:spoke-3:3		0/0		

```
-----
Interfaces: 1
=====
```

*A:ALA-48#

statistics

Syntax	statistics [interface <i>ip-int-name</i> <i>ip-address</i>]
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 <i>ip-int-name</i> <i>ip-address</i> — 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
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

Label	Description (Continued)
DHCP RELEASEs spoofed	Indicates the total number of DHCP release messages spoofed by the DHCP relay agent to the DHCP server.
DHCP FORCERENEWs 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 FORCERENEWs Spoofed : 0 *
=====
A:SUB-Dut-A#
```

summary

Syntax	summary
Context	show>router>dhcp
Description	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#
```

IGMP Snooping Show Commands

igmp-snooping

Syntax	igmp-snooping
Context	show>service>id>
Description	This command enables the context to display IGMP snooping information.

all

Syntax	all
Context	show>service>id>igmp-snooping
Description	Displays detailed information for all aspects of IGMP snooping on the VPLS service.
Output	Show All Service-ID — 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-48>show>service>id>igmp-snooping>snooping# all
=====
IGMP Snooping info for service 750
=====
IGMP Snooping Base info
-----
Admin State : Up
Querier      : No querier found
-----
Sap/Sdp      Oper    MRtr  Send    Max Num  Num
Id           State   Port  Queries Groups   Groups
-----
sap:1/1/7:0   Down    No    Disabled No Limit  0
sdp:1:22      Down    No    Disabled No Limit  0
sdp:8:750     Down    No    Disabled No Limit  0
-----
IGMP Snooping Querier info
-----
No querier found for this service.
-----
IGMP Snooping Multicast Routers
-----
MRouter      Sap/Sdp Id           Up Time           Expires           Version
-----
Number of mrouters: 0
-----
IGMP Snooping Proxy-reporting DB
-----
Group Address  Mode    Type    Up Time           Expires           Num Src
-----
Number of groups: 0
-----
IGMP Snooping SAP 1/1/7:0 Port-DB
-----
Group Address  Mode    Type    Up Time           Expires           Num Src
-----
Number of groups: 0
-----
IGMP Snooping SDP 1:22 Port-DB
-----
Group Address  Mode    Type    Up Time           Expires           Num Src
-----
Number of groups: 0
-----
IGMP Snooping SDP 8:750 Port-DB
-----
Group Address  Mode    Type    Up Time           Expires           Num Src
-----
Number of groups: 0
-----
IGMP Snooping Static Source Groups
-----
IGMP Snooping Statistics
-----
Message Type      Received      Transmitted      Forwarded
-----
General Queries    0             0                 0
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
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

```

```
=====
*A:ALA-48>show>service>id>snooping#

```

mrouter

Syntax	mrouter [detail]
Context	show>service>id>igmp-snooping
Description	Displays all multicast routers.
Parameters	detail — Displays detailed information.
Output	Show igmp-snooping mrouter — The following table describes the show igmp-snooping mrouter 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 mrouter
```

```
=====
IGMP Snooping Multicast Routers for service 700
=====
```

```

MRouter      Sap/Sdp Id      Up Time      Expires      Version
-----

```

Triple Play Service Configuration Commands

```
Number of mrouter: 0
```

```
*A:ALA-48#
```

mvr

Syntax	mvr
Context	show>service>id>igmp-snooping
Description	Displays Multicast VPLS Registration (MVR) information.
Output	Show igmp-snooping mvr — 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

Syntax

```

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

```

Context show>service>id>igmp-snooping

Description This command displays information on the IGMP snooping port database for the VPLS service.

Parameters **group** *grp-ip-address* — Displays the IGMP snooping port database for a specific multicast group address.

sap *sap-id* — Displays the IGMP snooping port database for a specific SAP. See [Common Service Commands on page 1020](#) for *sap-id* command syntax.

sdp *sdp-id* — 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.

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

group *grp-address* — Displays IGMP snooping statistics matching the specified group address.

source *ip-address* — Displays IGMP snooping statistics matching one particular source within the multicast group.

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      Up Time      Expires      Num Sources
-----
225.0.0.1          include   dynamic   0d 00:04:44   0s           2
-----
Number of groups: 1
=====

```

Triple Play Service Configuration Commands

```
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         0d 00:04:57  20s     dynamic   Fwd
1.1.1.2         0d 00:04:57  20s     dynamic   Fwd
-----
Number of groups: 1
=====
A:ALA-1>show>service>id>snooping#
```

proxy-db

Syntax	proxy-db [detail] proxy-db group grp-address
Context	show>service>id>igmp-snooping
Description	Displays information on the IGMP snooping proxy reporting database for the VPLS service.
Parameters	group grp-ip-address — 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         0d 00:05:54
```

```

1.1.1.2          0d 00:05:54
-----
Number of groups: 1
=====
A:ALA-1>show>service>id>snooping#

```

querier

Syntax	querier
Context	show>service>id>igmp-snooping
Description	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

Syntax	static [sap <i>sap-id</i> sdp <i>sdp-id:vc-id</i>]
Context	show>service>id>igmp-snooping
Description	Displays information on static IGMP snooping source groups for the VPLS service.
Parameters	sap <i>sap-id</i> — Displays static IGMP snooping source groups for a specific SAP. See Common Service Commands on page 1020 for <i>sap-id</i> command syntax. sdp <i>sdp-id</i> — Displays the IGMP snooping source groups for a specific spoke or mesh SDP. Values 1 — 17407 <i>vc-id</i> — 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

Triple Play Service Configuration Commands

```
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 1020](#) 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
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

- Syntax** **mc-ecmp-balance [detail]**
- Context** show>router>pim
- Description** This command displays multicast balance information.
- Parameters** **detail** — 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

- Syntax** **mcast-management**
- Context** show
- Description** This command shows multicast path management related information.

bandwidth-policy

Syntax	bandwidth-policy <i>policy-name</i> [detail]
Context	show>mcast-management
Description	This command displays multicast path management bandwidth policy information.
Parameters	<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

Syntax	channel [router <i>router-instance</i> vpls <i>service-id</i>] [<i>mda slot[/mda]</i>] [group <i>ip-address</i> [source <i>ip-address</i>]] [path <i>path-type</i>] [detail]
Context	show>mcast-management

Description This command displays multicast path management channel related information.

Parameters **vpls** *service-id* — Specifies an existing VPLS service ID.

Values service-id: 1 — 214748364
svc-name: A string up to 64 characters in length.

ip-address — ipv4-address a.b.c.d

path-type — Specifies the path type.

Values primary, secondary, ancillary

Output

```
*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#
```

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#
```


Clear Commands

id

Syntax	id <i>service-id</i>
Context	clear>service clear>service>statistics
Description	This command clears the identification for a specific service.
Parameters	<i>service-id</i> — The ID that uniquely identifies a service.
Values	service-id: 1 — 214748364 svc-name: A string up to 64 characters in length.

arp-host

Syntax	arp-host arp-host { mac <i>ieee-address</i> sap <i>sap-id</i> ip-address <i>ip-address</i> [/ <i>mask</i>] } arp-host [port <i>port-id</i>] [inter-dest-id <i>intermediate-destination-id</i> no-inter-dest-id] arp-host statistics [sap <i>sap-id</i> interface <i>interface-name</i>]
Context	clear>service>id
Description	This command clears ARP host data.

authentication

Syntax	authentication
Context	clear>service>id
Description	This command enters the context to clear session authentication information.

msap

Syntax	msap <i>msap-id</i>
Context	clear>service>id
Description	This command clears Managed SAP information.
Parameters	<i>msap-id</i> — Specifies a Managed SAP ID.
Values	dot1q [port-id lag-id]:qtag1 qinq [port-id lag-id]:qtag1.qtag2

qtag1	0 — 4094
qtag2	0 — 4094

msap-policy

Syntax	msap-policy <i>msap-policy-name</i>
Context	clear>service>id
Description	This command clears Managed SAPs created by the Managed SAP policy.
Parameters	<i>msap-policy-name</i> — Specifies an existing MSAP policy.

statistics

Syntax	statistics
Context	clear>service>id>authentication
Description	This command clears session authentication statistics for this service.

statistics

Syntax	statistics
Context	clear>service
Description	This command clears the statistics for a service.

subscriber

Syntax	subscriber <i>sub-ident-string</i>
Context	clear>service>statistics
Description	This command clears the statistics for a particular subscriber.
Parameters	<i>sub-ident-string</i> — Clears statistics for the specified subscriber identification string.

fdb

Syntax	fdb { all mac <i>ieee-address</i> sap <i>sap-id</i>] mesh-sdp <i>sdp-id[:vc-id]</i> spoke-sdp <i>sdp-id:vc-id</i> }																
Context	clear>service>id																
Description	This command clears FDB entries for the service.																
Parameters	all — Clears all FDB entries. mac <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 Common Service Commands on page 1020 for <i>sap-id</i> command syntax. mesh-sdp — Clears only service FDB entries associated with the specified mesh SDP ID. For a mesh SDP, the VC ID is optional. spoke-sdp — 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>Values</td><td>sdp-id[:vc-id]</td><td><i>sdp-id</i></td><td>1 — 17407</td></tr><tr><td></td><td></td><td>vc-id</td><td>1 — 4294967295</td></tr><tr><td></td><td>sdp-id:vc-id</td><td>sdp-id</td><td>1 — 17407</td></tr><tr><td></td><td></td><td>vc-id</td><td>1 — 4294967295</td></tr></table>	Values	sdp-id[:vc-id]	<i>sdp-id</i>	1 — 17407			vc-id	1 — 4294967295		sdp-id:vc-id	sdp-id	1 — 17407			vc-id	1 — 4294967295
Values	sdp-id[:vc-id]	<i>sdp-id</i>	1 — 17407														
		vc-id	1 — 4294967295														
	sdp-id:vc-id	sdp-id	1 — 17407														
		vc-id	1 — 4294967295														

mesh-sdp

Syntax	mesh-sdp <i>sdp-id[:vc-id]</i> ingress-vc-label						
Context	clear>service>id						
Description	Clears and resets the mesh SDP bindings for the service.						
Parameters	<i>sdp-id</i> — The mesh SDP ID to be reset. <table> <tr> <td>Values</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>Default</td><td>All VC IDs on the SDP ID.</td></tr> <tr> <td>Values</td><td>1 — 4294967295</td></tr> </table>	Values	1 — 17407	Default	All VC IDs on the SDP ID.	Values	1 — 4294967295
Values	1 — 17407						
Default	All VC IDs on the SDP ID.						
Values	1 — 4294967295						

spoke-sdp

Syntax	spoke-sdp <i>sdp-id:vc-id</i> ingress-vc-label
Context	clear>service>id
Description	Clears and resets the spoke SDP bindings for the service.
Parameters	<i>sdp-id</i> — The spoke SDP ID to be reset. Values 1 — 17407 <i>vc-id</i> — The virtual circuit ID on the SDP ID to be reset. Values 1 — 4294967295

sap

Syntax	sap <i>sap-id</i> {all counters stp}
Context	clear>service>statistics
Description	Clears SAP statistics for a SAP.
Parameters	<i>sap-id</i> — Specifies the physical port identifier portion of the SAP definition. See Common Service Commands on page 1020 for <i>sap-id</i> command syntax. all — Clears all SAP queue statistics and STP statistics. counters — Clears all queue statistics associated with the SAP. stp — Clears all STP statistics associated with the SAP.

sdp

Syntax	sdp <i>sdp-id</i> [keep-alive]
Context	clear>service>statistics
Description	Clears keepalive statistics associated with the SDP ID.
Parameters	<i>sdp-id</i> — The SDP ID for which to clear statistics. Values 1 — 17407 keep-alive — Clears the keepalive history.

counters

Syntax	counters
Context	clear>service>statistics>id
Description	Clears all traffic queue counters associated with the service ID.

sap

Syntax	sap sap-id {all counters stp}
Context	clear>service>statistics>id
Description	Clears statistics for the SAP bound to the service.
Parameters	<i>sap-id</i> — Specifies the SAP ID for which to clear statistics. See Common Service Commands on page 1020 for <i>sap-id</i> command syntax. all — Clears all queue statistics and STP statistics associated with the SAP. counters — Clears all queue statistics associated with the SAP. stp — Clears all STP statistics associated with the SAP.

spoke-sdp

Syntax	spoke-sdp sdp-id[:vc-id] {all counters stp}
Context	clear>service>statistics>id
Description	Clears statistics for the spoke SDP bound to the service.
Parameters	<i>sdp-id</i> — The spoke SDP ID for which to clear statistics. Values 1 — 17407 <i>vc-id</i> — The virtual circuit ID on the SDP ID to be reset. Values 1 — 4294967295 all — Clears all queue statistics and STP statistics associated with the SDP. counters — Clears all queue statistics associated with the SDP. stp — Clears all STP statistics associated with the SDP.

stp

Syntax	stp
Context	clear>service>statistics>id
Description	Clears all spanning tree statistics for the service ID.

detected-protocols

Syntax	detected-protocols {all sap <i>sap-id</i> spoke-sdp <i>sdp-id</i>[:<i>vc-id</i>]}
Context	clear>service>id>stp
Description	RSTP automatically falls back to STP mode when it receives an STP BPDU. The clear detected-protocols command forces the system to revert to the default RSTP mode on the SAP or spoke SDP.
Parameters	all — Clears all detected protocol information. <i>sap-id</i> — Clears the specified lease state SAP information. See Common Service Commands on page 1020 for <i>sap-id</i> command syntax. <i>sdp-id</i> — The SDP ID to be cleared. Values 1 — 17407 <i>vc-id</i> — The virtual circuit ID on the SDP ID to be cleared. Values 1 — 4294967295

lease-state

Syntax	lease-state[no-dhcp-release] lease-state [port <i>port-id</i>] [inter-dest-id <i>intermediate-destination-id</i>] [no-dhcp-release] lease-state [port <i>port-id</i>] no-inter-dest-id [no-dhcp-release] lease-state ip-address <i>ip-address</i>[/<i>mask</i>] [no-dhcp-release] lease-state mac <i>ieee-address</i> [no-dhcp-release] lease-state sap <i>sap-id</i> [no-dhcp-release] lease-state sdp <i>sdp-id</i>[:<i>vc-id</i>] [no-dhcp-release]
Context	clear>service>id>dhcp
Description	Clears DHCP lease state information for this service.
Parameters	no-dhcp-release — Specifies that the node will clear the state without sending the DHCP release message. port <i>port-id</i> — 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.

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 1020](#) 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

Syntax	statistics [sap <i>sap-id</i> sdp [<i>sdp-id</i> [: <i>vc-id</i>]]]
Context	clear>service>id>dhcp
Description	Clears DHCP statistics for this service.

port-db

Syntax	port-db { sap <i>sap-id</i> sdp <i>sdp-id</i> [: <i>vc-id</i>]} [group <i>grp-address</i> [source <i>ip-address</i>]]
Context	clear>service>id>igmp-snooping
Description	Clears the information on the IGMP snooping port database for the VPLS service.
Parameters	sap <i>sap-id</i> — Clears IGMP snooping statistics matching the specified SAP ID and optional encapsulation value. See Common Service Commands on page 1020 for <i>sap-id</i> command syntax. sdp <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. Values 1 — 17407 <i>vc-id</i> — The virtual circuit ID on the SDP ID for which to clear information. Default For mesh SDPs only, all VC IDs. Values 1 — 4294967295 group <i>grp-address</i> — Clears IGMP snooping statistics matching the specified group address.

source *ip-address* — Clears IGMP snooping statistics matching one particular source within the multicast group.

querier

Syntax	querier
Context	clear>service>id>igmp-snooping
Description	Clears the information on the IGMP snooping queriers for the VPLS service.

statistics

Syntax	statistics {all sap <i>sap-id</i> sdp <i>sdp-id:vc-id</i>}
Context	clear>service>id>igmp-snooping
Description	Clears IGMP snooping statistics for the VPLS service.
Parameters	sap <i>sap-id</i> — Clears IGMP snooping statistics matching the specified SAP ID and optional encapsulation value. See Common Service Commands on page 1020 for <i>sap-id</i> command syntax. sdp <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. Values 1 — 17407 <i>vc-id</i> — The virtual circuit ID on the SDP ID for which to clear statistics. Default For mesh SDPs only, all VC IDs Values 1 — 4294967295

mfib

Syntax	mfib
Context	clear>service>id>
Description	Enter the context to clear multicast FIB info for the VPLS service.

statistics

Syntax	statistics {all group <i>grp-address</i>}
Context	clear>service>id>mfib
Description	Clears multicast FIB statistics for the VPLS service.
Parameters	<i>grp-address</i> — Specifies an IGMP multicast group address that receives data on an interface.

mld-snooping

Syntax	mld-snooping
Context	clear>service>id
Description	This command enables the context to clear MLD snooping-related data.

port-db

Syntax	port-db sap <i>sap-id</i> [group <i>grp-ipv6-address</i>] port-db sap <i>sap-id</i> group <i>grp-ipv6-address</i> source <i>src-ipv6-address</i> port-db sdp <i>sdp-id:vc-id</i> [group <i>grp-ipv6-address</i>] port-db sdp <i>sdp-id:vc-id</i> group <i>grp-ipv6-address</i> source <i>src-ipv6-address</i>
Context	clear>service>id>mld-snooping
Description	This command clears MLD snooping port-db group data.

querier

Syntax	querier
Context	clear>service>id>mld-snooping
Description	This command clears MLD snooping querier information.

statistics

Syntax	statistics all statistics sap <i>sap-id</i> statistics sdp <i>sdp-id:vc-id</i>
Context	clear>service>id>mld-snooping
Description	This command clears MLD snooping statistics.
Parameters	all — Clears all MLD snooping statistics sap <i>sap-id</i> — Clears all MLD snooping SAP statistics. See Common Service Commands on page 1020 for <i>sap-id</i> command syntax. sdp <i>sdp-id:vc-id</i> — Clears all MLD snooping SDP statistics. See Common Service Commands on page 1020 for <i>sap-id</i> command syntax.

arp

Syntax	arp { all <i>ip-address</i> } arp interface [<i>ip-int-name</i> <i>ip-address</i>]
Context	clear>router
Description	This command clears all or specific ARP entries. The scope of ARP cache entries cleared depends on the command line option(s) specified.
Parameters	all — Clears all ARP cache entries. <i>ip-address</i> — Clears the ARP cache entry for the specified IP address. interface <i>ip-int-name</i> — Clears all ARP cache entries for the IP interface with the specified name. interface <i>ip-addr</i> — Clears all ARP cache entries for the specified IP interface with the specified IP address.

dhcp

Syntax	dhcp
Context	clear>router
Description	This command enables the context to clear and reset DHCP entities.

statistics

Syntax	statistics [interface <i>ip-int-name</i> <i>ip-address</i>]
Context	clear>router>dhcp
Description	Clears DHCP statistics.

Debug Commands

arp-host

Syntax	[no] arp-host
Context	debug>service>id
Description	This command enables and ARP host debugging. The no form of the command disables ARP host debugging

mld-snooping

Syntax	[no] mld-snooping
Context	debug>service>id
Description	This command enables and configures MLD-snooping debugging. The no form of the command disables MLD-snooping debugging

detail-level

Syntax	detail-level {low medium high} no detail-level
Context	debug>service>id>mld
Description	This command enables and configures the MLD tracing detail level. The no form of the command disables the MLD tracing detail level.

mac

Syntax	[no] mac <i>ieee-address</i>
Context	debug>service>id>mld
Description	This command shows MLD packets for the specified MAC address. The no form of the command disables disables the MAC debugging.

mode

Syntax	mode {dropped-only ingr-and-dropped egr-ingr-and-dropped} no mode
Context	debug>service>id>mld
Description	This command enables and configures the MLD tracing mode. The no form of the command disables the configures the MLD tracing mode.

sap

Syntax	[no] sap <i>sap-id</i>
Context	debug>service>id>mld
Description	This command shows MLD packets for a specific SAP. The no form of the command disables the debugging for the SAP.

sdp

Syntax	[no] sdp <i>sdp-id:vc-id</i>
Context	debug>service>id>mld
Description	This command shows MLD packets for a specific SDP. The no 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 302](#)
- [DHCP Features on page 304](#)
 - [DHCP Relay on page 304](#)
 - [Subscriber Identification Using Option 82 Field on page 304](#)
 - [DHCP Snooping on page 307](#)
 - [DHCP Lease State Table on page 309](#)
 - [DHCPv6 on page 315](#)
- [Configuring DHCP with CLI on page 321](#)

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 16](#).




Figure 16: 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 304](#)
 - [Subscriber Identification Using Option 82 Field on page 304](#)
 - [DHCP Snooping on page 307](#)
 - [DHCP Lease State Table on page 309](#)
-

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 IES and VPRN services. On VPRN interfaces, however they will only forward to DHCP servers that also participate in that VPRN.

While DHCP relay is not implemented in a VPLS, it is still possible to insert or modify Option 82 information

In a Routed CO environment, the subscriber interface's group interface DHCP relay is stateful.

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 7750 SR and can be used separately or together.

Inserting Option 82 information is supported independently of DHCP relay. However in a VPLS (when the 7750 SR 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 7750 SR, it can take following values:

- *sap-id*: The SAP index (only under a IES or VPRN service)
- *ifindex*: The index of the IP interface (only under a IES or VPRN service)
- *ascii-tuple*: An ASCII-encoded concatenated tuple, consisting of [*system-name*|*service-id*|*interface-name*] (for VPRN or 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.

Note that for VPRN the *ifindex* is unique only within a VRF. The DHCP relay function automatically prepends the VRF ID to the *ifindex* before relaying a DHCP Request.

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 7750 SR. 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 7750 SR 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 or VPRN. 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 7750 SR 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 304](#));
 - 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 466](#).

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 309](#)
- to perform Enhanced Subscriber Management, see section [Triple Play Enhanced Subscriber Management on page 603](#)

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 and VPRN IP 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 and VPRN services, if ARP populate is configured, no statics ARPs are allowed. For IES and VPRN services, if ARP populate is not configured, then statics 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 or VPRN IP interfaces terminated on a 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 and VPRN SAP IP 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. 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 the Routed CO model, see [Routed Central Office \(CO\) on page 660](#).
- To support enhanced Subscriber management, see section [RADIUS Authentication of Subscriber Sessions on page 605](#).

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 WiiMAX 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.

Typically, the DHCP server can be configured to perform as follows:

- When a user-db is specified, a host lookup is performed in the local-user-db and the local-user-db host defines how to get the IP-address or pool to get the IP-address from:
 - A fixed IP-address.
 - A pool — A free address of any subnet in that pool is offered.
 - use-pool-from-client — The pool name is taken from Option 82 vendor-specific suboption. (If not present, proceed to the next bullet.)
 - use-gi-address — The gi-address is used to find a matching subnet.
- When no user-db is specified or when the user-db does not return pool or IP-address information, then a pool or subnet is searched in following manner:
 - use-pool-from-client — If set, the pool name is taken from Option 82 vendor-specific suboption. If the pool name is not present, go to next bullet.
 - use-gi-address — If set, the gi-address is used to find a matching subnet. If no subnet is found, then no address is offered.
- If no IP-address, subnet, or pool is found in the first 2 bullets (above), no IP-address is offered to the client. If a subnet is found, an available address from the subnet is offered to the client. If a pool is found, an available address from any subnet of the pool is offered to the client.

NOTE: When the **no use-gi**, **no use-pool-from-client** and **no user-db** commands are specified, the DHCP server does not act.

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, 7750 SR-12, and 7710 SR-Series 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.
- Business services — A server can be defined in a VPRN service and associated with different interfaces. Locally attached hosts can get an address directly from the servers. Routed hosts will receive addresses through a relay point in the customer's network.
- 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 including overlapped private addressing in the form of VPRNs in the same SR-series router.

DHCP servers are configurable and can be used in the bridged CO, routed CO, VPRN wholesaler, 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 315](#)
 - [DHCPv6 Prefix Options on page 315](#)
 - [Neighbor Resolution via DHCPv6 Relay on page 315](#)
 - [DHCPv6 Lease Persistency on page 316](#)
 - [Local Proxy Neighbor Discovery on page 316](#)
-

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 re-populate 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.

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.

Configuring DHCP with CLI

This section provides information to configure DHCP using the command line interface.

- [Common Configuration Tasks on page 322](#)
 - [Enabling DHCP Snooping on page 322](#)
 - [Configuring Option 82 Handling on page 324](#)
 - [Enabling DHCP Relay on page 325](#)
 - [Configuring Local User Database Parameters on page 326](#)

Common Configuration Tasks

Topics in this section are:

- [Enabling DHCP Snooping on page 322](#)
 - [Configuring Option 82 Handling on page 324](#)
 - [Enabling DHCP Relay on page 325](#)
 - [Configuring Local User Database Parameters on page 326](#)
-

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 and VPRN IP 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 IES 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 626](#).

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:

- **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"
        pppoe
            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 profl"
                    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

```

```

        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"
        pppoe
            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 profl"
                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 PPPoE for authentication

```

A:pe5>config>service>vprn#
-----
    subscriber-interface "tomylinux" create
    address 10.2.2.2/16
    group-interface "grp_pppoe3" create
    pppoe
    pap-chap-user-db "pppoe"
    exit
exit
-----
A:pe5>config>service>vprn#

```

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 333](#)
- [VPLS DHCP Commands on page 335](#)
- [IES DHCP Commands on page 336](#)
- [VPRN DHCP Commands on page 339](#)
- [Local User Database Commands on page 342](#)
- [Show Commands on page 345](#)
- [Clear Commands on page 346](#)
- [DHCP Debug Commands on page 347](#)

Router DHCP Commands

```

config
  — router
    — dhcp
      — local-dhcp-server server-name [create]
      — no local-dhcp-server server-name
        — description description-string
        — no description
        — failover
          — 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-name
          — no peer ip-address
          — [no] shutdown
          — no startup-wait-time [min minutes] [sec seconds]
          — startup-wait-time [min minutes] [sec seconds]
      — [no] force-renew
      — pool pool-name [create]
      — no pool pool-name
        — description description-string
        — no description
        — 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
        — 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**
- **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**] **exclude-addresses** *start-ip-address [end-ip-address]*
 - **maximum-declined** *maximum-declined*
 - **no maximum-declined**
 - **minimum-free** *minimum-free*
 - **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**
- [**no**] **use-gi-address**
- [**no**] **use-pool-from-client**
- **user-db** *local-user-db-name*
- **no user-db**

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]
        — 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
    — mesh-sdp sdp-id[:vc-id] [vc-type {ether | vlan}]
      — dhcp
        — description description-string
        — no description
        — [no] snoop
    — 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] pool-name
              — [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] [radius-override]
            — no lease-time
            — [no] shutdown
          — 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]
            — 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
— [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
— lease-populate [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**
- **server** *server1* [*server2*...(up to 8 max)]
- **no server**
- [no] **shutdown**
- [no] **trusted**

VPRN DHCP Commands

```

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 | 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
          — 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
          — [no] relay-plain-bootp
          — server server1 [server2...(up to 8 max)]
          — no server
          — [no] shutdown
          — [no] trusted
          — [no] use-arp

```

VPRN DHCP Subscriber Interface Commands

```

config
  — service
    — vprn
      — subscriber-interface ip-int-name [fwd-service service-id fwd-subscriber-inter-
        face 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] [radius-override]
              — no lease-time
              — [no] shutdown
            — server server1 [server2...(up to 8 max)]
            — no server
            — [no] shutdown
      — [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] [radius-override]
  — no lease-time
  — [no] shutdown
— server server1 [server2...(up to 8 max)]
— no server
— [no] shutdown
— [no] trusted

```

Local User Database 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
— dhcp
— mask {[prefix-string prefix-string | prefix-length prefix-length] [suffix-
string suffix-string | suffix-length suffix-length]}
— no mask
— host host-name [create]
— no host host-name
— address gi-address
— address ip-address
— address pool pool-name
— no address
— authentication-policy policy-name
— no authentication-policy
— domain-name domain-name
— no domain-name
— host-identification
— circuit-id string ascii-string
— circuit-id hex hex-string
— no circuit-id
— mac ieee-address
— no mac
— option60 hex-string
— no option60
— remote-id remote-id
— 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
— 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
— msap-defaults
— group-interface ip-int-name
— no group-interface

```

- **policy** *msap-policy-name*
- **no policy**
- **service** *service-id*
- **no service**
- **[no] retail-service** *service-id*
- **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)]
 - **no dns-server**
 - **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**
 - **subnet-mask** *ip-address*
 - **no subnet-mask**
- **server** *ip-address*
- **no server**
- **[no] shutdown**
- **mask** *dhcp-match-type* {[**prefix-string** *prefix-string* | **prefix-length** *prefix-length*] [**suffix-string** *suffix-string* | **suffix-length** *suffix-length*]}
- **no mask** *type dhcp-match-type*
- **match-list** *dhcp-match-type-1* [*dhcp-match-type-2...*(up to 4 max)]
- **no match-list**
- **pppoe**
 - **mask** {[**prefix-string** *prefix-string* | **prefix-length** *prefix-length*] [**suffix-string** *suffix-string* | **suffix-length** *suffix-length*]}
 - **no mask**
 - **host** *host-name* [**create**]
 - **no host** *host-name*
 - **address** *gi-address*
 - **address** *ip-address*
 - **address** **pool** *pool-name*
 - **no address**
 - **authentication-policy** *policy-name*
 - **no authentication-policy**
 - **host-identification**
 - **circuit-id** **string** *ascii-string*
 - **circuit-id** **hex** *hex-string*
 - **no circuit-id**

```

— mac ieee-address
— no mac
— remote-id remote-id
— no remote-id
— service-name service-name
— no service-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
   — 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
— l2tp
   — group tunnel-group-name
   — no group
— msap-defaults
   — group-interface ip-int-name
   — 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 [ip-address...(up to 4 max)]
   — no dns-server
   — netbios-name-server ip-address [ip-address...(up to 4 max)]
   — no netbios-name-server
— pado-delay deci-seconds
— no pado-delay
— password {ignore | chap string | pap string} [hash|hash2]
— no password
— [no] shutdown
— match-list pppoe-match-type-1 [pppoe-match-type-2...(up to 3 max)]
— no match-list
— [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] pppoe-host pppoe-host-name [detail]
        — leases ip-address[/mask] [detail]
        — server-stats
        — subnet-stats ip-address[/mask]
        — subnet-stats pool pool-name
        — summary
      — servers
    — router
      — dhcp6
        — statistics
        — summary
dhcp6show
  — 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
show
  — service
    — id service-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

```

Tools Commands

```
tools
— 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]
— server-stats
— statistics [ip-int-name | ip-address]
— dhcp6
— 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
      — dhcp6 [ip-int-name]
      — no dhcp6
      — [no] local-dhcp-server server-name [lease-address ip-address]
        — 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
      — edit-pppoe-session sap sap-id ip ip-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]
      — eval-lease-state [svc-id service-id] [sap sap-id] [subscriber sub-ident-string] [ip ip-address]
      — local-user-db local-user-db-name
        — pppoe
          — 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

Global Commands

shutdown

Syntax	[no] shutdown
Context	<pre>config>service>ies>if>dhcp config>service>vpls>sap>dhcp config>service>vpls>sap>dhcp>option>vendor config>service>vprn>if>dhcp config>service>vprn>if>dhcp>proxy-server config>subscr-mgmt>loc-user-db config>subscr-mgmt>loc-user-db>dhcp>host config>subscr-mgmt>loc-user-db>dhcp>host>options config>subscr-mgmt>loc-user-db>pppoe>host config>router>dhcp>server>failover</pre>
Description	This command administratively disables an entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics. 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. The no form of this command places the entity into an administratively enabled state.

description

Syntax	description <i>description-string</i> no description
Context	<pre>config>service>vpls>sap>dhcp config>service>ies>if>dhcp config>service>ies>if>ipv6>dhcp6-relay config>service>vprn>if>dhcp config>router>dhcp>server config>router>dhcp>server>pool config>subscr-mgmt>loc-user-db</pre>
Description	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 no form of this command removes the string from the configuration.
Default	No description associated with the configuration context.

Triple Play DHCP Configuration Commands

Parameters *description-string* — 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.

Router DHCP Configuration Commands

local-dhcp-server

Syntax	local-dhcp-server <i>server-name</i> [create] no local-dhcp-server <i>server-name</i>
Context	config>router>dhcp config>service>vprn>dhcp
Description	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.
Default	none
Parameters	<i>server-name</i> — Specifies the name of local DHCP server.

failover

Syntax	failover
Context	config>router>dhcp>server
Description	This command enables the context to configure failover parameters.

maximum-client-lead-time

Syntax	maximum-client-lead-time [hrs <i>hours</i>] [min <i>minutes</i>] [sec <i>seconds</i>] no maximum-client-lead-time
Context	config>router>dhcp>server>failover
Description	This command specifies the maximum client lead time (MCLT). The MCLT is the maximum amount of time that one server can extend a lease for a client's binding beyond the time known by the partner server.
Parameters	hrs <i>hours</i> — Specifies the hour parameter of the MCLT. Values 1 — 23 min <i>minutes</i> — Specifies the minute parameter of the MCLT. Values 1 — 59 sec <i>seconds</i> — Specifies the seconds parameter of the MCLT. Values 1 — 59

partner-down-delay

Syntax	partner-down-delay [<i>hrs hours</i>] [<i>min minutes</i>] [<i>sec seconds</i>] no partner-down-delay
Context	config>router>dhcp>server>failover
Description	This command specifies the minimum safe-time after the beginning of a communications-interrupted state. After the expiry of this time, the server enters a partner-down state.
Parameters	hrs <i>hours</i> — Specifies the hour parameter of the partner down delay feature. Values 1 — 23 min <i>minutes</i> — Specifies the minute parameter of the partner down delay feature. Values 1 — 59 sec <i>seconds</i> — Specifies the seconds parameter of the partner down delay feature. Values 1 — 59

peer

Syntax	peer <i>ip-address</i> tag <i>sync-tag-name</i> no peer <i>ip-address</i>
Context	config>router>dhcp>server>failover
Description	This command creates a synchronization tag.
Parameters	sync-tag <i>sync-tag</i> — Specifies a synchronization tag to be used while synchronizing with the multi-chassis peer.

startup-wait-time

Syntax	[no] startup-wait-time [<i>min minutes</i>] [<i>sec seconds</i>]
Context	config>router>dhcp>server>failover
Description	This command specifies the startup wait time. The startup wait time is the time that one server attempts to contact the partner server. During this time, the server is unresponsive to DHCP client requests.
Parameters	min <i>minutes</i> — Specifies the minute parameter of the startup wait time feature. Values 1 — 59 sec <i>seconds</i> — Specifies the seconds parameter of the startup wait time feature. Values 1 — 59

force-renew

Syntax	[no] force-renew
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

pool

Syntax	pool <i>pool-name</i> [create] no pool <i>pool-name</i>
Context	config>router>dhcp>server
Description	This command configures a DHCP address pool on the router.
Default	none
Parameters	<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.

max-lease-time

Syntax	max-lease-time [days <i>days</i>] [hrs <i>hours</i>] [min <i>minutes</i>] [sec <i>seconds</i>] no max-lease-time								
Context	config>router>dhcp>server>pool								
Description	This command configures the maximum lease time. The no form of the command returns the value to the default.								
Default	10 days								
Parameters	<i>time</i> — Specifies the maximum lease time.								
Values	<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								

min-lease-time

Syntax	min-lease-time [days <i>days</i>] [hrs <i>hours</i>] [min <i>minutes</i>] [sec <i>seconds</i>] no min-lease-time
Context	config>router>dhcp>server>pool

Triple Play DHCP Configuration Commands

Description	This command configures the minimum lease time. The no form of the command returns the value to the default.		
Default	10 minutes		
Parameters	<i>time</i> — Specifies the minimum lease time.		
Values	days :	0 — 3650	
	hours	0 — 23	
	minutes:	0 — 59	
	seconds	0 — 59	

offer-time

Syntax	offer-time [min <i>minutes</i>] [sec <i>seconds</i>] 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	<i>time</i> — Specifies the offer time.		
Values	minutes:	0 — 10	
	seconds	0 — 59	

msap-defaults

Syntax	msap-default		
Context	config>sub-mgmt>lu-db>dhcp>hos config>sub-mgmt>lu-db>pppoe>host		
Description	This command configures MSAP authentication defaults.		

group-interface

Syntax	group-interface <i>ip-int-name</i> no group-interface		
Context	config>sub-mgmt>lu-db>dhcp>hos config>sub-mgmt>lu-db>pppoe>host		
Description	This command configures the group interface.		
Parameters	<i>ip-int-name</i> — Specifies the IP interface name.		
Values	32 chars max (must start with a letter)		

service

Syntax	service <i>service-id</i> no service
Context	config>sub-mgmt>lu-db>dhcp>hos config>sub-mgmt>lu-db>pppoe>host
Description	This command sets retail-service for a given subscriber host.
Parameters	<i>service-id</i> — Specifies the service ID as an interger.
Values	1-2147483648

policy

Syntax	policy <i>msap-policy-name</i> no policy
Context	config>sub-mgmt>lu-db>dhcp>hos config>sub-mgmt>lu-db>pppoe>host
Description	This command configures the MSAP policy.
Parameters	<i>msap-policy-name</i> — Specifies the policy name.

retail-service

Syntax	[no] retail-service <i>service-id</i>
Context	config>sub-mgmt>lu-db>dhcp>hos config>sub-mgmt>lu-db>pppoe>host
Description	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.
Default	no retail-service

options

Syntax	options
Context	config>router>dhcp>local-dhcp-serve>pool config>router>dhcp>local-dhcp-serve>pool>subnet config>subscr-mgmt>loc-user-db>dhcp>host config>subscr-mgmt>loc-user-db>pppoe>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 <i>option-number</i> address [<i>ip-address...</i> (up to 4 max)] custom-option <i>option-number</i> hex <i>hex-string</i> custom-option <i>option-number</i> string <i>ascii-string</i> no custom-option <i>option-number</i>
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>pppoe>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	<i>option-number</i> — specifies the option number that the DHCP server uses to send the identification strings to the DHCP client. Values 1 — 254 address <i>ip-address</i> — Specifies the IP address of this host. hex <i>hex-string</i> — Specifies the hex value of this option. Values 0x0..0xFFFFFFFF...(maximum 254 hex nibbles) string <i>ascii-string</i> — Specifies the value of this option. Values Up to 127 characters maximum.

dns-server

Syntax	dns-server address [<i>ip-address...</i> (up to 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>pppoe>host>options
Description	This command configures the IP address of the DNS server.
Default	none
Parameters	<i>ip-address</i> — 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).

domain-name

Syntax	domain-name <i>domain-name</i> no domain-name
Context	config>router>dhcp>server>pool>options config>subscr-mgmt>loc-user-db>dhcp>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	<i>domain-name</i> — Specifies the domain name for the client. Values Up to 127 characters

lease-rebind-time

Syntax	lease-rebind-time [days <i>days</i>] [hrs <i>hours</i>] [min <i>minutes</i>] [sec <i>seconds</i>] no lease-rebind-time
Context	config>router>dhcp>server>pool>subnet>options config>subscr-mgmt>loc-user-db>dhcp>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	<i>time</i> — Specifies the lease rebind time. Values days: 0 — 3650 hours: 0 — 23 minutes: 0 — 59 seconds 0 — 59

lease-renew-time

Syntax	lease-renew-time [days <i>days</i>] [hrs <i>hours</i>] [min <i>minutes</i>] [sec <i>seconds</i>] no lease-renew-time
Context	config>router>dhcp>server>pool>options config>subscr-mgmt>loc-user-db>dhcp>host>options
Description	This command configures the time the client transitions to a renew state. The no form of the command removes the time from the configuration.
Default	none

Triple Play DHCP Configuration Commands

Parameters	<i>time</i> — Specifies the lease renew time.		
Values	days:	0 — 3650	
	hours:	0 — 23	
	minutes:	0 — 59	
	seconds	0 — 59	

lease-time

Syntax	lease-time [days <i>days</i>] [hrs <i>hours</i>] [min <i>minutes</i>] [sec <i>seconds</i>] no lease-time		
Context	config>router>dhcp>server>pool>options config>subscr-mgmt>loc-user-db>dhcp>host>options		
Description	This command configures the amount of time that the DHCP server grants to the DHCP client permission to use a particular IP address. The no form of the command removes the lease time parameters from the configuration.		
Default	none		
Parameters	<i>time</i> — Specifies the lease time.		
Values	days :	0 — 3650	
	hours	0 — 23	
	minutes:	0 — 59	
	seconds	0 — 59	

netbios-name-server

Syntax	netbios-name-server ip-address [<i>ip-address</i> ...(up to 4 max)] no netbios-name-server		
Context	config>router>dhcp>server>pool>options config>subscr-mgmt>loc-user-db>dhcp>host>options config>subscr-mgmt>loc-user-db>pppoe>host>options		
Description	This command configures up to four Network Basic Input/Output System (NetBIOS) name server IP addresses.		
Default	none		
Parameters	<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

Syntax	netbios-node-type <i>netbios-node-type</i> no netbios-node-type

Context	config>router>dhcp>server>pool>options config>subscr-mgmt>loc-user-db>dhcp>host>options
Description	This command configures the Network Basic Input/Output System (NetBIOS) node type.
Default	none
Parameters	<i>netbios-node-type</i> — Specifies the netbios node type.
Values	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.

subnet

Syntax	subnet { <i>ip-address/mask</i> <i>ip-address netmask</i> } [create] no subnet { <i>ip-address/mask</i> <i>ip-address netmask</i> }
Context	config>router>dhcp>server>pool
Description	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.
Default	none
Parameters	<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

Syntax	[no] address-range <i>start-ip-address end-ip-address</i> [failover { local remote }]
Context	config>router>dhcp>server>pool>subnet
Description	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).
Default	none

Triple Play DHCP Configuration Commands

Parameters	<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). <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). failover local — Specifies that the DHCP server failover control type is in control under normal operation. failover remote — Specifies that the remote DHCP server failover system is in control under normal operation.

exclude-addresses

Syntax	[no] exclude-addresses <i>start-ip-address</i> [<i>end-ip-address</i>]
Context	config>router>dhcp>server>pool>subnet
Description	This command specifies a range of IP addresses that excluded from the pool of IP addresses in this subnet.
Default	none
Parameters	<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). <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).

maximum-declined

Syntax	maximum-declined <i>maximum-declined</i> no maximum-declined
Context	config>router>dhcp>server>pool>subnet
Description	This command configures the maximum number of declined addresses allowed.
Default	64
Parameters	<i>maximum-declined</i> — Specifies the maximum number of declined addresses allowed. Values 0 — 4294967295

minimum-free

Syntax	minimum-free <i>minimum-free</i> no minimum-free

Context	config>router>dhcp>server>pool>subnet
Description	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.
Default	1
Parameters	<i>minimum-free</i> — Specifies the minimum number of free addresses in this subnet.
Values	0 — 255

default-router

Syntax	default-router <i>ip-address</i> [<i>ip-address...</i> (up to 4 max)] no default-router
Context	config>router>dhcp>server>pool>subnet>options
Description	This command configures the IP address of the default router for a DHCP client. Up to four IP addresses can be specified. The no form of the command removes the address(es) from the configuration.
Default	none
Parameters	<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

Syntax	subnet-mask <i>ip-address</i> no subnet-mask
Context	config>router>dhcp>local-dhcp-serve>pool>subnet>options config>subscr-mgmt>loc-user-db>dhcp>host>options
Description	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 no form of the command removes the address from the configuration.
Default	none
Parameters	<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).

use-gi-address

Syntax	[no] use-gi-address
Context	config>router>dhcp>server

Triple Play DHCP Configuration Commands

Description	This command enables the use of gi-address matching. If the gi-address flag is enabled, a pool can be used even if a subnets is not found. If the local-user-db-name is not used, the gi-address flag is used and addresses are handed out by GI only. If a user must be blocked from getting an address the server maps to a local user database and configures the user with no address. A pool can include multiple subnets. Since the GI is shared by multiple subnets in a subscriber-interface the pool may provide IP addresses from any of the subnets included when the GI is matched to any of its subnets. This allows a pool to be created that represents a sub-int.
Default	no use-gi-address

use-pool-from-client

Syntax	[no] use-pool-from-client
Context	config>router>dhcp>server
Description	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 is specified by the value of the gi-address.
Default	no use-pool-from-client

user-db

Syntax	user-db <i>local-user-db-name</i> no user-db
Context	config>router>dhcp>server
Description	This command configures a local user database for authentication.
Default	not enabled
Parameters	<i>local-user-db-name</i> — Specifies the name of a local user database.

Service Commands

dhcp

Syntax	dhcp
Context	config>service>vpls>sap config>service>vpls>spoke-sdp config>service>vpls>mesh-sdp config>service>ies>interface config>service>vprn>interface config>service>vprn>sub-if config>service>vprn>sub-if>group-interface config>service>ies>sub-if>group-interface config>service>ies>sub-if config>service>ies>sub-if>grp-if
Description	This command enables the context to configure DHCP parameters.

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>vprn>sub-if>grp-if>dhcp
Description	This command enables the clients that will try to contact the DHCP server(s). The no form of the command removes the server client type from the configuration.
Parameters	dhcp — Specifies that the DHCP relay will forward requests to the DHCP server(s). pppoe — Specifies that PPPoE will attempt to request an IP address for a PPPoE client from the DHCP server(s)ly assigned to PPPoE node.

lease-populate

Syntax	lease-populate [<i>nbr-of-leases</i>] lease-populate [<i>nbr-of-leases</i>] l2-header [<i>mac ieee-address</i>] no lease-populate
Context	config>service>vprn>interface>dhcp config>service>vprn>sub-if>group-if>dhcp config>service>vprn>interface>dhcp

Description	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 is used to define the number lease state table entries allowed for this SAP 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 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. • Generate dynamic ARP replies if arp-reply-agent is enabled.
Default	no lease-populate
Parameters	<i>nbr-of-leases</i> — Specifies the number of DHCP leases allowed. Values 1 — 8000 l2-header — 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. mac — Specifies that the provisioned MAC 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.

lease-populate

Syntax	lease-populate [<i>nbr-of-entries</i>] lease-populate [<i>nbr-of-entries</i>] l2-header [<i>mac ieee-address</i>] no lease-populate
Context	config>service>vprn>if>dhcp config>service>vprn>sub-if>dhcp config>service>vprn>sub-if>grp-if>dhcp
Description	This command enables dynamic host lease state management for VPLS SAPs and VPRN or IES IP interfaces. Lease state information is extracted from snooped or relayed DHCP ACK messages to populate lease state table entries for the SAP or IP interface. The optional <i>number-of-entries</i> parameter defines the number lease state table entries allowed for this SAP or IP interface. If <i>number-of-entries</i> is not specified, only a single entry is allowed. Once the maximum number of entries has been reached, subsequent lease state entries are not allowed. If lease state population is enabled and an entry cannot be retained in the table, the DHCP Relay or DHCP snoop function will prevent the far-end host from receiving the DHCP ACK message. 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. Anti-spoof filtering is only available on VPLS SAPs, IES IP interfaces terminated on a SAP or VPRN IP interfaces terminated on a SAP.

The retained lease state information representing dynamic hosts may be used to populate the system's ARP cache based the arp-populate feature. ARP-populate functionality is only available for static and dynamic hosts associated with IES and VPRN SAP bound IP interfaces.

The retained lease state information representing dynamic hosts may be used to 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.

Default	not enabled
Parameters	<i>nbr-of-entries</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 DHCP ACK messages are discarded.
Values	1 — 8000

lease-populate

Syntax	lease-populate [<i>nbr-of-leases</i>] no lease-populate
Context	config>service>vpls>sap>dhcp config>service>ies>if>sap>dhcp config>service>ies>if>ipv6>dhcp-relay
Description	This command specifies the maximum number of DHCP or DHCP6 lease states allocated by the DHCP relay function, allowed on this interface. The no form of the command disables dynamic host lease state management.
Default	no lease-populate
Parameters	<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.
Values	1 — 8000

match-circuit-id

Syntax	[no] match-circuit-id
Context	config>service>ies>sub-if>grp-if>dhcp

Triple Play DHCP Configuration Commands

```
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 we use this as part of the key to guarantee correctness in our lookup. This is really only needed when we are dealing with an IP aware DSLAM that proxies the client HW mac address.
Default	no match-circuit-id

option

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

action

Syntax	action {replace drop keep} no action
Context	<pre>config>service>vpls>sap>dhcp>option config>service>ies>interface>dhcp>option config>service>vprn>interface>dhcp>option config>service>vprn>sub-if>grp-if>dhcp>option config>service>ies>sub-if>grp-if>dhcp</pre>
Description	This command configures the Relay Agent Information Option (Option 82) processing. The no form of this command returns the system to the default value.
Default	The default is to keep the existing information intact.

- Parameters**
- replace** — 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).
 - drop** — The DHCP packet is dropped if an Option 82 field is present, and a counter is incremented.
 - keep** — 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.
- 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.
- 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.

circuit-id

- Syntax** **circuit-id [ascii-tuple | vlan-ascii-tuple]**
no circuit-id
- Context** config>service>vpls>sap>dhcp>option
- Description** When enabled, the router sends an ASCII-encoded tuple in the **circuit-id** sub-option of the DHCP packet. This ASCII-tuple consists of the access-node-identifier, service-id, and SAP-ID, separated by “|”.
- In order to send a tuple in the circuit ID, the **action replace** command must be configured in the same context.
- If disabled, the **circuit-id** sub-option of the DHCP packet will be left empty.
- The **no** form of this command returns the system to the default.
- Default** circuit-id
- Parameters**
- ascii-tuple** — Specifies that the ASCII-encoded concatenated tuple consisting of the access-node-identifier, service-id, and interface-name is used.
 - 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 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

- Syntax** **circuit-id [ascii-tuple | ifindex | sap-id | vlan-ascii-tuple]**
no circuit-id
- Context** 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

Triple Play DHCP Configuration Commands

Description	When enabled, the router sends an ASCII-encoded tuple in the circuit-id sub-option of the DHCP packet. This ASCII-tuple consists of the access-node-identifier, service-id, and SAP-ID, separated by “ ”. In order to send a tuple in the circuit ID, the action replace command must be configured in the same context. If disabled, the circuit-id sub-option of the DHCP packet will be left empty. The no form of this command returns the system to the default.
Default	circuit-id ascii-tuple — Specifies that the ASCII-encoded concatenated tuple will be used which consists of the access-node-identifier, service-id, and interface-name, separated by “ ”. ifindex — Specifies that the interface index will be used. (The If Index of a router interface can be displayed using the command <code>show>router>interface>detail</code>) sap-id — Specifies that the SAP identifier will be used. 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

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

vendor-specific-option

Syntax	[no] vendor-specific-option
Context	<code>config>service>vpls>sap>dhcp>option</code> <code>config>service>ies>if>dhcp>option</code>


```
config>service>vprn>if>dhcp>option
config>service>ies>sub-if>grp-if>dhcp>option
config>service>vprn>sub-if>grp-if>dhcp>option
```

Description This command configures the Alcatel-Lucent vendor specific sub-option of the DHCP relay packet.

client-mac-address

Syntax [no] client-mac-address

Context 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

Description 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

Syntax [no] pool-name

Context 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

Description This command sends the pool name in the Alcatel vendor specific suboption of the DHCP relay packet.

The **no** form of the command disables the sending.

sap-id

Syntax [no] sap-id

Context 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

Description 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

Syntax	[no] service-id
Context	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
Description	This command enables the sending of the service 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 service ID in the Alcatel-Lucent vendor specific sub-option of the DHCP relay packet.

string

Syntax	[no] string text
Context	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
Description	This command specifies the string in the Alcatel-Lucent vendor specific sub-option of the DHCP relay packet. The no form of the command returns the default value.
Parameters	<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

Syntax	[no] system-id
Context	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
Description	This command specifies whether the system-id is encoded in the Alcatel-Lucent vendor specific sub-option of Option 82.

proxy-server

Syntax	proxy-server
Context	config>service>vpls>sap>dhcp config>service>vprn>sub-if>grp-if>dhcp

Description This command configures the DHCP proxy server.

emulated-server

Syntax **emulated-server** *ip-address*
no emulated-server

Context 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

Description 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.

The **no** form of this command reverts to the default setting. The local proxy server will not become operational without the emulated-server address being specified.

Parameters *ip-address* — 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

Syntax **emulated-server** *ip-address*
no emulated-server

Context config>service>vprn>if>dhcp>proxy
 config>service>vprn>sub-if>grp-if>dhcp>proxy-server

Description 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.

Parameters *ip-address* — Specifies the emulated server's IP address.

lease-time

Syntax **lease-time** [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*] [**radius-override**]
no lease-time

Context 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

Triple Play DHCP Configuration Commands

Description	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 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.
Default	7 days 0 hours 0 seconds
Parameters	radius-override — 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. Values 0 — 3650 <i>hours</i> — Specifies the number of hours that the given IP address is valid. Values 0 — 23 <i>minutes</i> — Specifies the number of minutes that the given IP address is valid. Values 0 — 59 <i>seconds</i> — Specifies the number of seconds that the given IP address is valid. Values 0 — 59

snoop

Syntax	snoop no snoop
Context	config>service>vpls>sap>dhcp config>service>vpls>spoke-sdp>dhcp config>service>vpls>mesh-sdp>dhcp config>service>vprn>if>dhcp>option config>service>vprn>if>dhcp>option
Description	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. Use the no form of the command to disable DHCP snooping on the specified SAP or SDP binding.
Default	no snoop

gi-address

Syntax	gi-address <i>ip-address</i> [<i>src-ip-addr</i>] no gi-address
Context	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
```

Description 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.

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.

Default no gi-address

Parameters *ip-address* — Specifies the host IP address to be used for DHCP relay packets.
src-ip-address — Specifies the source IP address to be used for DHCP relay packets.

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.

Default no server

Parameters *server* — Specify the DHCP server IP address.

relay-plain-bootp

Syntax [**no**] **relay-plain-bootp**

Context config>service>vprn>if>dhcp

Description This command enables the relaying of plain BOOTP packets.
 The **no** form of the command disables the relaying of plain BOOTP packets.

use-arp

Syntax	[no] use-arp
Context	config>service>vprn>if>dhcp
Description	This command enables the use of ARP to determine the destination hardware address. The no form of the command disables the use of ARP to determine the destination hardware address

trusted

Syntax	[no] trusted
Context	config>service>ies>if>dhcp config>service>vprn>if>dhcp config>service>vprn>sub-if>grp-if>dhcp config>service>ies>sub-if>grp-if>dhcp
Description	This command enables relaying of untrusted packets. The no form of this command disables the relay.
Default	not enabled

host-connectivity-verify

Syntax	host-connectivity-verify [interval <i>interval</i>] [action {remove alarm}]
Context	config>service>vprn>if>sap config>service>vprn>sub-if>grp-if config>service>vprn>sub-if>grp-if>dhcp
Description	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.
Default	no host-connectivity-verify
Parameters	interval <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. Values 1— 6000) Note that a zero value can be used by the SNMP agent to disable host-connectivity-verify.) 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.). DHCP-RELEASE will be signaled to

corresponding DHCP server. Static hosts will never be removed. The **alarm** keyword raises an alarm indicating that the host is disconnected.

dhcp

Syntax	dhcp
Context	config>service>vprn>interface config>service>vprn> config>service>vprn>sub-if>grp-if
Description	This command enables the context to configure DHCP parameters.

action

Syntax	action {replace drop keep} no action
Context	config>service>vprn>if>dhcp>option config>service>vprn>sub-if>grp-if>dhcp>option
Description	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. The no form of this command returns the system to the default value.
Default	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.
Parameters	replace — 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). drop — The packet is dropped, and an error is logged. keep — 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. 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. 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.

match-circuit-id

Syntax	[no] match-circuit-id

Triple Play DHCP Configuration Commands

Context	config>service>vprn>sub-if>grp-if>dhcp
Description	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. 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. When the match-circuit-id command is enabled we use this as part of the key to guarantee correctness in our lookup. This is really only needed when we are dealing with an IP aware DSLAM that proxies the client hardware MAC address.
Default	no match-circuit-id

option

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

Syntax	[no] string <i>text</i>
Context	config>service>vprn>if>dhcp>option>vendor config>service>vprn>sub-if>grp-if>dhcp>option>vendor
Description	This command specifies the vendor specific suboption string of the DHCP relay packet. The no form of the command returns the default value.
Parameters	<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

Syntax	[no] system-id
Context	config>service>vprn>if>dhcp>option>vendor config>service>vprn>sub-if>grp-if>dhcp>option>vendor

Triple Play DHCP Configuration Commands

Description	This command specifies whether the system-id is encoded in the Alcatel-Lucent vendor specific sub-option of Option 82.
Default	None

proxy-server

Syntax	proxy-server
Context	config>service>if>dhcp config>service>vprn>sub-if>grp-if>dhcp
Description	This command configures the DHCP proxy server.

emulated-server

Syntax	emulated-server <i>ip-address</i> no emulated-server
Context	config>service>vprn>if>dhcp>proxy config>service>vprn>sub-if>grp-if>dhcp>proxy-server
Description	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. The no form of this command reverts to the default setting. The local proxy server will not become operational without a specified emulated server address.
Parameters	<i>ip-address</i> — Specifies the emulated server address.
Default	Note that for a retail interface, the default is the local interface.

lease-time

Syntax	lease-time [days <i>days</i>] [hrs <i>hours</i>] [min <i>minutes</i>] [sec <i>seconds</i>] [radius-override] no lease-time
Context	config>service>vprn>if>dhcp>proxy config>service>vprn>sub-if>grp-if>dhcp>proxy-server
Description	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 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.
Default	7 days 0 hours 0 seconds
Parameters	radius-override — 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.

Values 0 — 3650

hours — Specifies the number of hours that the given IP address is valid.

Values 0 — 23

minutes — Specifies the number of minutes that the given IP address is valid.

Values 0 — 59

seconds — Specifies the number of seconds that the given IP address is valid.

Values 0 — 59

server

Syntax	server <i>server1</i> [<i>server2</i> ...(up to 8 max)]
Context	config>service>vprn>if>dhcp config>service>vprn>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. 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.
Default	no server
Parameters	<i>server</i> — Specify the DHCP server IP address.

host-connectivity-verify

Syntax	host-connectivity-verify [<i>interval interval</i>] [<i>action {remove alarm}</i>]
Context	config>service>vprn>if>sap config>service>vprn>sub-if>grp-if config>service>vprn>sub-if>grp-if>dhcp
Description	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.
Default	no host-connectivity-verify
Parameters	interval <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.

Values 1— 6000) Note that a zero value can be used by the SNMP agent to disable host-connectivity-verify.)

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.). DHCP-RELEASE will be signaled to corresponding DHCP server. Static hosts will never be removed. The **alarm** keyword raises an alarm indicating that the host is disconnected.

Interface Commands

local-proxy-arp

Syntax	[no] local-proxy-arp
Context	config>service>vprn>interface config>service>vprn>sub-if>grp-if
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

mac

Syntax	[no] mac <i>ieee-mac-address</i>
Context	config>service>vprn>interface config>service>vprn>if>vrrp config>service>vprn>sub-if>grp-if
Description	This command assigns a specific MAC address to a VPRN IP interface. The no form of this command returns the MAC address of the IP interface to the default value.
Default	The physical MAC address associated with the Ethernet interface that the SAP is configured on.
Parameters	<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

Syntax	[no] proxy-arp <i>policy-name</i> [<i>policy-name</i>...(up to 5 max)]
Context	config>service>vprn>interface config>service>vprn>sub-if>grp-if
Description	This command enables a proxy ARP policy for the interface. The no form of this command disables the proxy ARP capability.
Default	no proxy-arp

Triple Play DHCP Configuration Commands

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. 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-interface

Syntax **redundant-interface** *red-ip-int-name*
no redundant-interface

Context config>service>vpn
 config>service>vpn>sub-if>grp-if

Description This command configures a redundant interface used for dual homing.

Parameters *red-ip-int-name* — Specifies the redundant IP interface name.

remote-proxy-arp

Syntax **[no] remote-proxy-arp**

Context config>service>vpn>interface
 config>service>vpn>sub-if>grp-if

Description This command enables remote proxy ARP on the interface.

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.

Default no remote-proxy-arp

Subscriber Interface Commands

subscriber-interface

Syntax	[no] subscriber-interface <i>ip-int-name</i>
Context	config>service>ies config>service>vprn
Description	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. Use the no form of the command to remove the subscriber interface.
Parameters	<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

Syntax	[no] group-interface <i>ip-int-name</i>
Context	config>service>ies>sub-if
Description	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.
Default	none
Parameters	<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

Syntax	authentication-policy <i>name</i> no authentication-policy
Context	config>service>ies>sub-if>grp-if
Description	This command assigns a RADIUS authentication policy to the interface. The no form of this command removes the policy name from the group interface configuration.
Default	no authentication-policy
Parameters	<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

Syntax	local-user-db <i>local-user-db-name</i> [create] no local-user-db <i>local-user-db-name</i>
Context	config>subscr-mgmt
Description	This command enables the context to configure a local user database.
Default	not enabled
Parameters	<i>local-user-db-name</i> — Specifies the name of a local user database.

dhcp

Syntax	dhcp
Context	config>subscr-mgmt>loc-user-db
Description	This command configures DHCP host parameters.

pppoe

Syntax	pppoe
Context	config>subscr-mgmt>loc-user-db
Description	This command configures PPPoE hosts.

mask

Syntax	mask type <i>dhcp-match-type</i> {[prefix-string <i>prefix-string</i> prefix-length <i>prefix-length</i>] [suffix-string <i>suffix-string</i> suffix-length <i>suffix-length</i>]} no mask type <i>dhcp-match-type</i>				
Context	config>subscr-mgmt>loc-user-db>dhcp config>subscr-mgmt>loc-user-db>pppoe				
Description	This command configures the mask.				
Parameters	<i>dhcp-match-type</i> — Specifies up to four matching types to identify a host. <table> <tr> <td>Values</td><td>DHCP: circuit-id, option60, remote-id, sap-id, string, system-id PPPOE: circuit-id, remote-id, service-name, username</td></tr> <tr> <td>Values</td><td><i>prefix-string</i> <i>prefix-string</i></td></tr> </table>	Values	DHCP: circuit-id, option60, remote-id, sap-id, string, system-id PPPOE: circuit-id, remote-id, service-name, username	Values	<i>prefix-string</i> <i>prefix-string</i>
Values	DHCP: circuit-id, option60, remote-id, sap-id, string, system-id PPPOE: circuit-id, remote-id, service-name, username				
Values	<i>prefix-string</i> <i>prefix-string</i>				

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 DHCP 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 DHCP 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 DHCP circuit ID.

Values 1— 127

host

Syntax	host <i>host-name</i> [create] no host <i>host-name</i>
Context	config>subscr-mgmt>loc-user-db>dhcp config>subscr-mgmt>loc-user-db>pppoe
Description	This command defines a DHCP or PPPOE subscriber.
Parameters	<i>host-name</i> — create — Keyword used to create the host name. The create keyword requirement can be enabled/disabled in the environment>create context.

address

Syntax	address <i>gi-address</i> address <i>ip-address</i> address pool <i>pool-name</i> no address
Context	config>subscr-mgmt>loc-user-db>dhcp>host config>subscr-mgmt>loc-user-db>pppoe>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.

Triple Play DHCP Configuration Commands

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. <i>ip-address</i> — Specifies the fixed IP address to use for this host. <i>pool-name</i> — Specifies the 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.

authentication-policy

Syntax	authentication-policy <i>policy-name</i> no authentication-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	<i>policy-name</i> — Specifies the authentication policy name.

domain-name

Syntax	domain-name <i>domain-name</i> no domain-name
Context	config>subscr-mgmt>loc-user-db>dhcp>host
Description	This command sets the domain name which can be appended to user-name in RADIUS-authentication-request message for the given host.
Parameters	<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

Syntax	host-identification
Context	config>subscr-mgmt>loc-user-db>dhcp>host config>subscr-mgmt>loc-user-db>pppoe>host
Description	This command enables the context to configure host identification parameters.

server

Syntax	server <i>ip-address</i> no server
Context	config>subscr-mgmt>loc-user-db>dhcp>host
Description	This command configures the IP address of the DHCP server in which to relay. The no form of the command removes the value from the configuration.
Default	no server
Parameters	<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

Syntax	circuit-id string <i>ascii-string</i> circuit-id hex <i>hex-string</i> no circuit-id
Context	config>subscr-mgmt>loc-user-db>dhcp>host>host-ident config>subscr-mgmt>loc-user-db>pppoe>host>host-ident
Description	This command specifies the circuit-id to match.
Parameters	<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.
Values	0x0..0xFFFFFFFF (maximum 254 hex nibbles)

mac

Syntax	mac <i>ieee-address</i> no mac
Context	config>subscr-mgmt>loc-user-db>dhcp>host>host-ident config>subscr-mgmt>loc-user-db>pppoe>host>host-ident
Description	This command specifies the MAC address to match.

Triple Play DHCP Configuration Commands

Parameters *ieee-address* — 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.

option60

Syntax **option60** *hex-string*
no option60

Context config>subscr-mgmt>loc-user-db>dhcp>host>host-ident

Description This command specifies the Vendor-Identifying Vendor Option to match. Option 60 is encoded as Type-Length-Value (TLV). The *hex-string* portion of Option 60 in the received DHCP request is used for matching. Only the first 32 bytes can be defined here. If Option 60 from the message is longer, those bytes are ignored.

Default no option60

Parameters *hex-string* — Specifies the hex value of this option.

Values 0x0..0xFFFFFFFF...(maximum 254 hex nibbles)

remote-id

Syntax **remote-id** *remote-id*
no remote-id

Context config>subscr-mgmt>loc-user-db>dhcp>host>host-ident
config>subscr-mgmt>loc-user-db>pppoe>host>host-ident

Description This command specifies the remote id of this host.
The **no** form of this command returns the system to the default.

Default no remote-id

Parameters *remote-id* — Specifies the remote-id.

service-name

Syntax **service-name** *service-name*
no service-name

Context config>subscr-mgmt>loc-user-db>pppoe>host>host-ident

Description This command specifies the service-name tag in PADI and/or PADR packets to match for PPPoE hosts.

Parameters *service-name* — Specifies a PPPoE service name, up to 255 characters maximum.

sap-id

Syntax	sap-id <i>sap-id</i> no sap-id
Context	config>subscr-mgmt>loc-user-db>dhcp>host>host-ident
Description	This command specifies the SAP ID from the Alcatel Vendor Specific Sub-option in Option 82 to match.
Parameters	<i>sap-id</i> — Specifies a SAP ID, up to 255 characters maximum.

service-id

Syntax	service-id <i>service-id</i> no service-id
Context	config>subscr-mgmt>loc-user-db>dhcp>host>host-ident
Description	This command specifies an existing service ID from the Alcatel Vendor Specific Sub-Option in Option 82 to match.
Parameters	<i>service-id</i> — Specifies an existing service ID.
Values	1 — 2147483647

string

Syntax	string <i>string</i> no string
Context	config>subscr-mgmt>loc-user-db>dhcp>host>host-ident
Description	This command specifies the string from the Alcatel Vendor Specific Sub-Option in Option 82 to match.
Parameters	<i>string</i> — Specifies the string, up to 255 characters maximum.

system-id

Syntax	system-id <i>system-id</i> no system-id
Context	config>subscr-mgmt>loc-user-db>dhcp>host>host-ident
Description	This command specifies the system ID from the Alcatel Vendor Specific Sub-Option in Option 82 to match.
Parameters	<i>system-id</i> — Specifies the system ID, up to 255 characters maximum.

username

Syntax	username <i>user-name</i> [no-domain] username <i>user-name</i> domain-only no username
Context	config>subscr-mgmt>loc-user-db>pppoe>host>host-ident
Description	This command specifies how the username is specified.
Parameters	<i>username</i> — Specifies the user name of this host. no-domain — No username is specified. domain-only — Only the domain part of the username is specified, for example, alcatel-lucent.com

identification-strings

Syntax	identification-strings <i>option-number</i> [create] no identification-strings
Context	config>subscr-mgmt>loc-user-db>dhcp>host config>subscr-mgmt>loc-user-db>pppoe>host
Description	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 config>subscriber-mgmt>sub-ident-policy 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.
Default	254
Parameters	<i>option-number</i> — Specifies identification strings for the subscriber Values 1 — 254

ancp-string

Syntax	ancp-string <i>ancp-string</i> no ancp-string
Context	config>subscr-mgmt>loc-user-db>dhcp>host>ident-strings config>subscr-mgmt>loc-user-db>pppoe>host>ident-strings
Description	This command specifies the ANCP string which is encoded in the identification strings.
Parameters	<i>ancp-string</i> — Specifies the the ANCP string, up to 63 characters, maximum.

app-profile-string

Syntax	app-profile-string <i>app-profile-string</i> no app-profile-string
Context	config>subscr-mgmt>loc-user-db>dhcp>host>ident-strings config>subscr-mgmt>loc-user-db>pppoe>host>ident-strings
Description	This command specifies the application profile string which is encoded in the identification strings.
Parameters	<i>app-profile-string</i> — Specifies the the application profile string, up to 16 characters, maximum.

inter-dest-id

Syntax	inter-dest-id <i>intermediate-destination-id</i> no inter-dest-id
Context	config>subscr-mgmt>loc-user-db>dhcp>host>ident-strings config>subscr-mgmt>loc-user-db>pppoe>host>ident-strings
Description	This command specifies the intermediate destination identifier which is encoded in the identification strings.
Parameters	<i>intermediate-destination-id</i> — Specifies the intermediate destination identifier, up to 32 characters, maximum.

sla-profile-string

Syntax	sla-profile-string <i>sla-profile-string</i> no sla-profile-string
Context	config>subscr-mgmt>loc-user-db>dhcp>host>ident-strings config>subscr-mgmt>loc-user-db>pppoe>host>ident-strings
Description	This command specifies the SLA profile string which is encoded in the identification strings.
Parameters	<i>sla-profile-string</i> — Specifies the SLA profile string, up to 16 characters, maximum.

sub-profile-string

Syntax	sub-profile-string <i>sub-profile-string</i> no sub-profile-string
Context	config>subscr-mgmt>loc-user-db>dhcp>host>ident-strings config>subscr-mgmt>loc-user-db>pppoe>host>ident-strings
Description	This command specifies the subscriber profile string which is encoded in the identification strings.
Parameters	<i>sub-profile-string</i> — Specifies the subscriber profile string, up to 16 characters, maximum.

subscriber-id

Syntax	subscriber-id <i>sub-ident-string</i> no subscriber-id
Context	config>subscr-mgmt>loc-user-db>dhcp>host>ident-strings config>subscr-mgmt>loc-user-db>pppoe>host>ident-strings
Description	This command specifies the subscriber ID which is encoded in the identification strings.
Parameters	<i>sub-ident-string</i> — Specifies the subscriber ID string, up to 32 characters, maximum.

l2tp

Syntax	l2tp
Context	config>subscr-mgmt>loc-user-db>pppoe>host
Description	This command configures L2TP for the host.

group

Syntax	group <i>tunnel-group-name</i> no group
Context	config>subscr-mgmt>loc-user-db>pppoe>host>l2tp
Description	This command configures the L2TP tunnel group. The tunnel-group-name is configured in the config>router>l2tp context. Refer to the 7750 SR OS Router Configuration Guide.
Parameters	<i>tunnel-group-name</i> — Specifies an existing tunnel L2TP group up to 63 characters in length.

authentication-policy

Syntax	authentication-policy <i>policy-name</i> no authentication-policy
Context	config>subscr-mgmt>loc-user-db>pppoe>host
Description	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

Syntax	pado-delay <i>deci-seconds</i> no pado-delay
Context	config>subscr-mgmt>loc-user-db>pppoe>host

Description	This command configures the delay timeout before sending a PPPoE Active Discovery Offer. (PADO)
Parameters	<i>deci-seconds</i> — Specifies the delay timeout before sending a PADO.
Values	1 — 30

mask

Syntax	mask type <i>pppoe-match-type</i> {[prefix-string <i>prefix-string</i> prefix-length <i>prefix-length</i>] [suffix-string <i>suffix-string</i> suffix-length <i>suffix-length</i>]} no mask type <i>pppoe-match-type</i>
Context	config>subscr-mgmt>loc-user-db>pppoe
Syntax	mask type <i>dhcp-match-type</i> {[prefix-string <i>prefix-string</i> prefix-length <i>prefix-length</i>] [suffix-string <i>suffix-string</i> suffix-length <i>suffix-length</i>]} no mask type <i>dhcp-match-type</i>
Context	config>subscr-mgmt>loc-user-db>dhcp
Description	This command configures the mask.
Parameters	<i>pppoe-match-type</i> — Specifies the sub-option inserted by the PPPoE intermediate agent. Values circuit-id, remote-id, service-name, username <i>dhcp-match-type</i> — 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 <i>prefix-string</i> — 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 <i>prefix-length</i> — 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 <i>suffix-string</i> — 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 <i>suffix-length</i> — 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

Syntax	match-list <i>dhcp-match-type-1</i> [<i>dhcp-match-type-2...</i> (up to 4 max)] no match-list
Context	config>subscr-mgmt>loc-user-db>dhcp config>subscr-mgmt>loc-user-db>pppoe
Description	This command specifies the type of matching done to identify a host. There are different match-types for PPPoE hosts of which a maximum of 3 can be specified (4 for DHCP)
Default	no match-list
Parameters	<i>dhcp-match-type-x</i> — Specifies up to four matching types to identify a DHCP host. Values circuit-id, mac, remote-id, sap-id, service-id, string, system-id, option60

password

Syntax	password { ignore chap <i>string</i> pap <i>string</i> } [hash hash2] no password
Context	config>subscr-mgmt>loc-user-db>pppoe>host
Description	This command specifies a password type or configures password string for pap or chap. The pap and chap passwords are stored in a hashed format in the config files. The hash hash2 optional keywords are used for config execution. 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.
Parameters	ignore — 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. chap <i>string</i> — Specifies that the password for Challenge-Handshake Authentication Protocol (CHAP) is used. Only a password received with the CHAP protocol will be accepted. pap <i>string</i> — 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. hash hash2 — Specifies hashing scheme.

retail-service-id

Syntax	retail-service-id <i>service-id</i> no retail-service-id
Context	config>subscr-mgmt>loc-user-db>pppoe>host

Description	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 no form of the command removes the service ID from the configuration.		
Default	no retail-service-id		
Parameters	<i>service-id</i> — Specifies the the retailer service ID.		
	Values	service-id:	1 — 2147483647
		service-name:	Service name up to 64 characters in length.

Show Commands

id

Syntax	id <i>service-id</i>
Context	show>service
Description	This command displays information for a particular service-id.
Parameters	<i>service-id</i> — The unique service identification number that identifies the service in the service domain. all — Display detailed information about the service. base — Display basic service information. fdb — Display FDB entries. labels — Display labels being used by this service. sap — Display SAPs associated to the service. sdp — Display SDPs associated with the service. split-horizon-group — Display split horizon group information. stp — Display STP information.

dhcp

Syntax	dhcp
Context	show>router show>service>id
Description	This command enables the context to show DHCP statistics.

dhcp6

Syntax	dhcp6
Context	show>router show>service>id
Description	This command enables the context to show DHCP6 statistics.

lease-state

Syntax	lease-state [detail] lease-state [detail] interface <i>interface-name</i> lease-state [detail] ipv6-address <i>ipv6-prefix[/prefix-length]</i> lease-state [detail] mac <i>ieee-address</i>
Context	show>service>id>dhcp6
Description	Shows 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      Stdbby
-----
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

Syntax	statistics [sap <i>sap-id</i>] [sdp [<i>sdp-id</i> [: <i>vc-id</i>]]] interface <i>ip-int-name</i>]]
Context	show>service>id>dhcp 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	<i>sap-id</i> — Specifies the physical port identifier portion of the SAP definition. See Common Service Commands on page 1020 for <i>sap-id</i> command syntax. <i>sdp-id</i> — The SDP ID to be shown. Values 1 — 17407 <i>vc-id</i> — The virtual circuit ID on the ID to be shown. Values 1 — 4294967295 <i>ip-int-name</i> <i>ip-address</i> — Displays statistics for the specified IP interface.
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.

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

Syntax	summary
Context	show>router>dhcp show>service>id>dhcp
Description	Display the status of the DHCP Relay and DHCP Snooping functions on each interface.
OutputOutput	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

```

Triple Play DHCP Configuration Commands

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#

statistics

Syntax	statistics [interface <i>ip-int-name</i>]
Context	show>router>dhcp6 show>service>id>dhcp6
Description	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

Syntax	summary
Context	show>router>dhcp6 show>service>id>dhcp6

Description Display 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.
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
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
-----
```

Triple Play DHCP Configuration Commands

```
No Matching Entries
=====
*A:ALA-48>show>router>dhcp>local-dhcp-server#
```

associations

Syntax	associations
Context	show>router>dhcp
Description	This command displays the interfaces associated with this DHCP 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

Syntax	declined-addresses <i>ip-address[/mask]</i> [detail] declined-addresses pool <i>pool-name</i>
Context	show>router>dhcp>local-dhcp-server
Description	This command display information about declined addresses.
Parameters	pool <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). detail — 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

Syntax	free-addresses <i>ip-address[/mask]</i> free-addresses summary [subnet <i>ip-address[/mask]</i>] free-addresses pool <i>pool-name</i>
Context	show>router>dhcp>local-dhcp-server
Description	This command displays the free addresses in a subnet.
Parameters	pool <i>pool-name</i> — Specifies a DHCP pool name on the router. subnet <i>subnet</i> — Specifies a subnet of IP addresses that are served from the pool. summary — 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#
```

leases

Syntax	leases leases <i>ip-address[/mask]</i> address-from-user-db [detail] leases <i>ip-address[/mask]</i> dhcp-host <i>dhcp-host-name</i> [detail] leases <i>ip-address[/mask]</i> pppoe-host <i>pppoe-host-name</i> [detail] leases <i>ip-address[/mask]</i> [detail]
Context	show>router>dhcp>local-dhcp-server
Description	This command displays the DHCP leases.
Parameters	<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. Values 0 — 32 address-from-user-db [detail] — Displays only leases that have ip-addresses from the local-user-db. dhcp-host <i>dhcp-host-name</i> [detail] — Shows all leases that match a certain DHCP host from the local-user-db. pppoe-host <i>pppoe-host-name</i> [detail] — Displays all leases that match a certain PPPoE host from the local-user-db.

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

```
*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#
```

server-stats

Syntax	server-stats
Context	show>router>dhcp>server
Description	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
```

```
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
=====
```

Triple Play DHCP Configuration Commands

```
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
=====
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

- Syntax** `subnet pool pool-name [subnet subnet]`
- Context** `show>router>dhcp>local-dhcp-server>statistics`
- Description** This command displays subnet statistics.
- Parameters**
- pool *pool-name*** — Specifies the pool name on the router.
 - subnet *subnet*** — 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

Syntax	lease-state [[sap sap-id] [sdp dp-id:vc-id] [interface interface-name] [ip-address ip-address[/mask]>] [mac ieee-address]] [detail]
Context	show>service>id>dhcp
Description	This command displays DHCP lease state related information.
Parameters	sap-id — Specifies the physical port identifier portion of the SAP definition. See Common Service Commands on page 1020 for <i>sap-id</i> command syntax. sdp-id — The SDP ID to be shown. Values 1 — 17407 vc-id — The virtual circuit ID on the SDP ID to be shown. Values 1 — 4294967295

servers

Syntax	servers
Context	show>router>dhcp
Description	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
```



```

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
=====

```

remap-lease-state

Syntax	remap-lease-state old-mac <i>ieee-address</i> mac <i>ieee-address</i> remap-lease-state sap <i>sap-id</i> [<i>mac ieee-address</i>]
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 <i>ieee-address</i> 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 <i>ieee-address</i> old-mac <i>ieee-address</i> — specifies the old MAC address to remap. mac <i>ieee-address</i> — 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 1020](#) 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

Syntax	dhcp
Context	clear>router
Description	This command enables the context to clear and reset DHCP entities.

dhcp6

Syntax	dhcp6
Context	clear>router
Description	This command enables the context to clear and reset DHCP6 entities.

lease-state

Syntax	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]
Context	clear>service>id>dhcp
Description	This command clears DHCP lease state information.
Parameters	no-dhcp-release — Clears the state without sending the DHCP release message. ip-address ip-address — Clears the DHCP IP address lease state information. The <i>ip-address</i> 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). mac ieee-address — 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. sap sap-id — clears DHCP SAP lease state information. See Common Service Commands on page 1020 for <i>sap-id</i> command syntax.

Triple Play DHCP Configuration 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 1020](#)

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

Syntax	local-dhcp-server <i>server-name</i>
Context	clear>router>dhcp
Description	This command clears DHCP server data.
Parameters	<i>server-name</i> — Clears data for the specified local DHCP server.

declined-addresses

Syntax	declined-addresses <i>ip-address[/mask]</i> declined-addresses pool <i>pool-name</i>
Context	clear>router>dhcp>local-dhcp-server
Description	This command clears declined DHCP addresses.
Parameters	<i>pool-name</i> — Specifies the declined pool name. <i>ip-address[/mask]</i> — Specifies the declined IP address and mask.

leases

Syntax	leases <i>ip-address[/mask]</i> [offered]
Context	clear>router>dhcp>local-dhcp-server
Description	This command clears DHCP leases.
Parameters	<i>ip-address[/mask]</i> — Clears the specified IP address and mask. offered — Clears leases in offered state only.

server-stats

Syntax	server-stats
Context	clear>router>dhcp>local-dhcp-server
Description	This command clears all server statistics.

lease-state

Syntax	lease-state [ip-address <i>ipv6-address/prefix-length</i>] [mac <i>ieee-address</i>]
Context	clear>service>id>dhcp6
Description	This command clears DHCP6 lease state information.
Parameters	ip-address <i>ipv6-address</i> — Clears the DHCP6 IP address lease state information. The <i>ipv6-address</i> 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 IPv6 addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets). mac <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.

statistics

Syntax	statistics [sap <i>sap-id</i> sdp [<i>sdp-id</i> [: <i>vc-id</i>]] interface <i>ip-int-name</i> <i>ip-address</i>]
Context	clear>router>dhcp
Description	This command clears DHCP statistics.
Parameters	sap <i>sap-id</i> — clears DHCP statistics. See Common Service Commands on page 1020 for <i>sap-id</i> command syntax. <i>sdp-id</i> — Clears DHCP SDP statistics. Values 1 — 17407 <i>vc-id</i> — Clears DHCP the SDP VC ID statistics. Values 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.

statistics

Syntax	statistics
Context	clear>router>dhcp6
Description	This command clears DHCP6 statistics.

Debug Commands

dhcp

Syntax	[no] dhcp <i>[ip-int-name]</i>
Context	debug>router>ip
Description	This command enables DHCP debugging. The no form of the command disables debugging.
Parameters	<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

Syntax	dhcp6 <i>[ip-int-name]</i> no dhcp6
Context	debug>router>ip
Description	This command enables DHCP debugging. The no form of the command disables debugging.
Parameters	<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

Syntax	detail-level { low medium high } no detail-level
Context	debug>router>ip>dhcp debug>router>local-dhcp-server debug>router>ip>dhcp6
Description	This command debugs the DHCP tracing detail level.

local-dhcp-server

Syntax	[no] local-dhcp-server <i>server-name</i> [lease-address <i>ip-address</i>] [no] local-dhcp-server <i>server-name</i> [mac <i>ieee-address</i>]

Triple Play DHCP Configuration Commands

Context	debug>router
Description	This command enables, disables or configures debugging for a local DHCP server.
Parameters	<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

Syntax	mode {dropped-only ingr-and-dropped egr-ingr-and-dropped} no mode
Context	debug>router>ip>dhcp debug>router>local-dhcp-server debug>router>ip>dhcp6
Description	This command debugs the DHCP tracing detail level.

Tools Commands

tools

Syntax	tools
Context	<root>
Description	This command enables the context to enable useful tools for debugging purposes.
Default	none
Parameters	dump — Enables dump tools for the various protocols. perform — Enables tools to perform specific tasks.

perform

Syntax	perform
Context	tools
Description	This command enables the context to enable tools to perform specific tasks.
Default	none

subscriber-mgmt

Syntax	subscriber-mgmt
Context	tools>perform
Description	This command enables tools to control subscriber management.

edit-pppoe-session

Syntax	edit-pppoe-session sap <i>sap-id</i> ip <i>ip-address</i> [subscriber <i>sub-ident-string</i>] [sub-profile-string <i>sub-profile-string</i>] [sla-profile-string <i>sla-profile-string</i>] [inter-dest-id <i>intermediate-destination-id</i>] [ancp-string <i>ancp-string</i>] [app-profile-string <i>app-profile-string</i>] edit-pppoe-session svc-id <i>service-id</i> ip <i>ip-address</i> [subscriber <i>sub-ident-string</i>] [sub-profile-string <i>sub-profile-string</i>] [sla-profile-string <i>sla-profile-string</i>] [app-profile-string <i>app-profile-string</i>] [inter-dest-id <i>intermediate-destination-id</i>] [ancp-string <i>ancp-string</i>]
Context	tools>perform>subscriber-mgmt
Description	This command modifies PPPoE session information.

Triple Play DHCP Configuration Commands

Parameters	<i>sap-id</i> — Specifies the physical port identifier portion of the SAP definition. See Common Service Commands on page 1020 for <i>sap-id</i> command syntax. <i>ip-address</i> — Specifies the IP address. <i>sub-ident-string</i> — Specifies a subscriber identification profile. <i>sub-profile-string</i> — Specifies the subscriber profile string, up to 16 characters, maximum. <i>service-id</i> — The service identification number that identifies the service in the domain. <i>intermediate-destination-id</i> — Specifies the intermediate destination identifier which is encoded in the identification strings. ancp-string <i>ancp-string</i> — Specifies the ASCII string of the DSLAM circuit ID name. <i>app-profile-string</i> — Specifies an application profile string.

eval-lease-state

Syntax	eval-lease-state [svc-id <i>service-id</i>] [sap <i>sap-id</i>] [subscriber <i>sub-ident-string</i>] [ip <i>ip-address</i>]
Context	tools>perform>subscriber-mgmt
Description	This command evaluates lease state.
Parameters	<i>sap-id</i> — Specifies the physical port identifier portion of the SAP definition. See Common Service Commands on page 1020 for <i>sap-id</i> command syntax. <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). <i>sub-ident-string</i> — Specifies the subscriber ID string, up to 32 characters, maximum. <i>service-id</i> — Specifies an existing service ID. Values 1 — 2147483647

local-user-db

Syntax	local-user-db <i>local-user-db-name</i>
Context	tools>perform>subscriber-mgmt
Description	This command enables tools for controlling the local user database.
Parameters	<i>local-user-db-name</i> — Specifies the name of a local user database.

dhcp

Syntax	dhcp
Context	tools>perform>subscriber-mgmt>local-user-db
Description	This command contains the tools used for controlling DHCP entries in the local user database.

host-lookup

Syntax	host-lookup [mac <i>ieee-address</i>] [remote-id <i>remote-id</i>] [sap-id <i>sap-id</i>] [service-id <i>service-id</i>] [string <i>vso-string</i>] [system-id <i>system-id</i>] [option60 <i>hex-string</i>] [circuit-id <i>circuit-id</i> circuit-id-hex <i>circuit-id-hex</i>]
Context	tools>perform>subscriber-mgmt>local-user-db>dhcp
Description	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.
Parameters	mac <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. <i>remote-id</i> — specifies what information goes into the remote-id sub-option in the DHCP relay packet. Values Up to 255 characters maximum sap-id — Specifies a SAP identifier to be used. See Common Service Commands on page 1020 for <i>sap-id</i> command syntax. <i>service-id</i> — Specifies an existing subscriber service ID. Values 1 — 2147483647 <i>vso-string</i> — Specifies a Vendor Specific Option (VSO) string. <i>system-id</i> — Specifies the system ID. Values up to 255 characters maximum. option60 <i>hex-string</i> — Specifies the content of option 60 for this lookup. Values 0x0..0xFFFFFFFF (maximum 64 hex nibbles) circuit-id <i>circuit-id</i> — specifies the circuit ID from the Option 82. circuit-id-hex <i>circuit-id-hex</i> — Specifies the circuit ID in hexadecimal format from the Option 82. Values 0x0..0xFFFFFFFF (maximum 254 hex nibbles)

pppoe

Syntax	pppoe
Context	tools>perform>subscriber-mgmt>local-user-db
Description	This command contains the tools used to control PPPoE entries in the local user database.

authentication

Syntax	authentication password <i>password</i> [mac <i>ieee-address</i>] [remote-id <i>remote-id</i>] [circuit-id <i>circuit-id</i>] user-name <i>user-name</i> [service-name <i>service-name</i>] authentication password <i>password</i> [mac <i>ieee-address</i>] [remote-id <i>remote-id</i>] [circuit-id-hex <i>circuit-id-hex</i>] user-name <i>user-name</i> [service-name <i>service-name</i>]
Context	tools>perform>subscriber-mgmt>local-user-db>pppoe
Description	This command authenticates PPPoE user name. As local user database PAP/CHAP authentication can only be used when the local user database is connected to the PPPoE node under the group interface, the user lookup will be performed with match-list username.
Parameters	password <i>password</i> — specifies the password of this host up to 32 characters in length. mac <i>ieee-address</i> — Specifies information about the MAC address of the PPPoE session. <i>remote-id</i> — specifies what information goes into the remote-id sub-option in the DHCP relay packet. Values Up to 255 characters maximum circuit-id <i>circuit-id</i> — specifies the circuit ID from the Option 82. circuit-id-hex <i>circuit-id-hex</i> — Specifies the circuit ID in hexadecimal format from the Option 82. Values 0x0..0xFFFFFFFF (maximum 254 hex nibbles) user-name <i>user-name</i> — Specifies the PPPoE user name. service-name <i>service-name</i> —

host-lookup

Syntax	host-lookup [mac <i>ieee-address</i>] [remote-id <i>remote-id</i>] [user-name <i>user-name</i>] [service-name <i>service-name</i>] [circuit-id <i>circuit-id</i> circuit-id-hex <i>circuit-id-hex</i>]
Context	tools>perform>subscr-mgmt>loc-user-db>pppoe
Description	This command performs a lookup in the local user database. mac <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. remote-id <i>remote-id</i> — 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 PPPoE 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)

Point-to-Point Protocol over Ethernet (PPPoE) Management

In This Chapter

This chapter provides information about using PPPoE, including theory, supported features and configuration process overview.

Topics in this chapter include:

- [PPPoE on page 424](#)
- [PPPoE Authentication and Authorization on page 427](#)
 - [General Flow on page 427](#)
 - [RADIUS on page 428](#)
 - [Local User Database Directly Assigned to PPPoE Node on page 429](#)
 - [Local DHCP Server with Local User Database on page 430](#)
- [Multiple Sessions Per MAC Address on page 432](#)
- [Private Retail Subnets on page 433](#)

PPPoE

A Broadband Remote Access Server (BRAS) is a device that terminates PPPoE sessions. PPPoE sessions are supported on Alcatel-Lucent Broadband Services Router (BSR) on IOM2. The Point-to-Point Protocol (PPP) is used for communications between a client and a server. Point-to-Point Protocol over Ethernet (PPPoE) is a network protocol used to encapsulate PPP frames inside Ethernet frames.

Ethernet networks are packet-based, unaware of connections or circuits. Using PPPoE, Alcatel-Lucent users can dial from one router to another over an Ethernet network, then establish a point-to-point connection and transport data packets over the connection. In this application subscriber hosts can connect to the 7x50 using a PPPoE tunnel. There are two command available under PPPoE to limit the number of PPPoE hosts - one to set a limit that is applied on each SAP of the group-interface and one to set the limit per group-interface.

PPPoE is commonly used in subscriber DSL networks to provide point-to-point connectivity to subscriber clients running the PPP protocol encapsulated in Ethernet. IP packets are tunneled over PPP using Ethernet ports to provide the client's software or RG the ability to dial into the provider network. Most DSL networks were built with the use of PPPoE clients as a natural upgrade path from using PPP over dial-up connections. Because the PPP packets were used, many of the client software was reusable while enhancements were made such that the client could use an Ethernet port in a similar manner as it did a serial port. The protocol is defined by RFC 2516, *A Method for Transmitting PPP Over Ethernet (PPPoE)*.

PPPoE has two phases, the discovery phase and the session phase.

- **Discovery:** The client identifies the available servers. To complete the phase the client and server must communicate a session-id. During the discovery phase all packets are delivered to the PPPoE control plane (CPM or MDA). The IOM identifies these packets by their ethertype (0x8863).
 - PPPoE Active Discovery Initiation (PADI). This broadcast packet is used by the client to search for an active server (Access Concentrator) providing access to a service.
 - PPPoE Active Discovery Offer (PADO): If the access server can provide the service it should respond with a unicast PADO to signal the client it may request connectivity. Multiple servers may respond and the client may choose a server to connect to.
 - PPPoE Active Discovery Request (PADR): After the client receives a PADO it will use this unicast packet to connect to a server and request service.
 - PPPoE Active Discovery Session-confirmation (PADS) A server may respond to the client with this unicast packet to establish the session and provide the session-id. Once the PADS was provided the PPP phase begins.
- **Session:** Once the session ID is established connectivity is available for the duration of the session, using ethertype 0x8864. Either client or server can terminate a session.

During the life of the session the packets may be uniquely identified by the client's MAC address and session-id. The session can terminate either by PADT sent by the client or server or by an LCP Terminate-Request packet.

During session creation, the following occurs:

- PADI (control packet upstream): This packet is delivered to the control plane. The control plane checks the service tag for service name. In the case multiple nodes are in the same broadcast domain the service tag can be used to decide whether to respond to the client. A relay tag can also be present.
- PADO (control packet downstream): The packet is generated by the control plane as response to the PADI message. The packet is forwarded to the client using the unicast packet directed at the client's MAC address. The node populates the AC-name tag and service tag. The packet sources the forwarding Ethernet MAC address of the node. In the case SRRP is used on the interface, it uses the gateway address as the source MAC. When in a backup state, the packet is not generated.
- PADR (control packet upstream): This packet is delivered to the control plane. The packet is destined to the node's MAC address. The control plane then generates the PADS to create the session for this request.
- PADS: (control packet downstream): The control plane prepares for the session creation and sends it to the client using the client's MAC address. The session-id (16-bit value) is unique per client. The session-id is populated in the response. Once a session-id is generated, the client uses it in all packets. In cases that the server does not agree with the client's populated service tags, the PADS can be used to send a service error tag with a zero session-id to indicate the failure.
- PADT (control packet upstream/downstream): The packet is used to terminate a session. It can be generated by either the control plane or the client. The session-id must be populated. The packet is a unicast packet.
- PPP session creation supports the LCP authentication phase.

During a session, the following forwarding actions occur:

- Upstream, in the PPPoE before PPP phase, there is no anti-spoofing. All packets are sent to the CPM. During anti-spoof lookup with IP and MAC addressing, regular filtering, QoS and routing in context continue. All unicast packets are destined to the node's MAC address. Only control packets (broadcast) are sent to the control plane. Keep-alive packets are handled by the CPM.
- Downstream, packets are matched in the subscriber lookup table. The subscriber information provides queue and filter resources. The subscriber information also provides PPPoE information, such as the dest-mac-address and session-id, to build the packet sent to the client

PPPoE-capable interfaces can be created in a subscriber interface in both IES and VPRN services. Each SAP can support one or more PPPoE sessions depending on the configuration. A SAP can simultaneously have static hosts, DHCP leases and PPPoE sessions. The number of PPPoE sessions is limited per SAP, SLA profile or subscriber profile.

RADIUS can be used for authentication. IP addresses can be provided by both RADIUS and the local IP pool, with the possibility of choosing the IP pool through RADIUS.

DHCP clients and PPPoE clients are allowed on a single SAP or group interface. If DHCP clients are not allowed, the operator should not enable lease-populate and similarly if PPPoE clients are not allowed, the operator should not enable the PPPoE node. Note that the DHCP node can be enabled when only PPPoE clients are allowed since the DHCP relay function can be used for IP retrieval. The DHCP lease-populate is for DHCP leases only. A similar command host-limit is made available under PPPoE for limits on the number of PPPoE hosts. The existing per sla-profile instance host limit is for combined DHCP and PPPoE hosts for that instance.

- For authentication, local and RADIUS are supported.
 - RADIUS is supported through an existing policy. A username attribute has been added.
 - For PAP/CHAP, a local user database is supported and must be referenced from the interface configuration.
- The host configuration can come directly from the local user database or from the RADIUS or DHCP server. A local host configuration is allowed through a local DHCP server with a local user database.
- IP information can be obtained from the local user database, RADIUS, a DHCP server, or a local DHCP server.

If IP information is returned from a DHCP server. PPPoE options such as the DNS name are retrieved from the DHCP ACK and provided to the PPPoE client. An open authentication option is maintained for compatibility with existing DHCP-based infrastructure.

The DHCP server can be configured to run on a loopback address with a relay defined in the subscriber or group interfaces. The DHCP proxy functionality that is provided by the DHCP relay (getting information from RADIUS, lease-split, option 82 rewriting) cannot be used for requests for PPPoE clients.

PPPoE Authentication and Authorization

General Flow

When a new PPPoE session is setup, the authentication policy assigned to the group interface is examined to determine how the session should be authenticated.

If no authentication policy is assigned to the group interface or the **pppoe-access-method** is set to **none**, the local user database assigned to the PPPoE node under the group interface is queried either during the PADI phase or during the LCP authentication phase, depending on whether the match-list of the local user database contains the requirement to match on username. If the match-list does not contain the username option, PADI authentication will be performed and it is possible to specify an authentication policy in the local user database host for an extra RADIUS PAP-CHAP authentication point.

If an authentication policy is assigned and the pppoe-access-method is set to PADI, the RADIUS server will be queried for authenticating the session based on the information available when the PADI packet is received (any PPP user name and password are not known here). When it is set to PAP-CHAP, the RADIUS server will be queried during the LCP authentication phase and the PPP user name and password will be used for authentication instead of the user name and password configured in the authentication policy.

If this authentication is successful, the data returned by RADIUS or the local user database is examined. If no IP address was returned, the DHCP server is now queried for an IP address and possibly other information, such as other DHCP options and ESM strings.

The final step consists of complementing the available information with configured default values (ESM data), after which the host is created if sufficient information is available to instantiate it in subscriber management (at least subscriber ID, subscriber profile, SLA profile, and IP address).

The information that needs to be gathered is divided in three groups, subscriber ID, ESM strings, and IP data. Once one of the data sources has offered data for one of these groups, the other sources are no longer allowed to overwrite this data (except for the default ESM data). For example, if RADIUS provides an SLA profile but no subscriber ID and IP address, the data coming from the DHCP server (either through Python or directly from the DHCP option) can no longer overwrite any ESM string, only the subscriber ID and IP data. However, after the DHCP data is processed, a configured default subscriber profile will be added to the data before instantiating the host.

RADIUS

The following attributes are sent to the RADIUS server for PPPoE authentication (optional attributes can be configured under the **config>subscr-mgmt>auth-plcy>include-radius-attribute** context):

- WT-101 access loop options, DSL-Forum VSAs (optional):
 - Actual data rate Upstream (129)
 - Actual data rate Downstream (130)
 - Minimum data rate Upstream (131)
 - Minimum data rate Downstream (132)
 - Access loop encapsulation (144)
- Circuit-ID, DSL-Forum VSA 1 (optional)
- Remote-ID, DSL-Forum VSA 2 (optional)
- MAC address, Alcatel-Lucent VSA 27 (optional)
- PPPoE-Service-Name, Alcatel-Lucent VSA 35 (optional)
- Port identification attributes (optional)
 - NAS-Port-Id (87) — This attribute can optionally be prefixed by a fixed string and/or suffixed by the circuit-ID or remote-ID given in the PPPoE requests. If either the circuit-ID or remote-ID suffix are given, but the corresponding information is not available, the value 0/0/0/0/0 will be suffixed.
 - NAS-Port-Type (61). Values: 32 (null encap), 33 (dot1q), 34 (qinq), 15 (DHCP hosts), specified value (0 — 255)
- NAS identifier, attribute 32 (optional)
- User-Name, attribute 1
- User-Password, attribute 2
- Service-Type, attribute 6
- Framed-Protocol, attribute 7

In the reply from the RADIUS server, the following authorization attributes are accepted back for PPPoE hosts:

- Calling-station-id (optional)
- Framed-IP-Address, attribute 8
- Session-Timeout, attribute 27
- PADO-delay, Alcatel-Lucent VSA 34
- Framed-Pool, attribute 88 — If a DHCP request is done, this pool will be sent to the DHCP server in a vendor-specific sub-option under Option82 to indicate the pool from which the address from the client should be taken.

- Primary DNS, Alcatel-Lucent VSA 9
- Secondary DNS, Alcatel-Lucent VSA 10
- Primary NBNS, Alcatel-Lucent VSA 29
- Secondary NBNS, Alcatel-Lucent VSA 30
- Subscriber ID string, Alcatel-Lucent VSA 11
- Subscriber profile string, Alcatel-Lucent VSA 12
- SLA profile string, Alcatel-Lucent VSA 13
- ANCP string, Alcatel-Lucent VSA 16
- Intermediate Destination ID, Alcatel-Lucent VSA 28
- Application-profile string, Alcatel-Lucent VSA 45
- Service-Type, attribute 6 (must be correct if returned)
- Framed-Protocol, attribute 7 (must be correct if returned)
- PPPoE-Service-Name, Alcatel-Lucent VSA 35

For more information about Vendor-Specific Attributes and the Alcatel-Lucent dictionary, refer to the Alcatel-Lucent Dictionary in the 7750 SR OS System Management Guide

Local User Database Directly Assigned to PPPoE Node

The following are relevant settings for a local user database directly assigned to PPPoE node:

- Host identification parameters (user name only)
- Username
- Password
- Address
- DNS servers (under DHCP options)
- NBNS servers (under DHCP options)
- Identification (ESM) strings

Incoming PPPoE connections are always authenticated through the PPPoE tree in the local user database.

The matchlist for a local user database that is assigned directly to the PPPoE node under the group interface is always **user-name**, independent of the matchlist setting in the database.

For user-name matching, the incoming user name (user[@domain]) is always first converted to a user and a domain entity by splitting it on the first @-sign. If the no-domain parameter to the user name is given, the user component should be equal to the given user name, if the domain-only portion of the user name is given, the domain entity should be equal to the given user name and if

no extra parameters are given, the user and domain components are concatenated again and compared to the given user name.

The option number for the identification strings is not used if the local user database is assigned directly to the PPPoE node (it is only necessary if it is connected to a local DHCP server). Any valid value may be chosen in this case (if omitted, the default value chosen will be 254).

If a pool name is given for the address, this pool name will be sent to the DHCP server in a vendor-specific sub-option of Option 82 to indicate from which pool the server should take the address. If the gi-address option is given for the address, this will be interpreted as if no address was given.

Local DHCP Server with Local User Database

If a DHCP server is queried for IP or ESM information, the following information is sent in the DHCP request:

- Option 82 sub-options 1 and 2 (Circuit-ID and Remote-ID): These options contain the Circuit-ID and Remote-ID that were inserted as tags in the PPPoE packets).
- Option 82 sub-option 9 vendor-id 6527 VSO 6 (client type): this value is set to 1 to indicate that this information is requested for a PPPoE client. The local DHCP server uses this information to match a host in the list of PPPoE users and not in the DHCP list.
- Option 82 sub-option 6 (Subscriber-ID): This option contains the user name that was used if PAP/CHAP authentication was performed on the PPPoE connection.
- Option 82 sub-option 13 (DHCP pool): This option indicates to the DHCP server that the address from the client should be taken from the specified pool. The local DHCP server will only honor this request if the **use-pool-from-client** option is given in the server configuration.
- Option 82 sub-option 14 (PPPoE Service-Name): This option contains the service name that was used in the PADI packet during PPPoE setup.
- Option 60 (Vendor class ID): This option contains the string “ALU7XXXSBM” to identify the DHCP client vendor.
- The WT-101 access loop options are not sent for PPPoE clients

Local user database settings relevant to PPPoE hosts when their information is retrieved from the local DHCP server using this database:

- Host identification parameters (including user name)
- Address
- DNS servers (under DHCP options)
- NBNS servers (under DHCP options)
- Identification (ESM) strings

For user name matching, the incoming user name (user[@domain]) is always first converted to a user and a domain entity by splitting it on the first @-sign. If the no-domain parameter to the user name is given, the user component should be equal to the given user name, if the domain-only portion of the user name is given, the domain entity should be equal to the given user name and if no extra parameters are given, the user and domain components are concatenated again and compared to the given user name.

To prevent load problems, if DHCP lease times of less than 10 minutes are returned, these will not be accepted by the PPPoE server.

Multiple Sessions Per MAC Address

To support MAC-concentrating network equipment, which translates the original MAC address of a number of subscribers to a single MAC address towards the PPPoE service, the 77x0 supports at most 63 PPPoE sessions per MAC address. Each of these sessions will be identified by a unique combination of MAC address and PPPoE session ID.

To set up multiple sessions per MAC, the following limits should be set sufficiently high:

- Maximum sessions per MAC in the PPPoE policy
- The PPPoE interface session limit in the PPPoE node of the group interface
- The PPPoE SAP session limit in the PPPoE node of the group interface
- The multiple-subscriber-sap limit in the subscriber management node of the SAP

If host information is retrieved from a local DHCP server, care must be taken that, although a host can be identified by MAC address, circuit ID, remote ID or user name, a lease in this server is only indexed by MAC address and circuit ID. For example, multiple sessions per MAC address are only supported in this scenario if every host with the same MAC address has a unique Circuit-ID value.

Private Retail Subnets

PPPoE is commonly used in residential networks and has expanded into business applications. PPPoE session management allows PPPoE to be used for business VPRN access.

PPPoE provides the following:

- Control over the session
- De-muxing based on the session ID provides the ability for the SAP and the Layer 2 network to be common to all VPRNs.

The PPPoE subscriber host will terminate in a retail VPRN and provide a routed path to the customer site. The customer site may be connected to more than one 7750 SR/7710 SR for dual homing purposes. In such a network routing between the two BNGs is required and can be performed with either a direct spoke SDP between the two nodes or by MPBGP route learning.

Since the PPPoE session will terminate in the retail VPRN the node must learn which VPRN service ID will carry it. Both PADI and PAP/CHAP authentication are supported. If the local user database is used, the host configuration provides a reference to the VPRN service ID that will be used. If RADIUS is used, it returns the service ID VSA. The retail interface is determined by the connection between the wholesale and retail subscriber interfaces.

The PPPoE session is negotiated with the parameters defined by the retail VPRN interface. Because the IP address space of the sub-mgmt host may overlap between VPRN services the node must anti-spoof the packets at access ingress with the session-id.

When the **config>service>vprn>sub-if>private-retail-subnets** command is enabled on the subscriber interface, the node will not push the defined subnets in the retail context to the wholesale context. This allows IP overlap between PPPoE sessions. If an operator requires both residential and business services, two VPRNs connected to the same wholesaler can be created and use the flag in only one of them.

Private Retail Subnets

Figure 21 depicts a model with two participants in the service:

- The wholesale VPRN provides access to the SAP.
- The retail VPRN — A business VPRN that is actually routing the packets belonging to the PPPoE sessions terminating in it. The retail VPRNs may have overlapping IP addresses.




Figure 21: PPPoE VPRN with IP Overlap

PPPoE is commonly used in residential networks and has expanded into business applications. Session management can be used as an access technology into a business VPRN.

PPPoE provides the following:

- The PPPoE session is monitored and controlled.
- The session-id provides access de-muxing providing the ability for the SAP and the Layer 2 network to be common.

The PPPoE subscriber host will terminate in a retail VPRN and provide a routed path to the customer site. The customer site may be connected to more than one 7750 SR/7710 SR for dual homing purposes. Routing between the two BNGs is performed with either a direct spoke SDP between the two nodes or by MPBGP route learning.

Since the PPPoE session will terminate in the retail VPRN the node must learn which VPRN service ID will carry it. Both PADI and PAP/CHAP authentication are supported. If the local user database is used, the host configuration provides a reference to the VPRN service ID that will be used. If RADIUS is used, it returns the service ID VSA. The retail interface is determined by the connection between the wholesale and retail subscriber interfaces. RADIUS access from the retail VPRN allows the username/password check to be VPRN specific.

If MSAP is used, the SAP will be created in the wholesale VPRN using the information from RADIUS.

The PPPoE session is negotiated with the parameters defined by the retail VPRN interface.

Because the IP address space of the sub-mgmt host may overlap between VPRN services the node must anti-spoof the packets at access ingress with the session-id.

When the **private-retail-subnets** command is enabled on the subscriber interface, the node will not push the defined subnets in the retail context to the wholesale context. This allows overlap for PPPoE (MAC/sessions ID) and DHCP (MAC/IP) sessions. If an operator requires both residential and business services, two VPRNs connected to the same wholesaler can be created and use the flag in only one of them. This flag can be added and removed from an existing configuration and cannot be a create time parameter. If hosts are defined, they will have to be cleared before the flag is allowed.

PPPoE Command Reference

Configuration Commands

- [PPPoE Policy Configuration Commands on page 437](#)
- [PPPoE Service Commands on page 438](#)
- [PPPoE Local User Database Commands on page 439](#)
- [Show Commands on page 441](#)
- [Clear Commands on page 441](#)
- [Debug Commands on page 441](#)
- [Tools Commands on page 442](#)

PPPoE Policy Configuration Commands

```

config
— subscriber-mgmt
— pppoe-policy pppoe-policy-name [create]
— no pppoe-policy pppoe-policy-name
— description description-string
— no description
— [no] disable-cookies
— keepalive seconds [hold-up-multiplier multiplier]
— no keepalive
— max-sessions-per-mac sessions
— no max-sessions-per-mac
— pado-delay deci-seconds
— no pado-delay
— ppp-authentication {pap | chap | pref-chap}
— no ppp-authentication
— ppp-mtu mtu-bytes
— no ppp-mtu
— ppp-options
— custom-option protocol option-number address ip-address
— custom-option protocol option-number hex hex-string
— custom-option protocol option-number string ascii-string
— no custom-option protocol option-number
— [no] reply-on-padt

```

PPPoE Service Commands

```

config
  — service
    — ies service-id [customer customer-id] [vpn vpn-id] [create]
    — no ies service-id
      — [no] subscriber-interface ip-int-name
        — [no] group-interface ip-int-name
          — dhcp
            — client-applications dhcp
            — client-applications pppoe
            — client-applications dhcp pppoe
            — no client-applications
          — [no] pppoe
            — description description-string
            — no description
            — dhcp-client
              — [no] ccag-use-origin-sap
            — pap-chap-user-db local-user-db-name
            — no pap-chap-user-db
            — pppoe-policy pppoe-policy-name
            — no pppoe-policy
            — sap-session-limit sap-session-limit
            — no sap-session-limit
            — session-limit session-limit
            — no session-limit
            — [no] shutdown

config
  — service
    — vprn service-id [customer customer-id] [create]
    — no vprn service-id
      — subscriber-interface ip-int-name [fwd-service service-id fwd-subscriber-inter-
        face ip-int-name] [create]
      — no subscriber-interface ip-int-name
        — [no] group-interface ip-int-name
          — dhcp
            — client-applications dhcp
            — client-applications pppoe
            — client-applications dhcp pppoe
            — no client-applications
          — [no] pppoe
            — description description-string
            — no description
            — dhcp-client
              — [no] ccag-use-origin-sap
            — pap-chap-user-db local-user-db-name
            — no pap-chap-user-db
            — pppoe-policy pppoe-policy-name
            — no pppoe-policy
            — sap-session-limit sap-session-limit
            — no sap-session-limit
            — session-limit session-limit
            — no session-limit
            — [no] shutdown
      — [no] ip-overlap-allowed

```

PPPoE Local User Database Commands

```

config
— subscriber-mgmt
— local-user-db local-user-db-name [create]
— no local-user-db local-user-db-name
— pppoe
— mask {[prefix-string prefix-string | prefix-length prefix-length] [suffix-
string suffix-string | suffix-length suffix-length]}
— no mask
— host host-name [create]
— no host host-name
— address gi-address
— address ip-address
— address pool pool-name
— no address
— authentication-policy policy-name
— no authentication-policy
— host-identification
— circuit-id string ascii-string
— circuit-id hex hex-string
— no circuit-id
— mac ieee-address
— no mac
— remote-id remote-id
— no remote-id
— service-name service-name
— no service-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
— 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
— l2tp
— group tunnel-group-name
— no group
— 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)]
— no dns-server

```

- **netbios-name-server** **ip-address** [*ip-address...*(up to 4 max)]
- **no netbios-name-server**
- **pado-delay** *deci-seconds*
- **no pado-delay**
- **password** {**ignore** | **chap** *string* | **pap** *string*}
- **no password**
- **retail-service-id** *service-id*
- **no retail-service-id**
- **[no] shutdown**
- **mask type** *pppoe-match-type* {[**prefix-string** *prefix-string* | **prefix-length** *prefix-length*] [**suffix-string** *suffix-string* | **suffix-length** *suffix-length*]}
- **no mask type** *pppoe-match-type*
- **match-list** *pppoe-match-type-1* [*pppoe-match-type-2...*(up to 3 max)]
- **no match-list**
- **[no] shutdown**

Show Commands

```

show
— subscriber-mgmt
— pppoe-policy [pppoe-policy-name [association]]
— service
— id service-id
— pppoe
— session [interface ip-int-name | ip-address | sap sap-id] [session-id session-id] [mac ieee-address] [ip-address ip-address[/mask]] [port port-id] [no-inter-dest-id | inter-dest-id intermediate-destination-id] [detail | statistics]
— statistics [{sap sap-id | interface ip-int-name | ip-address}]
— summary

```

Clear Commands

```

clear
— service
— id service-id
— pppoe
— session all [no-padt]
— session {interface ip-int-name | ip-address | sap sap-id} [mac ieee-address] [session-id session-id] [ip-address ip-address[/mask]] [port port-id] [no-inter-dest-id | inter-dest-id intermediate-destination-id] [no-padt]
— statistics [{sap sap-id | interface ip-int-name | ip-address}]

```

Debug Commands

```

debug
— service
— id service-id
— [no] pppoe
— [no] event
— dhcp-client [terminate-only]
— no dhcp-client
— ppp [terminate-only]
— no ppp
— [no] mac ieee-address
— [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]
— no ppp
— [no] sap sap-id

```

Tools Commands

tools

- **perform**
 - **subscriber-mgmt**
 - **edit-pppoe-session** **sap** *sap-id* **ip** *ip-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*]
 - **edit-pppoe-session** **svc-id** *service-id* **ip** *ip-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*]
 - **eval-lease-state** [**svc-id** *service-id*] [**sap** *sap-id*] [**subscriber** *sub-ident-string*] [**ip** *ip-address*]
 - **local-user-db** *local-user-db-name*
 - **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*]
 - **pppoe**
 - **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*]
 - **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*]

PPPoE Configuration Commands

Global Commands

description

Syntax	description <i>description-string</i> no description
Context	config>subscr-mgmt>pppoe-policy config>service>ies>sub-if>grp-if>pppoe config>service>vprn>sub-if>grp-if>pppoe
Description	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 no form of this command removes the string from the configuration.
Default	No description associated with the configuration context.
Parameters	<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.

shutdown

Syntax	[no] shutdown
Context	config>service>ies>sub-if>grp-if>pppoe config>service>vprn>sub-if>grp-if>pppoe
Description	This command administratively disables an entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics. 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. The no form of this command places the entity into an administratively enabled state.

pppoe-policy

Syntax	pppoe-policy <i>pppoe-policy-name</i> [create] no pppoe-policy <i>pppoe-policy-name</i>
Context	config>subscr-mgmt

PPPoE Configuration Commands

Description	This command configures a PPPoE policy. These policies are referenced from interfaces configured for PPPoE. Multiple PPPoE policies may be configured. This default policy cannot be modified nor deleted.
Default	default
Parameters	<i>pppoe-policy-name</i> — Specifies the PPPoE policy name up to 32 characters in length. create — Keyword used to create the entity. The create keyword requirement can be enabled/disabled in the environment>create context.

disable-cookies

Syntax	[no] disable-cookies
Context	config>subscr-mgmt>pppoe-policy
Description	This command disables the use of cookies. The no form of the command enables cookies.
Default	no disable-cookies

keepalive

Syntax	keepalive seconds [hold-up-multiplier multiplier] no keepalive
Context	config>subscr-mgmt>pppoe-policy>
Description	This command defines the keepalive interval and the number of keepalives that can be missed before the session is declared down for this PPPoE policy. The no form of the command reverts to the default value.
Default	30 seconds 3 multiplier
Parameters	seconds Specifies the keepalive interval in seconds. Values 10 — 300 hold-up-multiplier multiplier — Specifies the number of keepalives that can be missed. Values 1 — 5

max-sessions-per-mac

Syntax	max-sessions-per-mac sessions no max-sessions-per-mac

Context	config>subscr-mgmt>pppoe-policy
Description	This command sets the maximum PPPoE sessions that can be opened for the given MAC address. The no form of the command reverts to the default value.
Default	1
	<i>sessions</i> — Specifies the maximum PPPoE sessions that can be opened for the given MAC address.
Values	1 — 63

pado-delay

Syntax	pado-delay <i>deci-seconds</i> no pado-delay
Context	config>subscr-mgmt>pppoe-policy
Description	This command configures the delay timeout before sending a PPPoE Active Discovery Offer (PADO) packet.
Default	no delay
Parameters	<i>deci-seconds</i> — Specifies the delay timeout before sending a PADO
Values	1 — 30

ppp-authentication

Syntax	ppp-authentication { pap chap pref-chap } no ppp-authentication
Context	config>subscr-mgmt>pppoe-policy
Description	This command configures the PPP protocol used to authenticate the PPPoE session.

ppp-mtu

Syntax	ppp-mtu <i>mtu-bytes</i> no ppp-mtu
Context	config>subscr-mgmt>pppoe-policy
Description	This command configures the the maximum PPP MTU size.
Default	no ppp-mtu
Parameters	<i>mtu-bytes</i> — Specifies the the maximum PPP MTU size.
Values	512 — 9212

ppp-options

Syntax	ppp-options
Context	config>subscr-mgmt>pppoe-policy
Description	This command enables the context to configure PPP options.

custom-option

Syntax	custom-option <i>protocol option-number address ip-address</i> custom-option <i>protocol option-number hex hex-string</i> custom-option <i>protocol option-number string ascii-string</i> no custom-option <i>protocol option-number</i>
Context	config>subscr-mgmt>pppoe-policy>ppp-options
Description	This command provides the ability to configure custom PPP options. Note that standard options such as the DNS name will be returned from DHCP or RADIUS and be converted to PPP automatically. Compression is not supported. The no form of the command removes the custom options from the configuration.
Parameters	<i>protocol</i> — Specifies a protocol for the custom option. Values lcp, ipcp <i>option-number</i> — Assigns an identifying number for the custom option. Values 0 — 255 <i>ip-address</i> — <i>ascii-string</i> — Specifies an ASCII format string for the custom option up to 127 characters long. <i>hex-string</i> — Specifies a hex value for the custom option. Values [0x0..0xFFFFF...(max 254 hex nibbles)]

reply-on-padt

Syntax	[no] reply-on-padt
Context	config>subscr-mgmt>pppoe-policy
Description	Some of the PPPoE clients expect reply on PPPoE Active Discovery Terminate (PADT) message before the context of the session is cleared up. To support such client, a command enabling reply to PADT is provided.
Default	no reply-on-padt

PPPoE Service Commands

pppoe

Syntax	[no] pppoe
Context	config>service>ies>sub-if>grp-if config>service>vprn>sub-if>grp-if
Description	This command configures PPPoE parameters. The no form of the command reverts all PPPoE parameters from the PPPoE context to their defaults.

dhcp-client

Syntax	dhcp-client
Context	config>service>vprn>sub-if>grp-if>pppoe config>service>ies>sub-if>grp-if>pppoe
Description	This command enables the context to configure the PPPoE-to-DHCP options.

ccag-use-origin-sap

Syntax	[no] ccag-use-origin-sap
Context	config>service>vprn>sub-if>grp-if>pppoe config>service>ies>sub-if>grp-if>pppoe
Description	This command enables the original VPLS SAP to be included in the circuit-id option to send to the DHCP server (in case this interface is connected to a VPLS by a CCA MDA). The no form of the command disables the feature.
Default	no ccag-use-origin-sap

pap-chap-user-db

Syntax	pap-chap-user-db <i>local-user-db-name</i> no pap-chap-user-db
Context	config>service>ies>sub-if>grp-if>pppoe config>service>vprn>sub-if>grp-if>pppoe
Description	This command configures the local user database to use for PPP Challenge-Handshake Authentication Protocol/Password Authentication Protocol (PAP/CHAP) authentication.

PPPoE Configuration Commands

If an authentication policy is also configured, **pppoe-access-method** must be set to none in this authentication policy to use the local user database (in that case RADIUS authentication will not be used for PPPoE hosts).

Default	no pap-chap-user-db
Parameters	<i>local-user-db-name</i> — Specifies the local user database to use for authentication.

pppoe-policy

Syntax	pppoe-policy <i>pppoe-policy-name</i> no pppoe-policy
Context	config>service>ies>sub-if>grp-if>pppoe config>service>vprn>sub-if>grp-if>pppoe
Description	This command associates a PPPoE policy on this interface. The no form of the command reverts to the default value.
Default	default
Parameters	<i>pppoe-policy-name</i> — Specifies a a PPPoE policy up to 32 characters in length on this interface.

sap-session-limit

Syntax	sap-session-limit <i>sap-session-limit</i> no sap-session-limit
Context	config>service>ies>sub-if>grp-if>pppoe config>service>vprn>sub-if>grp-if>pppoe
Description	This command specifies the number of PPPoE hosts per SAP allowed for this group-interface.
Default	1
Parameters	<i>sap-session-limit</i> — Specifies the number of PPPoE hosts per SAP allowed. Values 1 — 20000

session-limit

Syntax	session-limit <i>session-limit</i> no session-limit
Context	config>service>ies>sub-if>grp-if>pppoe config>service>vprn>sub-if>grp-if>pppoe
Description	This command specifies the number of PPPoE hosts allowed for this group interface.
Default	1

Parameters *session-limit* — Specifies the number of PPPoE hosts allowed

Values 1 — 20000

ip-overlap-allowed

Syntax [no] ip-overlap-allowed

Context config>service>vprn>sub-if

Description This command allows IP subnet overlap within the wholesale context. The overlap provides means to business services in PPPoE and DHCP. When used with PPPoE, the node will forward between the wholesale and retail contexts based on MAC and session-id. DHCP requires that the MAC address (mac/ip combinations) are unique. The command will fail if the sub-interface is not associated with a wholesale context.

Show Commands

pppoe-policy

Syntax	pppoe-policy [<i>pppoe-policy-name</i> [association]]
Context	show>subscr-mgmt
Description	This command displays PPPoE policy information.
Parameters	<i>pppoe-policy-name</i> — Specifies an existing PPPoE policy association — Displays the object the PPPoE policy is associated.

Sample Output

```
*A:ALA-49>show>subscr-mgmt# pppoe-policy policy1
=====
PPPoE Policy "policy1"
=====
Last Mgmt Change      : 11/16/2003 20:06:39      PPP-mtu              : N/A
Keepalive Interval    : 10s                    Keepalive Multiplier  : 1
Disable AC-Cookies    : No                      PADO Delay            : 0msec
Max Sessions-Per-Mac  : 63                      Reply-On-PADT         : No

-----
PPP Custom Options
-----
Protocol Number Value
-----
No options configured.
=====

*A:ALA-49>show>subscr-mgmt# pppoe-policy policy1 association
=====
PPPoE Policy "policy1"
=====
-----
Interface Associations
-----
Service-Id : 20 (IES)
- grp_pppoe1
- grp_pppoe2
- grp_pppoe3
=====
*A:ALA-49>show>subscr-mgmt#
```

pppoe

Syntax	pppoe
Context	show>service>id
Description	This command enables the context to display PPPoE information.

session

Syntax	session [interface <i>ip-int-name</i> <i>ip-address</i> sap <i>sap-id</i>] [session-id <i>session-id</i>] [mac <i>ieee-address</i>] [ip-address <i>ip-address[/mask]</i>] [port <i>port-id</i>] [no-inter-dest-id inter-dest-id <i>intermediate-destination-id</i>] [detail statistics]
Context	show>service>id>pppoe
Description	This command displays PPPoE session information.
Parameters	interface <i>ip-int-name</i> — <i>ip-address</i> — Displays information about the IP address of the PPPoE session. sap <i>sap-id</i> — Displays information about the specified SAP ID. session-id <i>session-id</i> — Displays information about the ID of the PPPoE session. mac <i>ieee-address</i> — Displays information about the MAC address of the PPPoE session. port <i>port-id</i> — Displays information about about the specified port ID. no-inter-dest-id — inter-dest-id <i>intermediate-destination-id</i> — Displays information about the specified intermediate destination ID. detail — Displays detailed information. statistics — Displays statistics about the PPPoE session.s

Sample Output

```
*A:ALA-49#show service id 20 pppoe session
=====
PPPoE sessions for svc-id 20
=====
Sap Id           Mac Address      Sid Up Time      IP Address
-----
1/1/3:200        00:00:00:00:00:03 1    1d 00:48:39      20.0.0.101
1/1/3:300        00:00:00:00:00:05 1    0d 00:01:08      30.0.0.119
-----
Number of sessions : 2
=====
*A:ALA-49#

*A:ALA-49# show service id 20 pppoe session ip-address 20.0.0.101 detail
=====
```

PPPoE Configuration Commands

```
PPPoE sessions for svc-id 20
=====
Sap Id           Mac Address      Sid Up Time      IP Address
-----
1/1/3:200        00:00:00:00:00:03 1    1d 00:49:46     20.0.0.101

LCP State        : Opened
IPCP State       : Opened
PPP MTU          : 1492
PPP Auth-Protocol : PAP
PPP User-Name    : user4@domain1

Subscriber-interface : sub_pppoe
Group-interface     : grp_pppoe2

Subscriber Origin   : Radius
Strings Origin      : Radius
IPCP Info Origin    : DHCP

Subscriber         : "radius_papchap4"
Sub-Profile-String : "sub1"
SLA-Profile-String : "sla1"
ANCP-String        : ""
Int-Dest-Id        : ""
App-Profile-String : ""

Primary DNS        : N/A
Secondary DNS       : N/A
Primary NBNS       : N/A
Secondary NBNS     : N/A

Circuit-Id        : 2
Remote-Id         :

Session-Timeout    : N/A
-----
Number of sessions : 1
=====
*A:ALA-49#

*A:ALA-49# show service id 20 pppoe session ip-address 20.0.0.101 statistics
=====
PPPoE sessions for svc-id 20
=====
Sap Id           Mac Address      Sid Up Time      IP Address
-----
1/1/3:200        00:00:00:00:00:03 1    1d 00:50:39     20.0.0.101

Packet Type      Received         Transmitted
-----
LCP Configure-Request  1              2
LCP Configure-Ack      1              1
LCP Configure-Nak      1              0
LCP Configure-Reject   0              0
LCP Terminate-Request  0              0
LCP Terminate-Ack      0              0
LCP Code-Reject        0              0
LCP Echo-Request       8927           866
LCP Echo-Reply         866           8927
LCP Protocol-Reject    0              0
```

```

LCP Discard-Request      0          0
-----
PAP Authenticate-Request 1          -
PAP Authenticate-Ack    -          1
PAP Authenticate-Nak    -          0
-----
CHAP Challenge          -          0
CHAP Response           0          -
CHAP Success            -          0
CHAP Failure            -          0
-----
IPCP Configure-Request   2          1
IPCP Configure-Ack       1          1
IPCP Configure-Nak       0          1
IPCP Configure-Reject    0          0
IPCP Terminate-Request   0          0
IPCP Terminate-Ack       0          0
IPCP Code-Reject         0          0
-----
Unknown Protocol         0          -
-----
Number of sessions : 1
=====
*A:ALA-49#

```

statistics

- Syntax** **statistics** [{**sap** *sap-id* | **interface** *ip-int-name* | *ip-address*}]
- Context** show>service>id>pppoe
- Description** This command displays PPPoE statistics.
- Parameters** **sap** *sap-id* — Displays information for the specified SAP. See [Common Service Commands on page 1020](#) for *sap-id* command syntax.
- interface** *ip-int-name* — Displays information about the specified interface.
- ip-address* — Displays information about the specified IP address.

Sample Output

```

*A:ALA-49# show service id 20 pppoe statistics
=====
PPPoE statistics for IES service 20
=====
Packet Type      Received      Transmitted
-----
PADI              2             -
PADO              -             2
PADR              2             -
PADS              -             2
PADT              0             0
session          9838          9839
-----
Drop Counters

```

PPPoE Configuration Commands

```
-----  
Rx Invalid Version      : 0  
Rx Invalid Type         : 0  
Rx Invalid Code         : 0  
Rx Invalid Session      : 0  
Rx Invalid Length       : 0  
Rx Invalid Tags         : 0  
Rx Invalid AC-Cookie    : 0  
Rx Dropped              : 0  
=====
```

```
*A:ALA-49#
```

summary

Syntax	summary
Context	show>service>id>pppoe
Description	This command displays PPPoE summary information

Clear Commands

pppoe

Syntax	pppoe
Context	clear>service>id
Description	This command enables the context to clear PPPoE-related data for the specified service.

session

Syntax	session all [no-padt] session { interface <i>ip-int-name</i> <i>ip-address</i> sap <i>sap-id</i> } [mac <i>ieee-address</i>] [session-id <i>session-id</i>] [ip-address <i>ip-address[/mask]</i>] [port <i>port-id</i>] [no-inter-dest-id inter-dest-id <i>intermediate-destination-id</i>] [no-padt]
Context	clear>service>id>ppoe
Description	This command clears PPPoE sessions.

statistics

Syntax	statistics [{ sap <i>sap-id</i> interface <i>ip-int-name</i> <i>ip-address</i> }]
Context	clear>service>id>ppoe
Description	This command clears PPPoE statistics.

Debug Commands

pppoe

Syntax	[no] pppoe
Context	debug>service>id
Description	This command enables and configures PPPoE debugging.

event

Syntax	[no] event
Context	debug>service>id>pppoe
Description	This command enables debugging for specific PPPoE events.

dhcp-client

Syntax	dhcp-client [terminate-only] no dhcp-client
Context	debug>service>id>pppoe>event
Description	This command enables debugging for specific DHCP client events.

ppp

Syntax	ppp [terminate-only] no ppp
Context	debug>service>id>pppoe>event
Description	This command enables debugging for specific PPP events.

mac

Syntax	[no] mac <i>ieee-address</i>
Context	debug>service>id>pppoe
Description	This command shows PPPoE packets for a particular MAC address.

packet

Syntax	[no] packet
Context	debug>service>id>pppoe
Description	This command enables debugging for specific PPPoE packets.

detail-level

Syntax	detail-level {low medium high} no detail-level
Context	debug>service>id>pppoe>packet
Description	This command configures the PPPoE packet tracing detail level.

dhcp-client

Syntax	[no] dhcp-client
Context	debug>service>id>pppoe>packet
Description	This command enables debugging for specific DHCP client packets.

discovery

Syntax	discovery [padi] [pado] [padr] [pads] [padt] no discovery
Context	debug>service>id>pppoe>packet
Description	This command enables debugging for specific PPPoE discovery packets.

mode

Syntax	mode {dropped-only ingr-and-dropped egr-ingr-and-dropped} no mode
Context	debug>service>id>pppoe>packet
Description	This command configures the PPPoE packet tracing mode.

ppp

Syntax	ppp [lcp] [pap] [chap] [ipcp] no ppp
Context	debug>service>id>pppoe>packet
Description	This command enables debugging for specific PPP packets

sap

Syntax	[no] sap <i>sap-id</i>
Context	debug>service>id>pppoe
Description	This command displays PPPoE packets for a particular SAP.

Tools Commands

tools

Syntax	tools
Context	<root>
Description	The context to enable useful tools for debugging purposes.
Default	none
Parameters	dump — Enables dump tools for the various protocols. perform — Enables tools to perform specific tasks.

perform

Syntax	perform
Context	tools
Description	This command enables the context to enable tools to perform specific tasks.
Default	none

subscriber-mgmt

Syntax	subscriber-mgmt
Context	tools>perform
Description	This command enables tools to control subscriber management.

local-user-db

Syntax	local-user-db <i>local-user-db-name</i>
Context	tools>perform>subscriber-mgmt
Description	This command enables tools for controlling the local user database.
Parameters	<i>local-user-db-name</i> — [32 chars max]

dhcp

Syntax	dhcp
Context	tools>perform>subscriber-mgmt>local-user-db
Description	This command contains the tools used for controlling DHCP entries in the local user database.

host-lookup

Syntax	host-lookup [mac <i>ieee-address</i>] [remote-id <i>remote-id</i>] [sap-id <i>sap-id</i>] [service-id <i>service-id</i>] [string <i>vso-string</i>] [system-id <i>system-id</i>] [option60 <i>hex-string</i>] [circuit-id <i>circuit-id</i> circuit-id-hex <i>circuit-id-hex</i>]
Context	tools>perform>subscriber-mgmt>local-user-db>dhcp
Description	This command performs a lookup in the local user database.
Parameters	mac <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. <i>remote-id</i> — Specifies what information goes into the remote-id sub-option in the DHCP relay packet. Values Up to 255 characters maximum sap-id — Specifies a SAP identifier to be used. <i>service-id</i> — Specifies an existing subscriber service ID. Values 1 — 2147483647 <i>vso-string</i> — Specifies a vendor-specific option string. Values Up to 255 characters maximum <i>system-id</i> — Specifies the system ID. Values Up to 255 characters maximum <i>hex-string</i> — [0x0..0xFFFFFFFF.. (max 64 hex nibbles)] <i>circuit-id</i> — Specifies the circuit-id string. Values Up to 127 characters maximum <i>circuit-id-hex</i> — [0x0..0xFFFFFFFF.. (max 254 hex nibbles)]

pppoe

Syntax	pppoe
Context	tools>perform>subscriber-mgmt>local-user-db

Description This command contains the tools used to control PPPoE entries in the local user database.

authentication

Syntax **authentication** *pppoe-user-name* [**password** *password*]

Context tools>perform>subscriber-mgmt>local-user-db>pppoe

Description This command authenticates PPPoE user name. As local user database PAP/CHAP authentication can only be used when the local user database is connected to the PPPoE node under the group interface, the user lookup will be performed with match-list username.

Parameters *pppoe-user-name* — Specifies the PPPoE username.
password — Specifies the password of this host up to 32 characters in length.

host-lookup

Syntax **host-lookup** [**mac** *ieee-address*] [**remote-id** *remote-id*] [**user-name** *user-name*] [**circuit-id** *circuit-id* | **circuit-id-hex** *circuit-id-hex*]

Context tools>perform>subscriber-mgmt>local-user-db>pppoe

Description This command performs a lookup in the local user database.

Parameters *ieee-address* — xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx
remote-id — [255 chars max]
user-name — Specifies the PPPoE username.
circuit-id — [127 chars max]
circuit-id-hex — [0x0..0xFFFFFFFF.. (max 254 hex nibbles)]

edit-pppoe-session

Syntax **edit-pppoe-session** **sap** *sap-id* **ip** *ip-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*]
edit-pppoe-session **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*]

Context tools>perform>subscriber-mgmt

Description This command modifies PPPoE session information.

Parameters *sap-id*: null [*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]
 dot1q [*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*
 QinQ [*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1.qtag2*

atm	[<i>port-id</i> <i>aps-id</i>][: <i>vpi/vci</i> <i>vpi</i> <i>vpi1.vpi2</i>]
frame	[<i>port-id</i> <i>aps-id</i>]: <i>dlci</i>
cisco-hdlc	<i>slot/mda/port.channel</i>
cem	<i>slot/mda/port.channel</i>
ima-grp	[<i>bundle-id</i>][: <i>vpi/vci</i> <i>vpi</i> <i>vpi1.vpi2</i>]
port-id	<i>slot/mda/port</i> [. <i>channel</i>]
bundle-id	<i>bundle-type-slot/mda.bundle-num</i> bundle keyword <i>type</i> ima, ppp <i>bundle-num1</i> — 256
bpgrp-id	<i>bpgrp-type-bpgrp-num</i> bpgrp keyword <i>type</i> ima, ppp <i>bpgrp-num1</i> — 1280
aps-id	<i>aps-group-id</i> [. <i>channel</i>] aps keyword <i>group-id1</i> — 64
ccag-id	<i>ccag-id.path-id</i> [<i>cc-type</i>]: <i>cc-id</i> ccag keyword <i>id</i> 1 — 8 <i>path-id</i> a, b <i>cc-type</i> .sap-net, .net-sap <i>cc-id</i> 0 — 4094
lag-id	<i>lag-id</i> lag keyword <i>id</i> 1 — 200
qtag1	0 — 4094
qtag2	*, 0 — 4094
vpi	NNI: 0 — 4095 UNI: 0 — 255
vci	1, 2, 5 — 65535
dlci	16 — 1022

ip-address — Displays information for the specified IP address.

sub-ident-string — Displays information for the specified subscriber identification profile.

sub-profile-string — Displays information for the specified subscriber profile.

service-id — Specifies the ID that uniquely identifies a service.

Values 1 — 2147483647

intermediate-destination-id — Specifies the intermediate destination identifier, up to 32 characters in length.

ancp-string *ancp-string* — Specifies the ASCII string of the DSLAM circuit ID name.

app-profile-string — Displays information about the specified application profile.

eval-pppoe-session

Syntax	eval-lease-state [svc-id service-id] [sap sap-id] [subscriber sub-ident-string] [ip ip-address]		
Context	tools>perform>subscriber-mgmt		
Description	This command evaluates lease state.		
Parameters	sap-id:	null	[port-id bundle-id bpgrp-id lag-id aps-id] dot1q [port-id bundle-id bpgrp-id lag-id aps-id]:qtag1 qinq [port-id bundle-id bpgrp-id lag-id]:qtag1.qtag2 atm [port-id aps-id][:vpi/vci vpi vpi1.vpi2] 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] port-id slot/mda/port[.channel] bundle-id bundle-type-slot/mda.bundle-num bundle keyword type ima, ppp bundle-num1 — 256 bpgrp-id bpgrp-type-bpgrp-num bpgrp keyword type ima, ppp bpgrp-num1 — 1280 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 lag-id lag-id lag keyword id 1 — 200 qtag1 0 — 4094 qtag2 *, 0 — 4094 vpi NNI: 0 — 4095 UNI: 0 — 255 vci 1, 2, 5 — 65535 dlci 16 — 1022 ip-address — a.b.c.d sub-ident-string — [32 chars max] service-id — [1..2147483647]

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 7750 SR supports many additional security features, which are described in the 7750 SR OS System Management Guide.

Topics in this chapter include:

- [Triple Play Security Features on page 466](#)
 - [Anti-Spoofing Filters on page 466](#)
 - [Layer 2 Triple Play Security Features on page 469](#)
 - [MAC Pinning on page 469](#)
 - [MAC Protection on page 469](#)
 - [DoS Protection on page 470](#)
 - [VPLS Redirect Policy on page 472](#)
 - [ARP Handling on page 473](#)
 - [Web Portal Redirect on page 475](#)
- [Configuring Triple Play Security with CLI on page 479](#)
- [Common Configuration Tasks on page 480](#)

Triple Play Security Features

Anti-Spoofing Filters

This section describes the Alcatel-Lucent 7750 SR 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, IES SAP-based IP interfaces, and VPRN SAP-based IP interfaces. Anti-spoof filtering is not available on IES or VPRN SDP bound IP interfaces. Anti-spoof filtering is not supported on Epipe and other VLL type 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 603](#)).

Topics in this section are:

- [Anti-spoofing Filter Types on page 467](#)
- [Filtering packets on page 467](#)

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 7750 SR OS 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 469](#)
 - [MAC Protection on page 469](#)
 - [DoS Protection on page 470](#)
 - [VPLS Redirect Policy on page 472](#)
-

MAC Pinning

This section describes the Alcatel-Lucent 7750 SR 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 7750 SR 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:

- **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, 7450 ESS, and 7710 SR 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, 7450 ESS, and 7710 SR 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 7750 SR 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 7750 SR OS Services Guide.

ARP Handling

ARP Reply Agent

This section describes the Alcatel-Lucent 7750 SR 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 7750 SR 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 7750 SR acting as a Broadband Subscriber Aggregator (BSA) with Layer 3 forwarding towards the network.

In an IES or VPRN 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 307](#))), 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 7750 SR acting as a Broadband Subscriber Aggregator (BSA) with Layer 3 forwarding towards the network.

Local proxy ARP allows the Alcatel-Lucent 7750 SR 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 or VPRN 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 7750 SR 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 7750 SR-7 or 7750 SR-12 acting as a Broadband Subscriber Aggregator (BSA). Web portal redirect is not supported on the 7750 SR-1 models.

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 7750 SR 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 22](#) 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 7750 SR intercepts the HTTP GET request and discards it.
- d. The 7750 SR 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 22: IP Illustration of Message Flow in Web Portal Redirect

The items in red text in [Figure 22](#) are messages the 7750 SR will send (masquerading as the destination), regardless of the destination IP address or type of service.

Information needed by the filter that may be sent to the portal:

- Subscriber's IP address
- Subscriber's MAC address
- Original requested URL
- Customer's SAP
- Customer's subscriber identification string (available when Enhanced Subscriber Management is enabled, see section [Triple Play Enhanced Subscriber Management on page 603](#))

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 466](#)).

Since most web sites are accessed using the domain name, the 7750 SR 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 495](#)).

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, IES and VPRN services.

Topics in this section include:

- [Common Configuration Tasks on page 480](#)
 - [Configuring Anti-Spoofing Filters on page 480](#)
 - [Configuring Triple Play Security features on page 481](#)
 - [Configuring ARP Handling on page 487](#)
 - [Configuring Web Portal Redirect on page 495](#)

Common Configuration Tasks

Topics in this section are:

- [Configuring Anti-Spoofing Filters on page 480](#)
 - [Configuring Triple Play Security features on page 481](#)
 - [Configuring MAC Pinning on page 481](#)
 - [Configuring MAC Protection on page 482](#)
 - [Configuring VPLS Redirect Policy on page 484](#)
 - [Configuring VPLS Redirect Policy on page 484](#)
 - [Configuring ARP Handling on page 487](#)
 - [Configuring Web Portal Redirect on page 495](#)
-

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 481](#)
 - [Configuring MAC Protection on page 482](#)
 - [Configuring VPLS Redirect Policy on page 484](#)
 - [Configuring VPLS Redirect Policy on page 484](#)
-

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 485](#)
- [Applying the Filter to a VPLS Service on page 486](#)

Figure 23 displays an IP filter entry configuration for VPLS redirect policy.




Figure 23: VPLS Redirect Policy Example

Information about defining and applying IP and MAC filters is described in the 7750 SR OS 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 487](#)
 - [Configuring Local Proxy ARP on page 488](#)
 - [Configuring ARP Reply Agent in a VPLS Service on page 489](#)
 - [Configuring Automatic ARP Table Population in an IES or VPRN Interface on page 491](#)
-

Configuring Proxy ARP

The implementation of proxy ARP with support for local proxy ARP allows the 7750 SR to respond to ARP requests in the subnet assigned to an IES or VPRN 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 7750 SR 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 or VPRN 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 7750 SR 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 or VPRN Interface

The following example displays the IES DHCP configuration to enable automatic population of the ARP table using snooped DHCP information on an IES or VPRN 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#

A:ALA-1>config>service>vprn>if# info
-----
      dhcp
        description "test"
        lease-populate 1
        no shutdown
      exit
-----
A:ALA-1>config>service>ies>if#
```

Configuring CPU Protection

CPU protection parameters are configured in the **config>sys>security>cpu-protection** context. CPU protection policies allow you to configure parameters the security options according to specifically defined policies. CPU protection policies control of discarding traffic for unconfigured protocols at the CPM network processor as well as rate limits.

CLI Syntax:

```
config>system>security>cpu-protection
  link-specific-rate packet-rate-limit
  policy cpu-protection-policy-id [create]
    alarm
    description description-string
    overall-rate packet-rate-limit
    per-source-rate packet-rate-limit
    port-overall-rate packet-rate-limit
    protocol-protection
```

The following displays CPU protection configuration.

```
A:ALA-48>config>sys>security>cpu-protection# info
-----
-----
A:ALA-48>config>sys>security>cpu-protection>policy#
```

Applying CPU Protection Policies

Refer to the [Subscriber Services Overview](#) section of the 7750 SR OS Services Guide for information about configuring service parameters.

Apply CPU protection policies to the following entities:

- [IES Service on page 493](#)
 - [VPLS Service on page 493](#)
 - [VPRN Service on page 493](#)
 - [Router Interface on page 494](#)
 - [MSAP Policy on page 494](#)
-

IES Service

Use the following CLI to apply a CPU policy to the IES service interface or group interface.

CLI Syntax: `config>service>ies service-id
interface interface-name
cpu-protection policy-id`

VPLS Service

Use the following CLI to apply a CPU policy to the VPLS service SAP.

CLI Syntax: `config>service>vpls service-id
sap sap-id
cpu-protection policy-id`

VPRN Service

Use the following CLI to apply a CPU policy to the VPRN service interface or group interface.

CLI Syntax: `config>service>vprn service-id
interface interface-name
cpu-protection policy-id
subscriber-interface
group-interface
cpu-protection policy-id`

Router Interface

CLI Syntax: `config>router>interface interface-name`
`cpu-protection policy-id`

MSAP Policy

CLI Syntax: `config>subscr-mgmt>msap-policy msap-policy-name`
`cpu-protection policy-id`

Configuring Web Portal Redirect

The generic CLI structure for defining and applying IP and MAC filters is described in the 7750 SR OS 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 7750 SR 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 7750 SR 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 497](#)
- [Layer 2 Security Commands on page 498](#)
- [ARP Handling Commands on page 498](#)
- [CPU-Protection Commands on page 500](#)
- [Security Show Commands on page 501](#)

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 7750 SR OS Services 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
      — subscriber-interface
        — group-interface
          — sap sap-id
            — anti-spoof {ip | 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]
```

```
config
— service
    — vprn
        — [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]
```

CPU-Protection Commands

```

config
  — system
    — security
      — cpu-protection
        — link-specific-rate packet-rate-limit
        — no link-specific-rate
        — policy cpu-protection-policy-id [create]
        — no policy cpu-protection-policy-id
          — [no] alarm
          — description description-string
          — no description
          — overall-rate packet-rate-limit
          — no overall-rate
          — per-source-rate packet-rate-limit
          — no per-source-rate
        — port-overall-rate packet-rate-limit
        — no port-overall-rate
        — [no] protocol-protection
      — router
        — interface interface-name
          — cpu-protection policy-id
          — no cpu-protection
      — service
        — ies
          — interface ip-int-name
            — cpu-protection policy-id
            — no cpu-protection
        — vpls
          — sap sap-id [create]
          — no sap sap-id
            — cpu-protection policy-id [mac-monitoring]
            — no cpu-protection
        — vprn
          — interface ip-int-name
            — cpu-protection policy-id
            — no cpu-protection
          — subscriber-interface
            — group-interface
              — sap sap-id [create]
              — no sap sap-id
                — cpu-protection policy-id [mac-monitoring]
                — no cpu-protection
config
  — subscriber-mgmt
    — msap-policy msap-policy-name [create]
    — no msap-policy msap-policy-name
      — cpu-protection policy-id [mac-monitoring]
      — no cpu-protection

```

Security Show Commands

```
show
  — service
    — id service-id
      — arp
  — system
    — security
      — cpu-protection
        — excessive-sources [service-id service-id sap-id sap-id]
        — policy policy-id [association]
        — protocol-protection
        — violators [port] [interface] [sap]
```

Triple Play Security Configuration Commands

Topics in this section include:

- [Triple Play Security Configuration Commands on page 503](#)
- [Triple Play Anti-Spoofing Commands on page 504](#)
- [Triple Play Layer 2 Security Commands on page 506](#)
- [ARP Handling Commands on page 509](#)
- [Show Commands on page 519](#)

Generic Commands

description

Syntax	description <i>description-string</i> no description
Context	config>sys>security>cpu-protection>policy
Description	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 no form of this command removes the string from the configuration.
Default	No description associated with the configuration context.
Parameters	<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.

Triple Play Anti-Spoofing Commands

anti-spoof

Syntax	anti-spoof {ip mac ip-mac} no anti-spoof
Context	config>service>vpls>sap config>service>ies>if>sap
Description	This command enables anti-spoof filtering and optionally changes the anti-spoof matching type for the SAP. 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 (ip, mac, ip-mac) defines the anti-spoof filter type enforced by the SAP when anti-spoof filtering is enabled. The no form of the command disables anti-spoof filtering on the SAP.
Default	no anti-spoof
Parameters	ip — 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 ip command will fail. mac — 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 anti-spoof mac command will fail. ip-mac — 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 ip-mac command will fail.

anti-spoof

Syntax	anti-spoof {ip ip-mac} no anti-spoof
Context	config>service>ies>subscriber-interface>grp-if>sap
Description	This command enables anti-spoof filtering and optionally change the anti-spoof matching type for the SAP. 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 (ip, ip-mac) defines the anti-spoof filter type enforced by the SAP when anti-spoof filtering is enabled. The no form of the command disables anti-spoof filtering on the SAP.
Default	no anti-spoof
Parameters	ip — 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 ip command will fail.

ip-mac — 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 ip-mac** command will fail.

Triple Play Layer 2 Security Commands

split-horizon-group

Syntax	[no] split-horizon-group [<i>group-name</i>] [<i>residential-group</i>]
Context	config>service>vpls
Description	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. A split horizon group must be created before SAPs and spoke SDPs can be assigned to the group. 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. Up to 30 split horizon groups can be defined per VPLS instance. The no form of the command removes the group name from the configuration.
Parameters	<i>group-name</i> — Specifies the name of the split horizon group to which the SDP belongs. <i>residential-group</i> — Defines a split horizon group as a residential split horizon group (RSHG). Doing so entails that: a) SAPs which are members of this Residential Split Horizon Group will have: – Double-pass queuing at ingress as default setting (can be disabled) – STP disabled (can <u>not</u> be enabled) – ARP reply agent enabled per default (can be disabled) – MAC pinning enabled per default (can be disabled) – Besides the multicast downstream also broadcast packets are discarded thus also blocking the unknown, flooded traffic. b) Spoke SDPs which are members of this Residential Split Horizon Group will have: – Downstream multicast traffic supported – Double-pass queuing is not applicable – STP is disabled (can be enabled) – ARP reply agent is not applicable (dhcp-lease-states are not supported on spoke SDPs) – MAC pinning enabled per default (can be disabled)
Default	A split horizon group is by default not created as a residential-group.

mac-protect

Syntax	mac-protect
Context	config>service>vpls
Description	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.
Default	disabled

mac

Syntax	[no] mac <i>ieee-address</i>
Context	config>service>vpls>mac-protect
Description	This command specifies the 48-bit IEEE 802.3 MAC address.
Parameters	<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.

restrict-protected-src

Syntax	[no] restrict-protected-src
Context	config>service>vpls>split-horizon-group config>service>vpls>sap
Description	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.
Default	no restrict-protected-src

restrict-unprotected-dst

Syntax	[no] restrict-unprotected-dst
Context	config>service>vpls>split-horizon-group config>service>vpls>sap
Description	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.
Default	no restrict-unprotected-dst

mac-pinning

Syntax	[no] mac-pinning
Context	config>service>vpls>sap config>service>vpls>spoke-sdp config>service>vpls>mesh-sdp
Description	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. The age of the MAC address entry in the FIB is set by the age timer. If mac-aging is disabled on a given VPLS service, any MAC address learned on a SAP/SDP with mac-pinning enabled will remain in the FIB on this SAP/SDP forever. Every event that would otherwise result in re-learning will be logged (MAC address; original-SAP; new-SAP). 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.
Default	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

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 host's 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.

arp-populate

Syntax	[no] arp-populate
Context	config>service>ies>interface config>service>vprn>interface config>service>ies>subscriber-interface>group-interface
Description	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 lease-populate), and optionally with static entries entered with the host command. Enabling the arp-populate command will remove any dynamic ARP entries learned on this interface from the ARP cache. The arp-populate command will fail if an existing static ARP entry exists for this interface. The arp-populate command will fail if an existing static subscriber host on the SAP does not have both MAC and IP addresses specified. Once arp-populate is enabled, creating a static subscriber host on the SAP without both an IP address and MAC address will fail. When arp-populate 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 arp-populate enabled. The arp-populate command can only be enabled on IES and VPRN interfaces supporting Ethernet encapsulation. Use the no 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.
Default	not enabled

arp-timeout

Syntax	arp-timeout <i>seconds</i> no arp-timeout
Context	config>service>ies>interface config>service>vprn>interface config>service>ies>subscriber-interface>group-interface
Description	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 **arp-timeout** is set to a value of zero seconds, ARP aging is disabled.

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.

Default	14400 seconds
Parameters	<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.
Values	0 — 65535

authentication-policy

Syntax	authentication-policy <i>name</i> no authentication-policy
Context	config>service>ies>sub-if>grp-if
Description	This command assigns a RADIUS authentication policy to the interface. The no form of this command removes the policy name from the group interface configuration.
Default	no authentication-policy
Parameters	<i>name</i> — Specifies the authentication policy name.

host-connectivity-verify

Syntax	host-connectivity-verify [interval <i>interval</i>] [action {remove alarm}]
Context	config>service>ies>sub-if>grp-if
Description	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 a UC ARP request. The subscriber host connectivity verification will maintain state (connected vs. not-connected) for all hosts.
Default	no host-connectivity-verify
Parameters	interval <i>interval</i> — The interval, expressed in minutes, which specifies the time interval at which all known sources should be verified. The actual rate is then dependent on number of known hosts and interval.
Values	1 — 6000

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

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 and etc.). DHCP release will be signaled to corresponding DHCP server. Static hosts will be never removed. The **alarm** keyword raises an alarm indicating that the host is disconnected.

local-proxy-arp

Syntax	[no] local-proxy-arp
Context	config>service>ies>interface config>service>vpn>interface config>service>ies>subscriber-interface>group-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

remote-proxy-arp

Syntax	[no] remote-proxy-arp
Context	config>service>ies>interface config>service>vpn>interface config>service>ies>subscriber-interface>group-interface
Description	This command enables or disables remote proxy ARP on the interface.
Default	no remote-proxy-arp

proxy-arp-policy

Syntax	proxy-arp-policy <i>policy-name</i> [<i>policy-name...</i> (up to 5 max)] no proxy-arp-policy
Context	config>service>ies>interface config>service>vpn>interface config>service>ies>subscriber-interface>group-interface
Description	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.
Default	none

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.

static-arp

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

Context config>service>ies>interface
 config>service>vprn>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 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 **no** form of the 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.

CPU Protection Commands

System Security

cpu-protection

Syntax	cpu-protection
Context	config>sys>security
Description	This command enables the feature to discard traffic for unconfigured protocols at the CPM network processor. As a result, all received protocol traffic are passed to the CPM CPUs for processing. This command does not change the configured or default rate limits. Note that disabling the protocol protection feature can result in higher CPU loads. The no form of the command reverts to the default behavior to discard unconfigured protocol traffic at the CPM network processor before it is passed to the CPUs.
Default	none

link-specific-rate

Syntax	link-specific-rate <i>packet-rate-limit</i> no link-specific-rate
Context	config>sys>security>cpu-protection
Description	This command configures a link-specific rate for CPU protection. This limit is applied to all interfaces within the system. The CPU will receive no more than the configured packet rate for all link level protocols such as LACP. The measurement is cleared each second and is based on the ingress port.
Default	no link-specific-rate
Parameters	<i>packet-rate-limit</i> — Specifies a packet arrival rate limit for link level protocols.
Values	1 — 65535

policy

Syntax	policy <i>cpu-protection-policy-id</i> [create] no policy <i>cpu-protection-policy-id</i>
Context	config>sys>security>cpu-protection
Description	This command configures CPU protection policies.

The **no** form of the command deletes the specified policy from the configuration.

Default none

Parameters *cpu-protection-policy-id* — Assigns a policy ID to the specific CPU protection policy.

Values 1 — 255

create — Keyword used to create CPU protection policy. The **create** keyword requirement can be enabled/disabled in the **environment>create** context.

alarm

Syntax [no] alarm

Context config>sys>security>cpu-protection>policy

Description This command enables the generation of an event that includes information about the protocol, the offending source and the measured rate. Only one event is generated per monitor period.

The **no** form of the command disables the notifications.

Default alarm

overall-rate

Syntax **overall-rate** *packet-rate-limit*
no overall-rate

Context config>sys>security>cpu-protection>policy

Description This command applies a packet arrival rate limit. The rate defined is a global rate limit for the protocol regardless of how many clients are present.

The **no** form of the command removes the parameter from the configuration.

Default no overall-rate

Parameters *packet-rate-limit* — Specifies the packet arrival rate limit.

Values 1 — 65535

per-source-rate

Syntax **per-source-rate** *packet-rate-limit*
no per-source-rate

Context config>sys>security>cpu-protection>policy

Description This command configures a per-source packet arrival rate limit. Use this command to apply a packet arrival rate limit on a per MAC/subscriber basis. The CPU will receive no more than the configured packet rate for each source. The measurement is cleared each second and is based on the source of the

packet such as SAP/IP/MAC (a known subscriber host for both DHCP and PPPoE), SAP/MAC/circuit-id (for DHCP and PPPoE session establishment), SAP/MAC/session-ID (for PPPoE without circuit-id or after the PADO phase), SAP/MAC.

This parameter is only applicable if the policy is assigned to a subscriber management interface and the **mac-monitor** keyword is specified in the **cpu-protection** command

Default	no per-source-rate
Parameters	<i>packet-rate-limit</i> — Specifies a per-source packet arrival rate limit.
Values	1 — 65535

port-overall-rate

Syntax	port-overall-rate <i>packet-rate-limit</i> no port-overall-rate
Context	config>sys>security>cpu-protection
Description	This command configures a per-port overall rate limit for CPU protection.
Parameters	<i>packet-rate-limit</i> — Specifies a per-port overall rate limit for CPU protection.
Values	1 — 65535, max

protocol-protection

Syntax	[no] protocol-protection
Context	config>sys>security>cpu-protection
Description	This command enables DOS protection of network control protocols.

Router CPU protection

cpu-protection

Syntax	cpu-protection <i>policy-id</i> no cpu-protection
Context	config>router>interface
Description	This command assigns an existing CPU protection policy to the associated IP interface. The CPU protection policies are configured in the config>sys>security>cpu-protection>policy <i>cpu-protection-policy-id</i> context. If no CPU protection policy is assigned to an interface, then a the default policy is used to limit the overall-rate. The default policy is policy number 1 and can be viewed in the show>system>security>cpu-protection>policy 1 command output.
Parameters	<i>policy-id</i> — Specifies an existing CPU protection policy.
Values	1 — 255

Service Interface CPU protection

cpu-protection

Syntax	cpu-protection <i>policy-id</i> no cpu-protection
Context	config>service>ies>if config>service>vprn>if
Description	This command assigns an existing CPU protection policy to the associated service interface. For these interface types, the per-source rate limit is not applicable. The CPU protection policies are configured in the config>sys>security>cpu-protection>policy <i>cpu-protection-policy-id</i> context. If no CPU protection policy is assigned to a service interface, then a the default policy is used to limit the overall-rate. The default policy is policy number 1 and can be viewed in the show>system>security>cpu-protection>policy 1 command output. The no form of the command removes CPU protection policy association from the interface, resulting in no default rate limiting of control packets.
Parameters	<i>policy-id</i> — Specifies an existing CPU protection policy.
Values	1 — 255

cpu-protection

Syntax	cpu-protection <i>policy-id</i> [mac-monitoring] no cpu-protection
Context	config>service>vpls>sap config>service>vprn>sub-if>grp-if>sap config>subscr-mgmt>msap-policy
Description	This command assigns an existing CPU protection policy to the associated service SAP, interface or MSAP policy. The CPU protection policies are configured in the config>sys>security>cpu-protection>policy <i>cpu-protection-policy-id</i> context. If no CPU protection policy is assigned to a service SAP, then a the default policy is used to limit the overall-rate. The default policy is policy number 1 and can be viewed in the show>system>security>cpu-protection>policy 1 command output.
Parameters	<i>policy-id</i> — Specifies an existing CPU protection policy. Values 1 — 255 mac-monitoring — When specified, the per MAC rate limiting should be performed, using the per-source-rate from the associated cpu-protection policy.

Show Commands

arp

Syntax	arp
Context	show>service>id
Description	This command displays the ARP cache entries for this service.
Output	Show All Service-ID Output — 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 static-arp commands. Managed — Learned from DHCP snooping or configured by host 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#
```

excessive-sources

Syntax	excessive-sources [service-id <i>service-id</i> sap-id <i>sap-id</i>]
Context	show>system>security>cpu-protection
Description	This command displays sources exceeding their per-source rate limit.
Parameters	service-id <i>service-id</i> — Displays the specified service ID exceeding the per-source rate limit.

Values 1 — 2147483647

Values: *sap-id*: null [port-id | bundle-id | bpgrp-id | lag-id | aps-id]
dot1q [port-id | bundle-id | bpgrp-id | lag-id | aps-id]:qtag1
qinq [port-id | bundle-id | bpgrp-id | lag-id]:qtag1.qtag2
atm [port-id | aps-id][:vpi/vci|vpi|vpi1.vpi2]
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]
port-id slot/mda/port[.channel]
bundle-id bundle-type-slot/mda.bundle-num
bundle keyword
type ima, ppp
bundle-num 1 — 256
bpgrp-id bpgrp-type-bpgrp-num
bpgrp keyword
type ima, ppp
bpgrp-num 1 — 1280
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
lag-id lag-id
lag keyword
id 1 — 200
qtag1 0 — 4094
qtag2 *, 0 — 4094
vpi NNI: 0 — 4095
UNI: 0 — 255
vci 1, 2, 5 — 65535
dlci 16 — 1022

policy

Syntax	policy <i>policy-id</i> [association]
Context	show>system>security>cpu-protection
Description	This command displays CPU protection policy information.
Parameters	<i>policy-id</i> — Displays information about the specified the CPU protection policy. association — Displays entities associated with specified the CPU protection policy

protocol-protection

Syntax	protocol-protection
Context	show>system>security>cpu-protection
Description	This command displays interfaces where protocol packets are dropped by protocol-protection.

violators

Syntax	violators [port] [interface] [sap]
Context	show>system>security>cpu-protection
Description	This command displays entry points where the overall rate limit is violated.

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 524](#)
- [Multicast in the BSR on page 525](#)
 - [Internet Group Management Protocol on page 525](#)
 - [Source Specific Multicast Groups on page 526](#)
 - [Protocol Independent Multicast Sparse Mode \(PIM-SM\) on page 527](#)
- [Multicast in the BSA on page 529](#)
 - [IGMP Snooping on page 529](#)
- [Configuring Triple Play Multicast Services with CLI on page 535](#)

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) 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 want to leave a multicast group, it will send a leave message.

Multicast in the BSR

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 525](#))
 - Source Specific Multicast Groups ([Internet Group Management Protocol on page 525](#))
 - Protocol Independent Multicast (Sparse Mode) ([PIM-SM on page 527](#))
-

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.

Source Specific Multicast Groups

IGMPv3 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.

The range of multicast addresses from 232.0.0.0 to 232.255.255.255 is currently set aside for source-specific multicast in IPv4. 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.

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 request can be translated into IGMPv3. 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 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

Note that this feature was formally called Ingress Multicast Enhancements and 9G. 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.

To support for on the IOM3-XP, two features control the way multicast traffic traverses the IOM3-XP:

- In the **config>mcast-mgmt>bw-pley** context, the **t2-paths** node allows the configuration of queuing parameters for the primary and secondary multicast paths on IOM3-XP. In addition, the number of secondary paths can also be defined. These configurations parameters are ignored if applied to an IOM-20g-b or IOM2-20g.
- In the **config>mcast-mgmt** context, the **fp** node is affiliated with the multicast path management configuration commands for IOM3-XP ingress multicast management. Ingress multicast management manages multicast switch fabric paths that are forwarding plane specific. It is within this CLI node that the multicast bandwidth policy is associated with the IOM3-XP forwarding plane.
- In addition, the **config>mcast-mgmt>chassis-level** commands configure the multicast plane limit which controls the per switch-fabric multicast planes bandwidth limit for managed multicast traffic. The switch-fabric multicast planes are the individual multicast spatial replication contexts available in the system. The total number of multicast planes depends on a combination of chassis type and chassis mode. The **per-mcast-plane-limit** command applies to all IOM types and is not specific to only IOM3-XP and IMMs.

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 7750 SR, 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 7750 SR implementation performs more than pure snooping of IGMP data, since it also summarizes upstream IGMP reports and responds to downstream queries.

The Alcatel-Lucent 7750 SR 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 Message processing

Figure 24 illustrates the basic IGMP message processing by the 7750 SR in several situations.




Figure 24: IGMP Message Processing

Scenario A: A host joins a multicast group (TV channel) which is not yet being received by other hosts on the 7750 SR, and thus is not yet present in the proxy database. The 7750 SR 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 7750 SR, and thus is already present in the proxy database. No upstream IGMP report is generated by the 7750 SR.

Scenario C: The multicast router will periodically send IGMP queries to the 7750 SR, requesting it to respond with generic membership reports. Upon receiving such a query, the 7750 SR will compile a report from its proxy database and send it back to the multicast router.

In addition, the 7750 SR 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 7750 SR 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 7750 SR 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 7750 SR 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 7750 SR 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.

IGMP 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 7750 SR supports IGMP filtering. Two types of filter can be defined:

- Filter IGMP 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 25: 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 MCAC is not, 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.

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 536](#)
 - [Enabling IGMP Snooping in a VPLS Service on page 536](#)
 - [Modifying IGMP Snooping Parameters on page 538](#)
 - [Configuring Static Multicast Groups on a SAP or SDP on page 541](#)
 - [Enabling IGMP Group Membership Report Filtering on page 542](#)
 - [Enabling IGMP Traffic Filtering on page 544](#)
 - [Configuring Multicast VPLS Registration \(MVR\) on page 545](#)
- [Configuring IGMP and PIM in the BSR on page 546](#)

IGMP

- [Enabling IGMP on page 546](#)
- [Configuring IGMP Interface Parameters on page 547](#)
- [Configuring Static Parameters on page 548](#)
- [Configuring SSM Translation on page 549](#)

PIM

- [Enabling PIM on page 550](#)
- [Configuring PIM Interface Parameters on page 551](#)
- [Importing PIM Join/Register Policies on page 554](#)
- [Configuring PIM Join/Register Policies on page 555](#)
- [Configuring Bootstrap Message Import and Export Policies on page 556](#)

Configuring IGMP Snooping in the BSA

- [Enabling IGMP Snooping in a VPLS Service on page 536](#)
 - [Configuring Static Multicast Groups on a SAP or SDP on page 541](#)
 - [Enabling IGMP Group Membership Report Filtering on page 542](#)
 - [Enabling IGMP Traffic Filtering on page 544](#)
 - [Configuring Multicast VPLS Registration \(MVR\) on page 545](#)
-

Enabling IGMP Snooping in a VPLS Service

- [With IGMPv3 Multicast Routers on page 536](#)
 - [With IGMPv3 Multicast Routers on page 536](#)
 - [With IGMPv1/2 Multicast Routers on page 537](#)
 - [Modifying IGMP Snooping Parameters on page 538](#)
 - [Modifying IGMP Snooping Parameters for a SAP or SDP on page 539](#)
-

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 SR-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 7750 SR OS 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 SR-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 7750 SR OS 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
```

Configuring IGMP and PIM in the BSR

Refer to the Multicast section in the 7750 SR OS Routing Protocols Guide for information about multicast and the commands required to configure basic IGMP and PIM parameters.

IGMP

- [Enabling IGMP on page 546](#)
- [Configuring IGMP Interface Parameters on page 547](#)
- [Configuring Static Parameters on page 548](#)
- [Configuring SSM Translation on page 549](#)

PIM

- [Enabling PIM on page 550](#)
 - [Configuring PIM Interface Parameters on page 551](#)
 - [Importing PIM Join/Register Policies on page 554](#)
 - [Configuring PIM Join/Register Policies on page 555](#)
 - [Configuring Bootstrap Message Import and Export Policies on page 556](#)
-

Enabling IGMP

The following displays an example of enabled IGMP.

```
A:LAX>>config>router# info detail
...
#-----
echo "IGMP Configuration"
#-----
      igmp
        query-interval 125
        query-last-member-interval 1
        query-response-interval 10
        robust-count 2
        no shutdown
      exit
#-----
...
A:LAX>>config>system#
```

Configuring IGMP Interface Parameters

The following example displays an IGMP configuration:

```
A:LAX>config>router>igmp# info
-----
      interface "lax-sjc"
      exit
      interface "lax-vls"
      exit
      interface "pl-ix"
      exit
-----
A:LAX>config>router>igmp# exit
```

Configuring Static Parameters

The following example displays a configuration to add IGMP a static multicast source::

```
A:LAX>config>router>igmp# info
-----
      interface "lax-sjc"
      exit
      interface "lax-vls"
      static
      group 229.255.0.2
      source 172.22.184.197
      exit
      exit
      exit
      interface "pl-ix"
      exit
-----
A:LAX>config>router>igmp#
```

The following example displays the configuration to add a IGMP static starg entry:

```
A:LAX>config>router>igmp# info
-----
      interface "lax-sjc"
      static
      group 230.1.1.1
      starg
      exit
      exit
      exit
      interface "lax-vls"
      static
      group 229.255.0.2
      source 172.22.184.197
      exit
      exit
      exit
      interface "pl-ix"
      exit
-----
A:LAX>config>router>igmp#
```

Configuring SSM Translation

The following displays an SSM translation configuration:

```
A:LAX>config>router>igmp# info
-----
    ssm-translate
      grp-range 229.255.0.1 231.2.2.2
      source 10.1.1.1
    exit
  exit
  interface "lax-sjc"
    static
      group 230.1.1.1
      starg
    exit
  exit
  exit
  interface "lax-vls"
    static
      group 229.255.0.2
      source 172.22.184.197
    exit
  exit
  exit
  interface "pl-ix"
  exit
-----
A:LAX>config>router>igmp# exit
```

Configuring PIM

Enabling PIM

When configuring PIM, make sure to enable PIM on all interfaces for the routing instance, otherwise multicast routing errors can occur.

The following example displays detailed output when PIM is enabled.

```
A:LAX>>config>router# info detail
...
#-----
echo "PIM Configuration"
#-----
    pim
    no import join-policy
    no import register-policy
    apply-to none
    rp
        no bootstrap-import
        no bootstrap-export
        static
        exit
        bsr-candidate
            shutdown
            priority 0
            hash-mask-len 30
            no address
        exit
        rp-candidate
            shutdown
            no address
            holdtime 150
            priority 192
        exit
    exit
    no shutdown
exit
#-----
...
A:LAX>>config>system#
```

Configuring PIM Interface Parameters

The following displays a PIM interface configuration:

```
A:LAX>config>router>pim# info
-----
      interface "system"
      exit
      interface "lax-vls"
      exit
      interface "lax-sjc"
      exit
      interface "pl-ix"
      exit
      rp
      static
        address 2.22.187.237
        group-prefix 224.24.24.24/32
      exit
      address 10.10.10.10
      exit
      exit
      bsr-candidate
      shutdown
      exit
      rp-candidate
      shutdown
      exit
      exit
```

```
A:LAX>config>router>pim#
```

```
A:SJC>config>router>pim# info
-----
      interface "system"
      exit
      interface "sjc-lax"
      exit
      interface "sjc-nyc"
      exit
      interface "sjc-sfo"
      exit
      rp
      static
        address 2.22.187.237
        group-prefix 224.24.24.24/32
      exit
      exit
      bsr-candidate
      shutdown
      exit
      rp-candidate
      shutdown
      exit
      exit
```

```
A:SJC>config>router>pim#
```

Configuring Triple Play Multicast Services with CLI

```
A:MV>config>router>pim# info
-----
interface "system"
exit
interface "mv-sfo"
exit
interface "mv-vlc"
exit
interface "p3-ix"
exit
rp
  static
    address 2.22.187.237
    group-prefix 224.24.24.24/32
  exit
exit
bsr-candidate
  address 2.22.187.236
  no shutdown
exit
rp-candidate
  address 2.22.187.236
  no shutdown
exit
exit
-----
```

```
A:MV>config>router>pim#
```

```
A:SFO>config>router>pim# info
-----
interface "system"
exit
interface "sfo-sjc"
exit
interface "sfo-was"
exit
interface "sfo-mv"
exit
rp
  static
    address 2.22.187.237
    group-prefix 224.24.24.24/32
  exit
exit
bsr-candidate
  address 2.22.187.239
  no shutdown
exit
rp-candidate
  address 2.22.187.239
  no shutdown
exit
exit
-----
```

```
A:SFO>config>router>pim#
```



```
A:WAS>config>router>pim# info
-----
      interface "system"
      exit
      interface "was-sfo"
      exit
      interface "was-vlc"
      exit
      interface "p4-ix"
      exit
      rp
        static
          address 2.22.187.237
          group-prefix 224.24.24.24/32
          exit
        exit
        bsr-candidate
          address 2.22.187.240
          no shutdown
        exit
        rp-candidate
          address 2.22.187.240
          no shutdown
        exit
      exit
-----
A:WAS>config>router>pim#
```

Importing PIM Join/Register Policies

The import command provides a mechanism to control the (*,g) and (s,g) state that gets created on a router. Import policies are defined in the **config>router>policy-options** context. See [Configuring PIM Join/Register Policies on page 555](#).

Note, in the import policy, if an action is not specified in the entry then the default-action takes precedence. If no entry matches then the default-action also takes precedence. If no default-action is specified, then the default default-action is executed.

The following example displays the command usage to apply the policy statement will not allow join messages for group 229.50.50.208/32 and source 192.168.0.0/16 but allows join messages for 192.168.0.0/16, 229.50.50.208:

Example:

```
config>router# pim
config>router>pim# import join-policy "foo"
config>router>pim# no shutdown
```

The following example displays the PIM configuration:

```
A:LAX>config>router>pim# info
-----
import join-policy "foo"
interface "system"
exit
interface "lax-vls"
exit
interface "lax-sjc"
exit
interface "p1-ix"
exit
rp
  static
    address 2.22.187.237
      group-prefix 224.24.24.24/3
    exit
    address 10.10.10.10
    exit
  exit
  bsr-candidate
    shutdown
  exit
  rp-candidate
    shutdown
  exit
exit
-----
A:LAX>config>router>pim#
```

Configuring PIM Join/Register Policies

Join policies are used in Protocol Independent Multicast (PIM) configurations to prevent the transportation of multicast traffic across a network and the dropping of packets at a scope at the edge of the network. PIM Join filters reduce the potential for denial of service (DoS) attacks and PIM state explosion—large numbers of Joins forwarded to each router on the RPT, resulting in memory consumption.

*g or s,g is the information used to forward unicast or multicast packets.

- **group-address** matches the group in join/prune messages
group-address 229.55.150.208/32 exact
- **source-address** matches the source in join/prune messages
source-address 192.168.0.0/16 longer
- **interface** matches any join message received on the specified interface
interface port 1/1/1
- **neighbor** matches any join message received from the specified neighbor
neighbor 1.1.1.1

The following configuration example will not allow join messages for group 229.50.50.208/32 and source 192.168.0.0/16 but allows join messages for 192.168.0.0/16, 229.50.50.208.

```
A:ALA-B>config>router>policy-options# info
-----
...
    policy-statement "foo"
        entry 10
            from
                group-address "229.50.50.208/32"
                source-address 192.168.0.0
            exit
            action reject
        exit
    exit
    policy-statement "reg-pol"
        entry 10
            from
                group-address "224.0.0.0/8"
            exit
            action accept
        exit
    exit
exit
...
-----
A:ALA-B>config>router>policy-options#
```

Configuring Bootstrap Message Import and Export Policies

Bootstrap import and export policies are used to control the flow of bootstrap messages to and from the RP.

The following configuration example specifies that no BSR messages received or sent out of interface port 1/1/1.

```
:A:ALA-B>config>router>policy-options# policy-statement pim-import
:A:ALA-B>config>router>policy-options>policy-statement$ entry 10
:A:ALA-B>config>router>policy-options>policy-statement>entry$ from
:A:ALA-B>config>router>policy-options>policy-statement>entry>from$ interface port1/1/1/
:A:ALA-B>config>router>policy-options>policy-statement>entry>from$ exit
:A:ALA-B>config>router>policy-options>policy-statement>entry# action reject
:A:ALA-B>config>router>policy-options>policy-statement>entry# exit
:A:ALA-B>config>router>policy-options>policy-statement# exit

:A:ALA-B>config>router>policy-options# policy-statement pim-export
:A:ALA-B>config>router>policy-options>policy-statement$ entry 10
:A:ALA-B>config>router>policy-options>policy-statement>entry$ to
:A:ALA-B>config>router>policy-options>policy-statement>entry>to$
```

Triple Play Multicast Command Reference

Command Hierarchies

- [MCAST Mangement Commands on page 557](#)
- [Multicast Info Policy Bundle Commands on page 559](#)
- [Triple Play Multicast Service Commands on page 561](#)
- [Ingress Multicast Path Management Commands on page 564](#)
- [Show Commands on page 566](#)
- [Clear Commands on page 567](#)
- [Debug Commands on page 567](#)

MCAST Mangement 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
        — 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**
- **chassis-level**
 - **[no] per-mcast-plane-limit**
 - **per-mcast-plane-limit** *megabits-per-second* [**secondary** *megabits-per-second*] [**dual-sfm** *megabits-per-second*] [**secondary-dual-sfm** *megabits-per-second*]

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-p2mp lsp-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-p2mp lsp-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-p2mp** *lsp-name* **sender** *ip-address*
- **no primary-tunnel-interface**
- **description** *description-string*
- **no description**

Refer to the 7750 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
          — 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
        — mld-snooping
          — 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] 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-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-address
            — [no] starg
            — mld-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
        — 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
    — [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] stargmld-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
  — 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 percent-of-resv-cbs
— no cbs
— hi-priority-only percent-of-mbs
— no hi-priority-only
— mbs percent-of-pool
— 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 percent-of-resv-cbs
— no cbs
— hi-priority-only percent-of-mbs
— no hi-priority-only
— mbs percent-of-pool
— no mbs
— secondary-path
— path-limit megabits-per-second
— no path-limit
— queue-parameters
— cbs percent-of-resv-cbs
— no cbs
— hi-priority-only percent-of-mbs
— no hi-priority-only
— mbs percent-of-pool

config
— router
— multicast-info-policy policy-name
— no multicast-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]
```

Show Commands

```

show
  — service
    — id service-id
      — igmp-snooping
        — base
        — all
        — 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]]
      — mld-snooping
        — all
        — base
        — mrouters [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]

```

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
        — 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
```

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

Syntax	description <i>description-string</i> no description
Context	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
Description	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 context in the configuration file. The no form of this command removes any description string from the context.
Default	No description is associated with the configuration context.
Parameters	<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

Syntax	[no] shutdown
Context	config>card>mda>ingress>mcast-mgmt config>card>mda>ingress>mcast-path-management
Description	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. 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

Syntax	mcast-management
Context	config
Description	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. 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. A chassis-level node within mulicast-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.

Bandwidth Policy Commands

bandwidth-policy

Syntax	bandwidth-policy <i>policy-name</i> [create] no bandwidth-policy <i>policy-name</i>
Context	config>mcast-mgmt
Description	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.

Multicast Management Configuration Commands

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 though 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

Syntax **admin-bw-threshold** *kilo-bits-per-second*
no admin-bw-threshold

Context config>mcast-management>bandwidth-policy

Description 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.

Values 1 — 40,000,000

Default 10

primary-path

Syntax	primary-path
Context	config>mcast-mgmt>bandwidth-policy
Description	This command enables the context to configure primary path parameters.

ancillary-path

Syntax	ancillary-path
Context	config>mcast-mgmt>bandwidth-policy
Description	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

Syntax	t2-paths
Context	config>mcast-management>bandwidth-policy
Description	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

Syntax	secondary-path
Context	config>mcast-mgmt>bandwidth-policy config>mcast-mgmt>bw-plcy>t2-paths
Description	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

Syntax	number-paths <i>number-of-paths</i> [dual-sfm <i>number-of-paths</i>] config>mcast-management>bandwidth-policy>t2-paths>secondary-paths>queue-parameters
Description	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 reciprocal number of primary paths is 15 to 1 resulting in a total of 16 multicast 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 high priority only percentage value.

Parameters	<i>number-of-paths</i> — The number-of-paths parameter is required when specifying hi-priority-only and is expressed as an integer between 0 and 100. The value 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. A value of 100 specifies that the LP-MBS threshold will be set to 0, resulting in a discard of all low congestion priority packets.

Values 0 — 100

Default 10

path-limit

Syntax	path-limit <i>megabits-per-second</i> no path-limit
Context	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
Description	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. 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. 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. 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. The no 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.
Default	no path-limit
Parameters	<i>megabits-per-second</i> — Specifies the primary path limit in megabits per seconds.
Values	1 — 5000
Default	primary-path: 2,000 secondary-path 1,500 ancillary-path: 5,000

queue-parameters

Syntax	queue-parameters
Context	<code>config>mcast-management>bandwidth-policy>primary-path</code> <code>config>mcast-management>bandwidth-policy>secondary-path</code> <code>config>mcast-management>bandwidth-policy>ancillary-path</code> <code>config>mcast-management>bandwidth-policy>t2-paths</code>
Description	This command defines the individual parameters for the queues through which multicast packets are forwarded into the switch fabric on each path. 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. 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. 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. 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 multicast-info-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. 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.

cbs

Syntax	cbs <i>percent-of-resv-cbs</i> no cbs
Context	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
Description	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. The cbs percent-of-resv-cbs 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. 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 cbs percent-of-resv-cbs 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. The no form of the command restores the path queues' default CBS value.
Parameters	<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). Values 0 — 100 Default Primary: 5 Secondary: 30 Ancillary: 65

hi-priority-only

Syntax	hi-priority-only <i>percent-of-mbs</i> no hi-priority-only
Context	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
Description	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. The no form of this command restores the default value.
Parameters	<i>percent-of-mbs</i> — The percentage parameter is required when specifying high-prior-only.
Values	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

Syntax	mbs <i>percent-of-pool</i> no mbs
Context	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
Description	This command is used to override the default Maximum Buffer Size (MBS) for each individual path in a queue. The queue's 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 queue's high-prior-only parameter.

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.

Parameters	<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.		
	Values	0 — 100	
	Default	Primary Default:	7
		Secondary Default:	40
		Ancillary Default:	80
	Default	10	

falling-percent-reset

Syntax	falling-percent-reset <i>percent-of-highest</i> no falling-percent-reset
Context	config>mcast-mgmt>bw-plcy
Description	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.

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.

Multicast Management Configuration Commands

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.

By default falling-percent-reset is disabled when a new bandwidth policy is created. The default bandwidth policy has a hard configured value of 50%. Setting falling-percent-reset to 100 is equivalent to specifying no falling-percent-reset.

The **no** form of the command restores the default value of 100% which disables the falling percent reset function.

Parameters	<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.
Values	1 — 100 percent
Default	100

mcast-pool

Syntax	mcast-pool percent-of-total percent-of-buffers resv-cbs percent-of-pool slope-policy policy-name no mcast-pool
Context	config>mcast-mgmt>bw-plcy
Description	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. (If path management is not enabled, the system does not allocate buffers to the multicast pool.)

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

Syntax	multicast-info-policy <i>policy-name</i> [create] no multicast-info-policy <i>policy-name</i>
Context	config>mcast-management
Description	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 egress multicast CAC function may use the policy information as a basis for allowing or disallowing downstream nodes to join multicast streams. 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 Egress Multicast CAC

The egress multicast CAC feature has its own multicast CAC policy that is applied to egress IP interfaces or egress VPLS interfaces. The policy contains bundles, each with their own sets of channel ranges defined. When a multicast join event occurs on the interface, the system searches the multicast CAC policy to determine how that join event should be processed. The information returned from the CAC lookup provides the bundles allowed bandwidth and the channels administrative bandwidth. Since the allowed bundle bandwidth may change between differing egress interfaces, multiple policies with the same channel information may be needed.

With the addition of the multicast information policy, managing the CAC feature is simplified. The CAC monitor for the egress interface first searches the multicast CAC policy to determine if the multicast join event matches a configured channel range. If a match is found, it simply uses the local policy information. If a match is not found, it then searches the multicast information policy associated with the service or routing context to which the join event is associated. The multicast information policy returns the following information to the interfaces CAC manager:

- Bundle name
- Administrative bandwidth ('0' if undefined)
- Congestion Priority Threshold (high or low)
- CAC Type (mandatory or optional)

The CAC manager evaluates the returned results according to the following rules:

- If the returned administrative bandwidth = '0', all results are ignored
- If the returned bundle name is not found in the local multicast CAC policy, all results are ignored
- The administrative bandwidth is interpreted as channel 'bw'
- A value of 'high' for congestion priority threshold is interpreted as 'class high'
- A value of 'low' for congestion priority threshold is interpreted as 'class low'
- A value of 'mandatory' for CAC type is interpreted as 'type mandatory'
- A value of 'optional' for CAC type is interpreted as 'type optional'
- Bundle bandwidth is always derived from the local multicast CAC policy

Multicast Management Configuration Commands

Using the multicast information policy to store the CAC information allows a single centralized managed policy for all channel information, allowing the multicast CAC policies to only have bundle defined with the appropriate bundle bandwidth. The multicast CAC policy still may be for channel information in exception cases.

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:

8. Administrative bandwidth ('0' if undefined)
9. 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

Syntax **multicast-info-policy** *policy-name*
no multicast-info-policy

Context config>service>vpls
config>router

Description 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.

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 <i>bundle-name</i> [create] no bundle <i>bundle-name</i>
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 <i>bundle-name</i> — 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 <i>ip-address</i> [<i>ip-address</i>] [create] no channel <i>ip-address</i> [<i>ip-address</i>]
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. [What do we do with 224.0.0.x addresses?] 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

Syntax	admin-bw <i>kbps</i> no admin-bw
Context	config>mcast-mgmt>mcast-info-plcy>bundle>channel config>mcast-mgmt>mcast-info-plcy>bundle>source-override

Description	This command specifies an administrative bandwidth for multicast channels. The specified bandwidth rate may be used by the multicast ingress path manger, multicast CAC manager or multicast ECMP manager. 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 bw-activity use-admin-bw 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. The no form of the command restores the default value for admin-bw. If the command is executed in the channel context, the channels administrative bandwidth value is set to null. If the command is executed in the source-override context, the source override administrative bandwidth value is set to null.
Default	none
Parameters	<i>kbps</i> — Specifies the administrative bandwidth for multicast channels. Values 1 — 40000000 kbps 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.

bw-activity

Syntax	bw-activity {use-admin-bw dynamic [falling-delay seconds]} [black-hole-rate kbps] no bw-activity
Context	config>mcast-mgmt>mcast-info-plcy>bundle>channel config>mcast-mgmt>mcast-info-plcy>bundle>source-override
Description	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 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.

Multicast Management Configuration Commands

The **no** form of the command restores the default bandwidth activity monitoring setting (dynamic or null depending on the context).

Default	no bw-activity
Parameters	use-admin-bw dynamic — The use-admin-bw and dynamic keywords are mutually exclusive and one must be specified when executing the bw-activity command. The use-admin-bw keyword indicates the channels current ingress bandwidth should be derived from the admin-bw setting. The admin-bw setting must not currently be set to 0 for the use-admin-bw setting to succeed. The dynamic 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. falling-delay seconds — 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. Values 10 — 3600 Default 30 Bundle default: dynamic Channel default: Null (undefined) Source-override default: Null (undefined) black-hole-rate kbps — 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. Values 0 — 40000000 Default None

cac-type

Syntax	cac-type {mandatory optional}
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 defines the channels multicast CAC channel type, either mandatory or optional . 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 the channels multicast CAC channel type. optional — This keyword is mutually exclusive with the mandatory keyword and specifies the 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. <table><tr><td>Default</td><td>Bundle :</td><td>None (no explicit-sf-path)</td></tr><tr><td></td><td>channel:</td><td>Null (undefined)</td></tr><tr><td></td><td>Source-override</td><td>Null (undefined)</td></tr></table> Override sequence — The channel setting overrides the bundle setting. The source-override setting overrides the channel and bundle settings.	Default	Bundle :	None (no explicit-sf-path)		channel:	Null (undefined)		Source-override	Null (undefined)
Default	Bundle :	None (no explicit-sf-path)								
	channel:	Null (undefined)								
	Source-override	Null (undefined)								

keepalive-override

Syntax	keepalive-override <i>keepalive-timer</i> no keepalive-override
Context	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

Description This command configures the keepalive timer override.

Parameters *keepalive-timer* — Specifies the keepalive timer override, in seconds.

Values 10 — 300

preference

Syntax **preference** *preference-level*
no preference

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 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.

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.

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.

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.

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.

Values

primary-tunnel-interface

Syntax	primary-tunnel-interface rsvp-p2mp <i>lsp-name</i> sender <i>ip-address</i> 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 <i>lsp-name</i> — Species a string of up to 32 characters identifying the LSP name as configured at the ingress LER. sender <i>ip-address</i> — Specifies a string of 15 characters representing the IP address of the ingress LER for the LSP.

source-override

Syntax	source-override <i>ip-address</i> [create] no source-override <i>ip-address</i>
Context	config>mcast-mgmt>mcast-info-plcy>bundle>channel

Multicast Management Configuration Commands

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	<i>ip-address</i> — 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 <i>preference-level</i> 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 <i>low</i> class and channels with a preference equal to or greater than the threshold will be treated as <i>high</i> 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.
Default	4

Parameters	<i>preference-level</i> — The preference-level parameter is required when specifying the cong-priority-threshold.
Values	0 — 7
Default	4

ecmp-opt-threshold

Syntax	ecmp-opt-threshold <i>preference-level</i> no ecmp-opt-threshold
Context	config>mcast-mgmt>mcast-info-plcy>bundle
Description	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-theshold, 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. 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. Changing the theshold causes all channels ECMP optimization eligibility to be reevaluated. The no form of the command restores the default ECMP optimization preference threshold value.
Parameters	<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. Values 0 — 7 Default 7

mc-ecmp-balance

Syntax	[no] mc-ecmp-balance
Context	config>router>pim
Description	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. 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. 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. In order to achieve a proper balance, multicast groups and their bandwidth should be configured in the config mcast-management context.

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.

If **mc-ecmp-balance** is enabled, a redistribution may be triggered whenever a interface is added to an ECMP link.

If **mc-ecmp-balance** 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.

If **mc-ecmp-balance** is enabled, then a threshold (**ecmp-opt-threshold**) can be configured to avoid moving multicast streams where interruption should be avoided.

The **ecmp-opt-threshold** 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 **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.

Changing the **ecmp-opt-threshold** causes all channels ECMP optimization eligibility to be reevaluated.

The **no** form of the command removes the re-balancing capability from the configuration.

mc-ecmp-balance-hold

Syntax	mc-ecmp-balance-hold <i>minutes</i> no mc-ecmp-balance-hold
Context	config>router>pim
Description	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. 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. This parameter should be used to avoid excessive changes to the multicast distribution. A rebalance will not occur to multicast streams that have a priority greater then the configured ecmp-opt-threshold. The no form of the command reverts to default.
Parameters	<i>minutes</i> — Specifies the hold down time in minutes.
Values	2 — 600

mc-ecmp-rebalance

Syntax	mc-ecmp-rebalance [ecmp-opt-threshold]
Context	tools>perform>router>pim
Description	This command triggers an immediate rebalance, regardless if the hold timer has triggered or if any changes have occurred.

Parameters **ecmp-opt-treshold** — 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 **ecmp-opt-threshold** value.

Values 1 — 7

Chassis Level Commands

chassis-level

Syntax	chassis-level
Context	config>mcast-management
Description	The chassis-level CLI node contains the multicast plane replication limit for each switch fabric multicast plane. 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.

mmrp-imp-override

Syntax	[no] mmrp-imp-override
Context	config>mcast-mgmt>chassis-level
Description	This command enables ingress Multicast Path Management (IMPM) from monitoring PIM and IGMP. The no form of the command disables the IMPM monitoring.
Default	no mmrp-imp-override

per-mcast-plane-limit

Syntax	per-mcast-plane-limit <megabits-per-second> [secondary<megabits-per-second>] [dual-sfm <megabits-per-second> [secondary-dual-sfm<megabits-per-second>]] no per-mcast-plane-limit
Context	config>mcast-mgmt>chassis-level
Description	This command configures the overall multicast (primary plus secondary) and specific secondary rate limits for each switch fabric multicast plane. The number of multicast planes available depends on the chassis type, the configured chassis mode (b, c, or d) and the installed switch fabric modules. Table 8 lists the number of available multicast planes per chassis type.

Table 8: Multicast Planes Based on Chassis Type

Chassis Type	Switch Fabric Type Number Installed	Chassis Mode	Multicast Planes Available
ESS-6	80G x 1	B or C	7
		D	14

Table 8: Multicast Planes Based on Chassis Type

Chassis Type	Switch Fabric Type Number Installed	Chassis Mode	Multicast Planes Available
ESS-7 / SR-7	80G x 2	B or C	14
		D	28
	200G/400G x 1	B or C	8
		D	24
ESS-12 / SR-12	200G/400G x 2	B or C	16
		D	48
	200G x 1	B or C	8
		D	24
	200G x 2	B or C	16
		D	48
	400G x 1	B or C	8
		D	24
	400G x 2	B or C	16
		D	48

Chassis multicast planes should not be confused with IOM multicast paths. The IOM uses multicast paths to reach multicast planes on the switch fabric. An IOM may have less or more multicast paths than the number of multicast planes available in the chassis.

Each IOM 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 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-limit command) at which time either flows are moved to other paths or potentially black-holed (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-limit command 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-limit` should be used to compensate for the non-multicast or secondary multicast switch fabric traffic.

The **no** form of the command restores the default multicast plane bandwidth limits for the chassis.

Parameters *megabits-per-second* — The megabits-per-second parameter directly following the `per-mcast-plane-limit` command is required and defines the overall multicast rate limit supported individually by all switch fabric multicast planes. If the optional `dual-sfm` keyword is not specified, the rate applies when either one or two switch fabric modules are installed. The parameter is defined as a whole integer value representing megabits per second.

Values 1 — 2,000

secondary *megabits-per-second* — Limits the amount of secondary multicast traffic allowed for each individual switch fabric multicast plane. The megabits-per-second parameter following the `secondary` keyword is required and is specified as a whole integer representing rate in megabits per second. As multicast bandwidth is mapped to a multicast plane, primary multicast flows consume] the overall multicast plane limit while secondary flows consume both overall and secondary plane limits. The specified megabits-per-second rate must be equal to or less than the configured overall rate. If the secondary limit is equal to the overall limit, the system removes the secondary limit.

dual-sfm *megabits-per-second* [*secondary megabits-per-second*] — The `dual-sfm` keyword is used to specify an alternate set of overall and secondary multicast plane rate limits when two switch fabric modules are installed. When specified, the first defined rate directly following `per-mcast-plane-limit` is used when only a single switch fabric module is installed. For ESS-6, 7 & 12 and SR-7 & 12 chassis, the switch fabric module is integrated into the CPM (Control Plane Module). The `secondary` keyword is optional when specifying dual-sfm multicast limits and is used to limit the aggregate rate for secondary multicast flows associated with a dual switch fabric module system.

secondary *megabits-per-second* — The `secondary` keyword following `dual-sfm` is optional and is used to limit the amount of secondary multicast traffic allowed for each individual switch fabric multicast plane when two switch fabric modules are installed and operational. The megabits-per-second parameter following the `secondary` keyword is required and is specified as a whole integer representing rate in megabits per second. As multicast bandwidth is mapped to a multicast plane, primary multicast flows consume only the overall multicast plane limit while secondary flows consume both overall and secondary plane limits. The specified megabits-per-second rate must be equal to or less than the configured overall rate. If the secondary limit is equal to the overall limit, the system removes the secondary limit.

Forwarding Plane Commands

fp

Syntax	fp [<i>fp-number</i>]				
Context	config>card				
Description	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. 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. 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.				
Parameters	<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. <table> <tr> <td>Values</td><td>1</td></tr> <tr> <td>Default</td><td>1</td></tr> </table>	Values	1	Default	1
Values	1				
Default	1				

ingress

Syntax	ingress
Context	config>card>fp
Description	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.

multicast-path-management

Syntax	multicast-path-management
Context	config>card>fp>ingress config>card>mda>ingress
Description	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

Syntax	bandwidth-policy <i>policy-name</i> no bandwidth-policy				
Context	config>card>fp>ingress>mcast-path-management config>card>mda>ingress>mcast-path-management				
Description	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. 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. The no form of the command removes an explicit bandwidth policy from a forwarding plane or MDA and restores the default bandwidth policy.				
Parameters	<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. <table><tr><td>Values</td><td>Any existing bandwidth policy name</td></tr><tr><td>Default</td><td>default</td></tr></table>	Values	Any existing bandwidth policy name	Default	default
Values	Any existing bandwidth policy name				
Default	default				

primary-override

Syntax	primary-override
Context	config>card>mda>ingress>mcast-mgmt
Description	This command enables the context to configure MDA ingress multicast path-limit overrides. The path override CLI nodes are not supported on IOM-3.

secondary-override

Syntax	secondary-override
Context	config>card>mda>ingress>mcast-mgmt
Description	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

Syntax	ancillary-override
Context	config>card>mda>ingress>mcast-mgmt
Description	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	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. The path-limit commands are not supported on IOM-3. 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.		
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		

Triple Play Enhanced Subscriber Management

In This Section

This section describes 7750 SR features which provide Enhanced Subscriber Management functions for Triple Play services.

Topics in this section include:

- [RADIUS Authentication of Subscriber Sessions on page 605](#)
- [RADIUS Authentication Extensions on page 607](#)
- [Enhanced Subscriber Management Overview on page 626](#)
 - [Enhanced Subscriber Management Basics on page 626](#)
 - [Dynamic Subscriber Host Processing on page 637](#)
 - [ESM for IPv6 on page 633](#)
 - [Enhanced Subscriber Management Entities on page 640](#)
 - [ANCP and GSMP on page 643](#)
 - [Using Scripts for Dynamic Recognition of Subscribers on page 648](#)
 - [Limiting Subscribers and Hosts on a SAP on page 655](#)
 - [Static Subscriber Hosts on page 655](#)
 - [QoS for Subscribers and Hosts on page 656](#)
 - [Layer 3 Subscriber Interfaces on page 662](#)
 - [Multi-Chassis Synchronization on page 657](#)
 - [Routed Central Office \(CO\) on page 660](#)
 - [Linked Subscriber Interface on page 674](#)
 - [Hub-and-Spoke Forwarding on page 675](#)
 - [Dual Homing on page 677](#)
 - [Local and Remote Dual Master Detected on page 666](#)

- [Interaction Notes on page 667](#)
 - [Managed SAP \(M-SAP\) on page 701](#)
 - [Multicast Management on page 704](#)
- [Configuring Enhanced Subscriber Management with CLI on page 711](#)
 - [Configuring RADIUS Authentication of DHCP Sessions on page 712](#)
 - [Configuring Enhanced Subscriber Management on page 713](#)
 - [Service Manager Scripting Language on page 1007](#)
 - [Sample Python Scripts on page 1015](#)

RADIUS Authentication of Subscriber Sessions

This section describes the Alcatel-Lucent 7750 SR 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 302](#) for an explanation of DHCP and [DHCP Snooping on page 307](#) for an explanation of DHCP snooping.

For information about the radius server selection algorithm, refer to the Security chapter in the 7750 SR OS System Management Guide.

If authentication is enabled, the 7750 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-----

```

RADIUS Authentication of Subscriber Sessions

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.

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




Figure 26: Triple Play Aggregation Network with RADIUS-Based DHCP Host Authentication

Figure 26 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 7750's subscriber management in TPSDA, the RADIUS client running on 7750 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 7750 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 7750 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, 7750 SR supports 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 7750 SR 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**.

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*.

Provisioning of Enhanced Subscriber Management (ESM) Objects

In the ESM concept on 7750 SR 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.

RADIUS Authentication of Subscriber Sessions

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.
- session-timeout — Defines the lease time.
- 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 7750 SR network element serves as a RADIUS client, which means the messages are originated by a 7750 SR. 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.

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) 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 sever will send CoA messages to the 7750 network element requesting the change in attributes included in the CoA request:

- framed-ip-address — This attribute is mandatory for any CoA message.
- nas-port-id — This attribute is mandatory for any CoA message.

Note that the framed-ip-address and nas-port-id are only used for identification of the subscriber and cannot be changed.

- alc-subscriber-profile-string
- alc-sla-profile-string
- alc-ancp-string
- alc-app-profile-string
- alc-int-dest-id-string
- alc-subscriber-id-string

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, 7750 SR 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:

- framed-ip-address— This attribute is mandatory for any CoA message.
- NAS-port-id— This attribute is mandatory for any CoA message.
- 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 is mandatory for new subscriber host creation any CoA message. It indicates the SAP where the host should be created.
- framed-ip-address — This attribute is mandatory for any CoA message.
- alc-client-hw-address — A string in the xx:xx:xx:xx:xx:xx format. This attribute is mandatory for new subscriber-host creation.
- session-timeout — If not present, a default of 7 days will be assigned.
- alc-retail-svc-id — This is only used in case of wholesaling for selection of the retail service
- 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), 7750 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:

- Support of CoA messages is disabled by default. They can be enabled on a per authentication-policy basis.
- When CoA is enabled, the 7750 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.
- A 7750 will first try to find an appropriate service (and therefore authentication policy) for an incoming CoA, and then see if the server is defined within this policy and whether the shared secret is correct.

In all cases (creation, modification, forcere-new) attributes nas-port-id and framed-ip-address are mandatory for CoA request as these are used as the key for subscriber-host identification. In case these two attributes are not present in CoA, the message will be NAK'd with corresponding error code.

- CoA messages received from the given RADIUS server affect only applicable subscriber-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 7750 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 7750 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 7750 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 configuration commands. Following RADIUS attributes will be supported in accounting messages:

- framed-ip-address
- framed-ip-netmask
- agent-circuit-id (as defined by DSL forum)
- agent-remote-id (as defined by DSL forum)
- calling-station-id
- alc-subscriber-id-string
- alc-subscriber-profile-string
- alc-sla-profile-string
- user-name
- NAS-identifier
- NAS-port-id
- NAS-port-type — Values: 32 (null encap), 33 (dot1q), 34 (qinq), 15 (DHCP hosts), specified value (0 — 255)
- MAC-address

In the 7750 subscriber management concept, a single subscriber-host can have multiple associated queues. Every queue represents a separate set of accounting statistics and therefore, the single accounting message will contain statistics for all queues associated with the given subscriber host.

RADIUS accounting defines accounting statistics in terms of the following attributes which contain the queue-id (the first 2 bytes of a 10B word) and the counter (the remaining 8B):

- alc-acct-input-inprof-octets-64 — ingress-in-profile-forwarded-bytes.

- alc-acct-input-outprof-octets-64 — ingress-out-of-profile-forwarded-bytes.
- alc-acct-input-inprof-packets-64 — ingress-in-profile-forwarded-packets.
- alc-acct-input-outprof-packets-64 — ingress-out-of-profile-forwarded-packets.
- alc-acct-output-inprof-octets-64 — egress-in-profile-forwarded-bytes.
- alc-acct-output-outprof-octets-64 — egress-out-of-profile-forwarded-bytes.
- alc-acct-output-inprof-packets-64 — egress-in-profile-forwarded-packets.
- alc-acct-output-outprof-packets-64 — egress-out-of-profile-forwarded-packets.

In addition to accounting-start, interim-accounting, and accounting-stop messages, a RADIUS client on a 7750 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

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 7750 SR OS System Management Guide⁷ for further information.

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.

RADIUS Accounting AVP	Include-radius-attrs Acct/ Auth	Host Accounting	SLA-Profile Accounting
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

RADIUS Authentication of Subscriber Sessions

RADIUS Accounting AVP	Include-radius-attrs Acct/ Auth	Host Accounting	SLA-Profile Accounting
Acct-Session-Id	No/No	Yes	Yes
Acct-Session-Time	No/No	Yes	Yes
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

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 7750 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.

Enhanced Subscriber Management Overview

Topics in this section:

- [Enhanced Subscriber Management Basics on page 626](#)
 - [Using Scripts for Dynamic Recognition of Subscribers on page 648](#)
 - [Limiting Subscribers and Hosts on a SAP on page 655](#)
 - [QoS for Subscribers and Hosts on page 656](#)
-

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 7750 SR 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 7750 SR 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 7750 SR 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.

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 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 628](#)), first, the 7750 SR 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 7750 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 7750 SR) 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 7750 SR. 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 648](#).

Subscriber Identification String

Subscribers are managed by the 7750 SR 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 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 652](#) 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.

IPv6 DNS server information should be requested separately from the ESR by sending a DHCPv6 Information-Request.

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.

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.

Dynamic Subscriber Host Processing

Dynamic Tables

To support all processing for Enhanced Subscriber Management, several tables are maintained in the 7750 SR (Figure 27).

- [Active Subscriber Table on page 638](#)
- [SLA Profile Instance Table on page 638](#)
- [Subscriber Host Table on page 638](#)
- [DHCP Lease State Table on page 639](#)



OSSG084

Figure 27: 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 631](#)).
 - In use subscriber profiles (see [Subscriber Profile on page 631](#)).
-

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 28 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).




Figure 28: 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.

ANCP and GSMP

- [ANCP on page 643](#)
 - [General Switch Management Protocol Version 3 \(GSMPv3\) on page 646](#)
-

ANCP

Access Node Control Protocol Management (ANCP) can provide the following information to the 7750 SR:

- ANCP can communicate the current access line rate to the 7750. 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 7750 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.

Static ANCP Management

As depicted in [Figure 29](#), 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 7750 to provide rate management.




Figure 29: Static ANCP Management Example

Enhanced Subscriber Management (ESM) Dynamic ANCP

In this application ANCP is used between the DSLAM and the 7750 SR to provide line control. There are multiple attributes defined as described below. [Figure 30](#) 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 30: 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.

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, 7750 SR 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 fulfills 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 301](#) for an explanation of DHCP and DHCP snooping in the 7750.

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 1015](#). Additional information about the service manager scripting language, see [Service Manager Scripting Language on page 1007](#).

One or more scripts can be written by the operator and stored centrally on a server (in a location accessible by the 7750). 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 31 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 31: 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 32: 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
 - queue mapping
- The IP filter to use.

The SLA profile also has a host-limit attribute which limits the total 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 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 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.

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 33: 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.

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.
- The SLA profile refers to specific queue/policer settings for each host (BTV, VoIP, PC).
- The subscriber profile also refers to CFHP ingress and egress policer-control-policies which define the overall treatment for hosts of this subscriber.

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.

Multi-Chassis Synchronization

Figure 34 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 34: 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.

Routed Central Office (CO)

The Routed CO feature allows a network operator to connect a DSLAM to a 7750 SR. [Figure 35](#) shows a DSLAM connected to a 7750 SR using a Layer 3 interface within an IES 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 35: 7750 SR DSLAM Connection

In early releases of the 7750 SR OS, only one SAP could be associated with each Layer 3 interface. Now 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 defined the subscriber subnets, the group interfaces are responsible for aggregating the SAPs.

As depicted in [Figure 36](#), an operator can create a new subscriber interface in the IES 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 37](#) shows the details of group interface A.




Figure 36: Subscriber Interface in an IES Service




Figure 37: Details of a Group-Interface

Layer 3 Subscriber Interfaces

This section describes the Alcatel-Lucent 7750 SR acting as a Broadband Service Router (BSR), with Enhanced Subscriber Management enabled.

In this model, a 7750SR is positioned directly behind a DSLAM. This design removes the need for a Layer 2 aggregation network between the 7750 SR and the DSLAM, however it does involve more routing entities in the network.

Figure 38 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 referred to as BBNG, Broadband Network Gateway, in the DSL Forum.





Figure 38: 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 39: Detailed View of Configurable Objects Related to Layer 3 Subscriber Interfaces

Figure 39 shows a detailed view of a 7750 SR 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 Figure 39) 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.

Enhanced Subscriber Management Overview

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 (see [Triple Play Enhanced Subscriber Management on page 603](#)).

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.

Routed Subscriber Hosts

In order to support forwarding of packets addressing hosts in this subnet (the subnet behind the actual host), an anti-spoof look-up (nh-mac) is used. 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.

There are two ways to associated routes with any given routed-subscriber, BGP peering and RADIUS-based assignment. BGP peering can be set up automatically (if enabled) every time a routed-subscriber is created. The parameters for the BGP peering come from bgp-peering policy and from RADIUS VSAs in authentication-response message. The ip-address of the routed-subscriber will be used as bgp-peer address. The associated routes are then learned from the BGP peering.

Ingress and egress ESM functionality is enforced on the same basis.

Subscriber hosts on a subscriber interface SAP with the nh-mac anti-spoof type can have one or more associated routes (these are referred to as managed routes). There is no limitation on subnet length (/32 routes are supported).

In case of a RADIUS-based assignment, following rules are observed:

- When the framed-route attribute in RADIUS response is absent or in a DHCP-ACK, the corresponding host will be installed as standalone host. Only traffic addressing host's address will be forwarded.
- Metrics contained in the RADIUS framed-route attribute will be ignored.
- The route included in the framed route attribute will be accepted as a managed route only if it points to the host's address as a next-hop or it contains 0.0.0.0 as next-hop address. In the latter case, the host's IP address will be used as the next-hop ip-address. If neither is applicable, the framed-route attribute is ignored and host is installed without the ill-defined route.
- An overlap between managed routes associated with different subscriber hosts in the context of the same IES/VPNRN service is not supported. In this case, the new host configuration with managed-route falling into the range of existing managed-route will be refused. Alarm and logging will accompany this event.
- A frame-routed attribute associated with the host that is not associated with a SAP with “nh-mac” anti-spoof type, will be ignored and the host will be installed as stand-alone host without any managed routed.
- Managed routes can be redistributed through corresponding routing protocols. Managed-routes are subject to policy statements. Filtering is based on the subnet/mask.

This feature is supported only on Layer 3 interfaces (IES and VPRN subscriber and group interfaces). It is also supported at network level in a combination with the managed SAP feature which allows assigning individual subscriber SAPs to different VPRNs.

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 will 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.

Interaction Notes

- [Routed CO for VPRN Interactions on page 667](#)
 - [DHCP Interactions on page 669](#)
 - [Routed CO Interactions on page 669](#)
-

Routed CO for VPRN Interactions

Much like the Routed CO model, 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 16 subscriber subnets and can be configured with a host address for each subnet. The wholesaler subscriber interface can contain addresses that are used with its own end-customers. 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 either the wholesaler's or retailer's 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 will require Enhanced Subscriber Management. The subscriber interface's subnets are allowed to be advertised to both IGP and BGP within a VPRN.

If a packet is received by a wholesaler instance sourced from a subscriber host it will be forwarded based on the routing table of the retailer's VPRN. The destination IP address can be part of the same retailer subnet of another local interface (subscriber or regular interface). If the destination IP address is part of a different VPRN (the destination subnet is not available in the retailer's VPRN) the packet will be dropped. However, if the destination IP address is of a subscriber that is within the same retailer's VPRN the packet can be delivered. When a subscriber-split-horizon is enabled, the system will divide the retailer VPRN into two VRFs. The upstream VRF will contain routes other than the subscriber interface while the downstream VRF will contain all routes.

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.

In the retailer's VPRN, DHCP settings (RADIUS and DHCP server assignments) are allowed. The entire DHCP node is allowed only for subscriber interfaces that are associated with a wholesaler context. RADIUS authentication is available to the retailer's subscriber interface (when the interface is defined to be linked to another).

The node can be defined with both a DHCP relay or proxy function in each routing instance. If the user configures a DHCP relay within the instance, the local-proxy-server command will enable

DHCP split leases. Note that the local DHCP proxy in this case can be different than the one defined in the wholesale context. 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.

The system does not allow subscriber interface chaining. A retailer subscriber interface can be associated with only one wholesaler subscriber interface. The interface cannot be associated with another subscriber interface that is referencing an interface already. Further, once a retailer subscriber interface is created and linked with a wholesale subscriber interface all subscriber interfaces in that retailer's VPRN cannot be used to forward (cannot be linked) to other VPRNs.

Because the subscriber interface in the retailer context does not use a group-interface, some group-interface DHCP functions are allowed. This functionality is available to a retailer subscriber interface only (an interface that was configured to reference another).

Note that the wholesaler DHCP settings can be configured with a vendor-specific option to send the service ID to the server. The service ID that is used must be the retailer's service ID returned from RADIUS. If the option is used in the retailer's context, the retailer's service ID will always be used.

Lease-populate must be enabled for dynamic hosts in the context of the wholesaler VPRN. In that context, the limit must include all hosts for all retailer's hosts seen. A second limit (identified as lease-populate-limit) can be defined in the context of the retailer to limit the number of leases that are allowed to be associated with its subscribers.

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 Interactions

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 sixteen 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.

VPRN Routed COs

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 interface. 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.

VPRN Wholesaler-Retail Routed CO

In the wholesaler/retailer model (see [Figure 40](#)), the wholesaler instance connections that are common to the access nodes are distributed to many retail instances. 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 other direction. The wholesale/retail traffic flow is controlled with minimal communication with RADIUS. A RADIUS policy is defined in the wholesaler instance. The RADIUS response used during the subscriber instantiation provides the service context of the retailer VPRN. If the wholesaler has a retail business, the operator can configure a separate VPRN for their retail services.

The retailer's subscriber interface controls the DHCP configuration wherever possible instead of the wholesaler's group interface DHCP configuration. The only exception is the lease-populate value. The lease-populate value in the wholesale context controls the per-SAP limits. The lease-populate value in the retail subscriber interface controls the limits for that retailer's interface. Both limits must be satisfied before a new subscriber can be instantiated.




Figure 40: Routed CO Example

As previously stated, in the wholesaler instance, the flow of traffic is possible with minimal communication with RADIUS. A RADIUS policy defined in the wholesaler instance provides the service context of the retailer VPRN.




Figure 41: Traffic Flow

Figure 41 shows the process starting with a DHCP DISCOVER message and, optionally, with a renew message.

- The packet enters the wholesaler instance which defines the wholesaler RADIUS authentication policy.
- The node contacts RADIUS to authorize access. If RADIUS denies access, the DHCP packet is discarded.
- If access is granted, RADIUS returns the service selection VSA.
- If RADIUS does not return this VSA but authenticates the host, service is allowed only in the wholesaler router instance. In this case, processing RADIUS and DHCP continues within the routing context configured.
- The RADIUS server can return all the attributes for the host such as the service ID, policy and IP address information. The node then provides the IP address to the client using the local DHCP proxy defined in the VPRN and install the host. If the wholesaler provides the IP configuration through RADIUS, proxy DHCP on the retailer must be enabled.

- The RADIUS server can optionally return the host's policy definition. If RADIUS replies with only the VPRN, the policy definition processing will continue by the DHCP configuration of the retailer VPRN. The packet is relayed using the VPRN DHCP configuration and the host will be installed after the ACK is received.
-

Linked Subscriber Interface

In a wholesale model that is using multiple VPRNs it is necessary to use two subscriber interfaces. The primary subscriber interface will be configured in the context of the wholesaler. This interface is defined with group interface parenting SAPs defined. A second linked subscriber interface must be defined in the retailer context. This linked subscriber interface defines the subnet and potentially the DHCP and RADIUS behavior of the subscribers in that context.

Since that linked subscriber interface does not parent any group interfaces or SAPs it cannot directly forward. It must be linked to another interface used for forwarding. A subscriber interface can be linked to a single subscriber interface and context. Multiple subscriber interfaces can be associated with a single primary subscriber interface in a 1:N (primary sub-int:linked sub-int) configuration. After a subscriber interface was linked to a context all other subscriber interfaces must be linked to the same context.

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. The new flag is "type wholesale-retail" which 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. When the flag is enabled, routes learned from MBGP, IGP through a regular interface, static routes through regular interfaces and locally attached regular interface routes will be considered HUB routes and allowed to be used for upstream traffic. Subscriber subnets cannot be used for upstream traffic. Each packet is examined against the HUB routes when the flag is set to eliminate subscriber to subscriber direct forwarding.

Figure 42 describes the forwarding model without a "type wholesale-retail" flag. Each packet is forwarded based on the VPRN (same as the VRF) Z forwarding table. Any packet that is destined to a known subscriber host will be routed directly.




Figure 42: Forwarding Model Without Wholesale-Retail Type Flag

Figure 43 describes forwarding following the “type wholesale-retail” flag. Packets coming from the subscriber interface (SAP is defined in the wholesale VPRN) will be routed based on VRF Z2 routes only. Each packet coming from an LSP or Interfaces B or C will be routes based on routes from both VRF Z1 and Z2. This forwarding model is sometimes referred to as subscriber-split-horizon.




Figure 43: Forwarding Model With Wholesale-Retail Type Flag

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 44: Dual-Homing to Two PEs

Figure 44 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 45](#)).

[Figure 45](#) 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 45: 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 45](#) (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 45 on page 679](#)). 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/VRPN 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 46 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 46: Dual Homing Ring Under Steady-State Condition

In [Figure 46](#), the ring is fully closed and every access node has two possible paths towards the VPLS core. [Figure 46](#) 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 47](#) 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 680](#). For all hosts behind the unreachable ring nodes, the corresponding subscriber host FDB entries point to the shunt SDP.




Fig_38

Figure 47: 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 680](#), 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 48 depicts a single DSLAM dual-homed to two BSRs.




Figure 48: Ring Topology Using Layer 3 CO Model

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 Router Redundancy Protocol on page 692](#).

MC Services

The typical implementation of MC services at the network level is shown in [Figure 49](#).




Figure 49: 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 49](#)). 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 50](#).

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 50: 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 51](#).

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 Router Redundancy Protocol on page 692](#).

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.




Figure 51: Dual Homing Example

Synchronization

To establish subscriber state the nodes must synchronize subscriber information. Refer to the 7750 SR OS 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.

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.

Subscriber Router Redundancy Protocol

Subscriber Router Redundancy Protocol (SRRP) is closely tied to the multi-chassis synchronization protocol used to synchronize information between redundant nodes. SRRP is valid only in the 7750 SR-12, SR-7, and 7710 SR models. 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). Out-of-band is performed by Multi-Chassis Synchronization (MCS). 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 2.2 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 will be considered:

- Associate the SRRP instance to a Multi-Chassis Synchronization (MCS) peering terminating on the neighboring router (the MCS peering will 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 9](#) lists the SRRP's state effect on subscriber hosts associated with group IP interfaces.

Table 9: 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	• Will respond to ARP for all owned subscriber subnet IP addresses. • Will not respond to ARP for subscriber subnet SRRP gateway IP addresses. • All ARP responses will contain the native MAC of the group IP interface (not the SRRP gateway MAC).	• Will respond to ARP for all subscriber hosts on the subscriber subnet.	• Will respond to ARP for all reachable remote IP hosts.	• All routing out the group IP interface will use the native group IP interface MAC address. • The group IP interface redundant IP interface will not be used. • Will not accept packets destined to the SRRP gateway MAC received on the group IP interface.

Table 9: SRRP State Effect on Subscriber Hosts Associated with Group IP Interface (Continued)

SRRP State	ARP	Local Proxy ARP Enabled	Remote Proxy ARP Enabled	Subscriber Host Routing
Becoming Master (In order to enter becoming master state, a master must currently exist)	- Will respond to ARP for all owned subscriber subnet IP addresses (hardware address and source MAC = group IP interface native MAC). - Will respond to ARP for subscriber subnet SRRP gateway IP addresses (hardware address = SRRP gateway IP address, source MAC = group IP interface native MAC).	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).	Will respond to ARP for all reachable remote IP hosts (hardware address = SRRP gateway MAC address, source MAC = group IP interface native MAC).	- All routing out the group IP interface will use the native group IP interface MAC address. - Subscriber hosts mapped to the redundant IP interface are remapped to the group IP interface. - Will not accept packets destined to the SRRP gateway MAC received on the group IP interface.
Master	- Will respond to ARP for all owned subscriber subnet IP addresses (hardware address and source MAC = group IP interface native MAC). - Will respond to ARP for subscriber subnet SRRP gateway IP addresses (hardware address = SRRP gateway IP address, source MAC = group IP interface native MAC).	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).	Will respond to ARP for all reachable remote IP hosts (hardware address = SRRP gateway MAC address, source MAC = group IP interface native MAC).	- All routing out the group IP interface will use the SRRP gateway MAC address. - Subscriber hosts mapped to the redundant IP interface are remapped to the group IP interface. - Will accept packets destined to the SRRP gateway MAC received on the group IP interface.

Table 9: SRRP State Effect on Subscriber Hosts Associated with Group IP Interface (Continued)

SRRP State	ARP	Local Proxy ARP Enabled	Remote Proxy ARP Enabled	Subscriber Host Routing
Becoming Backup (redundant IP interface operational)	• Will respond to ARP for all owned subscriber subnet IP addresses (hardware address and source MAC = group IP interface native MAC). • Will not respond to ARP for subscriber subnet SRRP gateway IP addresses	• Will not respond to ARP for any subscriber hosts on the subscriber subnet.	• Will not respond to ARP for any remote IP hosts.	• Will not route out the group IP interface for subscriber hosts associated with the subscriber subnet. • Subscriber hosts mapped to the group IP interface are remapped to the redundant IP interface. • Will accept packets destined to the SRRP gateway MAC received on the group IP interface.
Becoming Backup (redundant IP interface not available)	• Will respond to ARP for all owned subscriber subnet IP addresses (hardware address and source MAC = group IP interface native MAC). • Will not respond to ARP for subscriber subnet SRRP gateway IP addresses.	• Will not respond to ARP for any subscriber hosts on the subscriber subnet.	• Will not respond to ARP for any remote IP hosts.	• Will route out the group IP interface for subscriber hosts associated with the subscriber subnet using the group IP interface native MAC address. • Subscriber hosts mapped to the redundant IP interface are remapped to the group IP interface. • Will accept packets destined to the SRRP gateway MAC received on the group IP interface

Table 9: SRRP State Effect on Subscriber Hosts Associated with Group IP Interface (Continued)

SRRP State	ARP	Local Proxy ARP Enabled	Remote Proxy ARP Enabled	Subscriber Host Routing
Backup (redundant IP interface operational)	• Will respond to ARP for all owned subscriber subnet IP addresses (hardware address and source MAC = group IP interface native MAC). • Will not respond to ARP for subscriber subnet SRRP gateway IP addresses.	• Will not respond to ARP for any subscriber hosts on the subscriber subnet	• Will not respond to ARP for any remote IP hosts	• Will not route out the group IP interface for subscriber hosts associated with the subscriber subnet. • Subscriber hosts mapped to the group IP interface are remapped to the redundant IP interface. • Will not accept packets destined to the SRRP gateway MAC received on the group IP interface.
Backup (redundant IP interface not available)	• Will respond to ARP for all owned subscriber subnet IP addresses (hardware address and source MAC = group IP interface native MAC). • Will not respond to ARP for subscriber subnet SRRP gateway IP addresses.	• Will not respond to ARP for any subscriber hosts on the subscriber subnet.	• Will not respond to ARP for any remote IP hosts.	• Will route out the group IP interface for subscriber hosts associated with the subscriber subnet using the group IP interface native MAC address. • Subscriber hosts mapped to the redundant IP interface are remapped to the group IP interface. • Will not accept packets destined to the SRRP gateway MAC received on the group IP interface.

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.

Unlike VRRP, only two nodes may participate in an SRRP instance due 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 masters 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 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 698](#)
 - [Containing Service Type and ID on page 698](#)
 - [Containing Subscriber IP Interface Name on page 699](#)
 - [Subscriber Subnet Information on page 699](#)
 - [Containing Group IP Interface Information on page 699](#)
 - [Remote Redundant IP Interface Mismatch on page 699](#)
 - [Remote Sending Redundant IP Interface Unavailable on page 700](#)
 - [Remote SRRP Advertisement SAP Non-Existent on page 700](#)
 - [Remote Sending Local Receive SRRP Advertisement SAP Unavailable on page 700](#)
 - [Local and Remote Dual Master Detected on page 666](#)
-

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 advertisement SAP unavailable from the remote node, an SRRP Remote Receive advertisement SAP Unavailable event will be generated detailing 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.

Managed SAP (M-SAP)

Managed SAPs allow the use of policies and a SAP template for the creation of a SAP. Although the 7750 SR 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.

When enabled, receiving a triggering packet initiates 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 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.

A “capture SAP” triggers the process. This SAP is defined in a similar way to a default SAP but does not forward traffic. 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.

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 M-SAP created will bear both q-tags that arrived in the original packet if authenticated by RADIUS.

An M-SAP is created in the service-id context that is returned from RADIUS. The new RADIUS attribute Alc-MSAP-Serv-Id refers to the service in which the M-SAP is created. In a routed-CO scenario, the M-SAP is created in a group-interface context. The group-interface name is returned from RADIUS attribute Alc-MSAP-Interface and exists in the provided service for the M-SAP to

Managed SAP (M-SAP)

be installed. The M-SAP parameters are defined in the creation policy. The policy name is returned from RADIUS in the attribute Alc-MSAP-Policy in order for the M-SAP to be created.

While M-SAPs are supported in both routed-co and VPLS TPSDA models, the triggering SAP can be configured only in VPLS.

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 RADIUS does not provide all the information required to install the host (lacking an IP address), the MSAP is created with a short timer while waiting for the host acquire the necessary information and install the host. Default SAP policies are used unless the profiles are known.

In most cases, M-SAPs are allowed to have multiple subscribers to share an MSAP. In architectures that provide service access using a shared SAP, multiple subscribers can share the SAP. These environments require few SAPs and therefore are not supported. Multiple leases for the same subscriber are allowed. Only a single M-SAP policy is allowed. If an M-SAP was defined by a host and a new host installation is attempting to change the policy, the installation fails and an event is raised.

All trigger types can be combined on a SAP supporting DHCP and PPPoE hosts.

The authentication policy is defined in the M-SAP policy. Based on the configuration, the system will re-authenticate. If not used, or for PPPoE, the MSAPs remains active if the session renews.

When PPPoE is used with M-SAPs, the authentication-policy cannot use the username for the M-SAP creation.

The authentication policy used in the capture SAP is the same as the policy used for the managed SAP. In a Layer 3 scenario, the authentication policy is defined in the group-interface context. The managed SAP will not be created if the group-interface name returned from RADIUS points to a different authentication policy other than the policy defined by the capture SAP.

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 50 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):




Figure 52: 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
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 meantime) the out-of-credit-action 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.

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 and IES services.

Topics in this section include:

- [Configuring RADIUS Authentication of DHCP Sessions on page 712](#)
- [Configuring Enhanced Subscriber Management on page 713](#)
 - [Basic Configurations on page 713](#)
 - [Configuring Enhanced Subscriber Management Entities on page 714](#)
 - [Configuring a Subscriber Identification Policy on page 715](#)
 - [Configuring a Subscriber Profile on page 716](#)
 - [Configuring an SLA Profile on page 718](#)
 - [Configuring Explicit Mapping Entries on page 719](#)
 - [Applying the Profiles and Policies on page 721](#)
 - [Configuring Dual Homing on page 723](#)
 - [Configuring Wholesale/Retail Parameters on page 725](#)
- [Service Manager Scripting Language on page 1007](#)
- [Sample Python Scripts on page 1015](#)

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 715](#)
- [Configuring a Subscriber Profile on page 716](#)
- [Configuring a Subscriber Identification Policy on page 715](#)
- [Configuring Explicit Mapping Entries on page 719](#)
- [Applying the Profiles and Policies on page 721](#)

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#
```

Routed CO with Basic Subscriber Management Features

The following displays the output of an IES service configured with and without enhanced subscriber management.

```
A:term17>config>service>ies# inf
-----
subscriber-interface "s2" create
  address 11.20.1.1/16
  dhcp
    gi-address 11.20.1.1
  exit
group-interface "g3" create
  description "With Enhanced Subscriber Mgmt"
  arp-populate
  dhcp
    server 12.1.1.1
    trusted
    lease-populate 8000
    no shutdown
  exit
  sap lag-1:11 create
    sub-sla-mgmt
      def-sub-profile "subProf"
      def-sla-profile "slaProf"
      sub-ident-policy "foo"
      multi-sub-sap
      no shutdown
    exit
    host ip 11.20.1.10 mac 00:00:aa:aa:aa:dd subscriber "One" sub-profile
    "subProf" sla-profile "slaProf"
  exit
exit
subscriber-interface "s3" create
  address 11.39.1.1/16
  dhcp
    gi-address 11.39.1.1
  exit
group-interface "g5" create
  description "Without Enhanced Subscriber Mgmt"
  arp-populate
  dhcp
    server 12.1.1.1
    trusted
    lease-populate 8000
    no shutdown
  exit
  sap 4/1/1:24.4094 create
  exit
exit
exit
no shutdown
-----
A:term17>config>service>ies#
```


Applying the Profiles and Policies

NOTE: Subscriber interfaces operate only with basic (or enhanced) subscriber management. At the very least, a host, either statically configured or dynamically learned by DHCP must be present in order for the interface to be useful.

Apply the Enhanced Subscriber Management profiles and policies to the following entities:

- [SLA Profile on page 721](#)
- [Subscriber Identification Policy on page 728](#)
- [Subscriber Profile on page 728](#)

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]}
    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]
    subscriber-interface ip-int-name
      group-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]
        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 [sub-
                scriber sub-ident-string]
```

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>service>vprn 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>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`

Configuring Dual Homing

The following displays an example of a dual homing configuration. The configuration shows dual homing with a peer node with a system address of 1.1.1.23. The DHCP server returns a default route with a 11.21.1.3 next hop.

```
A:ALA-48#
#-----
echo "Redundancy Configuration"
#-----
    redundancy
        multi-chassis
            peer 1.1.1.23 create
            sync
                srrp
                sub-mgmt
                port lag-100 sync-tag "Tag1" create
            exit
            no shutdown
        exit
        no shutdown
    exit
exit
#-----
echo "Service Configuration"
#-----
    service
        customer 1 create
            description "Default customer"
        exit
        sdp 23 create
            far-end 1.1.1.23
            no shutdown
        exit
        ies 40 customer 1 create
            redundant-interface "r40-1" create
                address 2.1.1.1/31
                spoke-sdp 23:1 create
            exit
        exit
        subscriber-interface "s40-1" create
            address 11.21.1.1/16 gw-ip-address 11.21.1.3
            dhcp
                gi-address 11.21.1.1
            exit
        group-interface "g40-1" create
            dhcp
                server 12.1.1.1
                lease-populate 8000
                no shutdown
            exit
        redundant-interface r40-1
        remote-proxy-arp
        sap lag-100:1 create
            sub-sla-mgmt
                def-sub-profile "subProf"
                def-sla-profile "slaProf"
                sub-ident-policy "subIdentPolicy"
            multi-sub-sap
```

Configuring Enhanced Subscriber Management with CLI

```

                                no shutdown
                                exit
                                exit
                                sap lag-100:4094 create
                                exit
                                srrp 1 create
                                message-path lag-100:4094
                                no shutdown
                                exit
                                exit
                                exit
                                no shutdown
                                exit
                                exit
                                ...
                                -----
A:ALA-48#
```

Configuring Wholesale/Retail Parameters

In the following example, the wholesaler is configured in VPRN service 10. The retailers are configured in VPRN services 102 and 105.

```
A:ALA-48#
#-----
...
service
  customer 1 create
    description "Default customer"
  exit
  vprn 101 customer 1 create
    route-distinguisher 1:101
    vrf-target target:1:101
    subscriber-interface "s101-1" create
      address 101.1.1.1/24
      dhcp
        gi-address 101.1.1.1 src-ip-addr
      exit
      group-interface "g101-1" create
        dhcp
          server 12.1.1.1
          trusted
          lease-populate 8000
          gi-address 101.1.1.1
          no shutdown
        exit
        authentication-policy "rad101"
        remote-proxy-arp
        sap lag-1:101 create
          sub-sla-mgmt
            def-sub-profile "defSubProf"
            def-sla-profile "defSlaProf"
            sub-ident-policy "subIdentPol2"
            multi-sub-sap
            no shutdown
          exit
        host ip 102.1.2.100 mac 00:00:aa:aa:aa:aa subscriber "Static1" sub-
profile "subProf" sla-profile "slaProf"
        host ip 102.1.2.101 mac 00:00:aa:aa:aa:aa subscriber "Static2" sub-
profile "subProf" sla-profile "slaProf"
        host ip 102.1.2.102 mac 00:00:aa:aa:aa:aa subscriber "Static3" sub-
profile "subProf" sla-profile "slaProf"
      exit
    exit
  exit
  subscriber-interface "s101-2" create
    address 103.1.1.1/24
    address 103.3.1.1/24
    group-interface "g101-4" create
      arp-populate
      dhcp
        server 12.1.1.1
        trusted
        lease-populate 8000
        gi-address 103.3.1.1
        no shutdown
```

Configuring Enhanced Subscriber Management with CLI

```
        exit
        authentication-policy "rad101"
        sap lag-1:105 create
            sub-sla-mgmt
                def-sub-profile "defSubProf"
                def-sla-profile "defSlaProf"
                sub-ident-policy "subIdentPol2"
                multi-sub-sap
                no shutdown
            exit
        exit
    exit
    exit
    exit
    no shutdown
exit
vprn 102 customer 1 create
    route-distinguisher 1:102
    subscriber-interface "r102-1" fwd-service 101 fwd-subscriber-interface "s101-
1" create
        address 102.1.1.1/24
        address 102.1.2.1/24
        address 101.1.3.1/24
        dhcp
            proxy-server
                emulated-server 102.1.1.1
                lease-time min 5
                no shutdown
            exit
            option
                vendor-specific-option
                    sap-id
                    string "Retailer"
                exit
            exit
            server 102.1.1.2 12.1.1.1
            lease-populate 8000
            no shutdown
        exit
        authentication-policy "rad102"
    exit
    subscriber-interface "r102-2" fwd-service 101 fwd-subscriber-interface "s101-
2" create
        address 103.2.1.1/24
        dhcp
            lease-populate 8000
            no shutdown
        exit
    exit
    exit
    vprn 105 customer 1 create
        route-distinguisher 1:105
        subscriber-interface "s105-1" fwd-service 101 fwd-subscriber-interface "s101-
2" create
            address 103.3.2.1/24
            dhcp
                proxy-server
                    emulated-server 103.3.2.1
                    lease-time min 5
                    no shutdown
                exit
                lease-populate 8000
                no shutdown
            exit
        exit
```

```
        authentication-policy "rad105"  
    exit  
    no shutdown  
exit  
...  
#-----  
A:ALA-48#
```

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]
sub-sla-mgmt
single-sub-parameters
non-sub-traffic sub-profile sub-profile-name
sla-profile sla-profile-name [**subscriber sub-ident-string**]
sub-ident-policy **sub-ident-policy-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]
sub-sla-mgmt
def-sub-profile default-subscriber-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>service>vpls service-id
sap sap-id
host {[ip ip-address] [mac ieee-address]} [subscriber sub-ident-string] [**sub-profile sub-profile-name**] [sla-profile 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
```

Subscriber Management Command Reference

Configuration Commands

- [ANCP Commands on page 732](#)
 - [VPRN GSMP Configuration Commands on page 733](#)
- [Authentication Policy Commands on page 734](#)
- [BGP Peering Policy Commands on page 737](#)
- [Explicit Subscriber Mapping Commands on page 742](#)
- [Host Tracking Policy Commands on page 742](#)
- [Local User Database Commands on page 342](#)
- [Managed SAP Policy Commands on page 756](#)
- [PPPoE Policy Configuration Commands on page 437](#)
- [RADIUS Accounting Policy Commands on page 743](#)
- [SLA Profile Commands on page 747](#)
- [Subscriber Identification Policy Commands on page 749](#)
- [Subscriber Profile Commands on page 750](#)
- [Multi-Chassis Redundancy Commands on page 752](#)
- [VPLS Subscriber Management Configuration Commands on page 754](#)
- [VPRN Subscriber Interface Configuration Commands on page 758](#)
- [IES Subscriber Management Configuration Commands on page 763](#)
- [Redundant Interface Commands on page 768](#)
- [Show Commands on page 768](#)
- [Clear Commands on page 771](#)
- [Debug Commands on page 772](#)

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 61](#) 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-shev [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

```

VPRN GSMP Configuration Commands

```

config
  — service
    — [no] vprn service-id [customer customer-id] [vpn vpn-id] [m-vpls]
      — gsmp
        — [no] group name
          — ancp
            — [no] dynamic-topology-discover
            — [no] oam
          — description description-string
          — no description
          — hold-multiplier multiplier
          — no hold-multiplier
          — keepalive seconds
          — no keepalive
          — [no] neighbor ip-address
            — description description-string
            — no description
            — local-address ip-address
            — no local-address
            — priority-marking dscp dscp-name
            — priority-marking prec ip-prec-value
            — no priority-marking
            — [no] shutdown
          — [no] shutdown
        — [no] shutdown
      — [no] shutdown

```

Authentication Policy Commands

```

config
— subscriber-mgmt
    — [no] authentication-policy name
        — [no] accept-authorization-change
        — description description-string
        — no description
        — [no] include-radius-attribute
            — [no] access-loop-options
            — [no] called-station-id
            — calling-station-id {mac | remote-id | sap-id | sap-string}
            — no calling-station-id
            — [no] circuit-id
            — [no] dhcp-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] 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
            — fallback-action accept
            — fallback-action user-db local-user-db-name
            — no fallback-action
            — retry count
            — no retry
            — router router-instance
            — router service-name
            — no router
            — server server-index [address ip-address] [secret key] [port port]
            — no server server-index
            — source-address ip-address
            — no source-address
            — timeout seconds
            — no timeout
        — [no] re-authentication
        — user-name-format format [append domain-name]
        — no user-name-format

```

```

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
  — 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-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**
- **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-wmark** [**low** *percentage*] [**high** *percentage*]
- **no sub-insert-wmark**

BGP Peering Policy Commands

```

config
  — subscriber-mgmt
    — bgp-peering-policy policy-name [create]
    — no bgp-peering-policy policy-name
      — [no] advertise-inactive
      — [no] aggregator-id-zero
      — [no] as-override
      — auth-keychain name
      — no auth-keychain
      — authentication-key [authentication-key | hash-key] [hash | hash2]
      — no authentication-key
      — cluster cluster-id
      — no cluster
      — [no] connect-retry seconds
      — [no] damping
      — description description-string
      — no description
      — [no] disable-4byte-asn
      — [no] disable-client-reflect
      — disable-communities [standard] [extended]
      — no disable-communities
      — [no] disable-fast-external-failover
      — export policy-name [policy-name...(upto 5 max)]
      — no export
      — hold-time seconds
      — no hold-time
      — import policy-name [policy-name...(up to 5 max)]
      — no import
      — keepalive seconds
      — no keepalive
      — local-address ip-address
      — no local-address
      — local-as as-number [private]
      — no local-as
      — local-preference local-preference
      — no local-preference
      — loop-detect {drop-peer | discard-route | ignore-loop| off}
      — no loop-detect
      — med-out {number | igp-cost}
      — no med-out
      — min-as-origination seconds
      — no min-as-origination
      — min-route-advertisement seconds
      — no min-route-advertisement
      — multihop ttl-value
      — no multihop
      — [no] next-hop-self
      — [no] passive
      — peer-as as-number
      — preference preference
      — no preference
      — prefix-limit limit
      — no prefix-limit

```

- **[no] remove-private**
- **ttl-security** *min-ttl-value*
- **no ttl-security**
- **type** {**internal** | **external**}
- **no type**

Category Map Commands

```

config
— subscriber-mgmt
— category-map category-map-name [create]
— no category-map category-map-name
— 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
— [no] exhausted-credit-service-level
— [no] egress-ip-filter-entries
— entry entry-id [create]
— action drop
— action forward
— action http-redirect url
— no action
— 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 netmask}
— 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] egress-ipv6-filter-entries
— entry entry-id [create]
— action drop
— action forward
— no action
— 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**
- **icmp-code** *icmp-code*
- **no icmp-code**
- **icmp-type** *icmp-type*
- **no icmp-type**
- **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-ipv6-filter-entries**
 - **entry** *entry-id* [*create*]
 - **action drop**
 - **action forward**
 - **action http-redirect** *url*
 - **no action**
 - **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**
- **[no] ingress-ipv6-filter-entries**
 - **entry** *entry-id* [*create*]
 - **action drop**
 - **action forward**

```

— no action
— 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
  — icmp-code icmp-code
  — no icmp-code
  — icmp-type icmp-type
  — no icmp-type
  — 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
— ingress-ipv6-filter-entries pir-rate
— ingress-ipv6-filter-entries max
— no ingress-ipv6-filter-entries
— out-of-credit-action-override {continue | block-category | change-service-level}
— no out-of-credit-action-override
— 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
— description description-string
— no description

```

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
```

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
— [no] host-accounting
— include-radius-attribute
— [no] acct-authentic
— [no] acct-delay-time
— [no] called-station-id
— [no] calling-station-id
— [no] circuit-id
— [no] delegated-ipv6
— [no] framed-interface-id
— [no] framed-ip-addr
— [no] framed-ip-netmask
— [no] framed-ipv6-prefix
— [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
— [no] user-name
— 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]
— no server server-index
— source-address ip-address
— no source-address
— timeout seconds
— no timeout
— session-id-format {description | number}
— no session-id-format
— update-interval minutes
— no update-interval
— [no] use-std-acct-attributes

```

```
config
  — log
    — [no] accounting-policy
      — collection-interval
      — default
      — [no] description
      — record
      — [no] shutdown
      — to
    — event-control application-id [event-name | event-number] [generate [severity-level] [throttle]]
    — event-control application-id [event-name | event-number] suppress
    — no event-control application [event-name | event-number]
```

SLA Profile Commands

```

config
— subscriber-mgmt
— [no] sla-profile sla-profile-name
— description description-string
— no description
— egress
— [no] ip-filter filter-id
— qos sap-egress-policy-id
— no qos
— queue queue-id
— no queue queue-id
— avg-frame-overhead percent
— 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
— 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 {no-stats | minimal | offered-profile-no-cir | offered-priority-no-cir | offered-limited-profile-cir | offered-profile-cir | offered-priority-cir | offered-total-cir}
— no stat-mode
— [no] qos-marking-from-sap
— report-rate agg-rate-limit
— report-rate scheduler scheduler-name
— no report-rate
— host-limit max-number-of-hosts [remove-oldest]
— no host-limit
— 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
— 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**
- **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** | **kilo-bits-per-second**}]
 - **no rate**
 - **stat-mode** {**no-stats** | **minimal** | **offered-profile-no-cir** | **offered-priority-no-cir** | **offered-limited-profile-cir** | **offered-profile-cir** | **offered-priority-cir** | **offered-total-cir**}
 - **no stat-mode**
- **report-rate** **agg-rate-limit**
- **report-rate** **scheduler** *scheduler-name*
- **no report-rate**

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

```

Subscriber Profile Commands

```

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 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 egress-agg-rate-limit
      — egress-scheduler-policy scheduler-policy-name
      — no egress-scheduler-policy
        — scheduler scheduler-name rate [pir-rate] [cir cir-rate]
        — no scheduler scheduler-name
      — [no] host-tracking-policy
      — 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
      — hsmdda
        — egress-queues
          — agg-rate-limit agg-rate
          — no agg-rate-limit
          — qos policy-id
          — no qos
          — [no] queue-override
            — 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**
 - **slope-policy** *hsmda-slope-policy-name*
 - **no slope-policy**
- **ingress-queues**
 - **agg-rate-limit** *agg-rate*
 - **no agg-rate-limit**
 - **qos** *policy-id*
 - **no qos**
 - [**no**] **queue-override**
 - **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**
 - **slope-policy** *hsmda-slope-policy-name*
 - **no slope-policy**
- **ingress-scheduler-policy** *scheduler-policy-name*
- **no ingress-scheduler-policy**
 - **scheduler** *scheduler-name* **rate** [*pir-rate*] [**cir** *cir-rate*]
 - **no scheduler** *scheduler-name*
- **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**]
- **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**

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
— [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] 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

```

VPLS Subscriber Management Configuration Commands

```

config
  — service
    — [no] vpls service-id [customer customer-id] [vpn vpn-id] [m-vpls]
      — gsmp
        — [no] group name
          — ancp
            — [no] dynamic-topology-discover
            — [no] oam
          — description description-string
          — no description
          — hold-multiplier multiplier
          — no hold-multiplier
          — keepalive seconds
          — no keepalive
          — [no] neighbor ip-address
            — description description-string
            — no description
            — local-address ip-address
            — no local-address
            — priority-marking dscp dscp-name
            — priority-marking prec ip-prec-value
            — no priority-marking
            — [no] shutdown
          — [no] shutdown
        — [no] shutdown
      — sap sap-id [split-horizon-group group-name]
      — 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
        — 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
          — host-connectivity-verify source-ip ip-address [source-mac ieee-
            address] [interval interval] [action {remove | alarm}]
        — [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**

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
      — ies-vprn-only-sap-parameters
        — anti-spoof {ip-mac | nh-mac}
        — no anti-spoof
      — 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}
            — vendor-specific-option {mac | string string}
            — no vendor-specific-option
            — [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
config
— service
— vpls
— sap
— default-msap-policy policy-name
— no default-msap-policy
— eval-msap {policy msap-policy-name | msap sap-id }

```

VPRN Subscriber Interface Configuration Commands

```

config
  — service
    — vprn service-id [customer customer-id]
    — no vprn service-id
    — subscriber-interface ip-int-name [fwd-service service-id fwd-subscriber-inter-  

face ip-int-name] [create]
    — no subscriber-interface ip-int-name
      — address {ip-address/mask | ip-address netmask} [gw-ip-address ip-  

address]
      — no address
      — authentication-policy name
      — no authentication-policy
      — delayed-enable seconds [init-only]
      — no delayed-enable
      — description description-string
      — no description
      — dhcp
        — 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-entries]
        — 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] [radius-override]
            — no lease-time
            — [no] shutdown
          — server server1 [server2...(up to 8 max)]
          — no server
          — [no] shutdown
      — [no] group-interface ip-int-name
        — arp-host
          — host-limit max-num-hosts
          — no host-limit
          — min-auth-interval min-auth-interval
          — no min-auth-interval
          — sap-host-limit max-num-hosts-sap
          — no sap-host-limit
          — [no] shutdown
        — [no] arp-populate

```

```

— arp-timeout seconds
— no arp-timeout
— authentication-policy name
— no authentication-policy
— description description-string
— no description
— dhcp
  — client-applications {[dhcp] [pppoe]}
  — no client-applications
  — description description-string
  — no description
  — gi-address ip-address [src-ip-address]
  — no gi-address
  — lease-populate nbr-of-leases
  — no lease-populate
  — [no] match-circuit-id
  — [no] option
    — action {replace | drop | keep}
    — no action
    — circuit-id [ascii-tuple | ifindex | sap-id | vlan-ascii-tupl]
    — 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] [radius-override]
    — no lease-time
    — [no] shutdown
  — server server1 [server2...(up to 8 max)]
  — no server
  — [no] shutdown
  — [no] trusted
— host-connectivity-verify [interval interval] [action {remove|alarm}]
— icmp
  — [no] mask-reply
  — redirects [number seconds]
  — no redirects
  — ttl-expired number seconds]
  — no ttl-expired
  — unreachables [number seconds]
  — no unreachables
— [no] local-proxy-arp
— [no] loopback
— [no] mac ieee-address

```

- [no] **pppoe**
 - **description** *description-string*
 - **no description**
 - **pap-chap-user-db** *local-user-db-name*
 - **no pap-chap-user-db**
 - **pppoe-policy** *pppoe-policy-name*
 - **no pppoe-policy**
 - **sap-session-limit** *sap-session-limit*
 - **no sap-session-limit**
 - **session-limit** *session-limit*
 - **no session-limit**
 - [no] **shutdown**
- [no] **proxy-arp-policy** *policy-name* [*policy-name...*(up to 5 max)]
- [no] **remote-proxy-arp**
- [no] **sap** *sap-id*
 - **accounting-policy** *acct-policy-id*
 - **no accounting-policy** [*acct-policy-id*]
 - **anti-spoof** {**ip** | **ip-mac** | **nh-mac**}
 - **no anti-spoof**
 - **atm**
 - **egress**
 - **traffic-desc** *traffic-desc-profile-id*
 - **no traffic-desc**
 - **encapsulation** *atm-encap-type*
 - **ingress**
 - **traffic-desc** *traffic-desc-profile-id*
 - **oam**
 - [no] **alarm-cells**
 - [no] **periodic-loopback**
 - [no] **calling-station-id**
 - [no] **collect-stats**
 - **cpu-protection** *policy-id* [**mac-monitoring**]
 - **no cpu-protection**
 - **description** *description-string*
 - **no description**
 - **egress**
 - **agg-rate-limit** *agg-rate* [**queue-frame-based-accounting**]
 - **no agg-rate-limit**
 - **filter ip** *ip-filter-id*
 - **filter ipv6** *ipv6-filter-id*
 - **no filter** [**ip** *ip-filter-id*] [**ipv6** *ipv6-filter-id*]
 - **no filter**
 - [no] **qinq-mark-top-only**
 - [no] **qos** *policy-id*
 - **scheduler-policy** *scheduler-policy-name*
 - **no scheduler-policy**
 - **host ip** *ip-address* [**mac** *ieee-address*] [**subscriber** *sub-ident-string*] [**sub-profile** *sub-profile-name*] [**sla-profile** *sla-profile-name*] [**anccp-string** *anccp-string*]
 - **no host** {[**ip** *ip-address*] [**mac** *ieee-address*]}
 - **no host all**
 - **ingress**
 - **filter ip** *ip-filter-id*


```

— filter ipv6 ipv6-filter-id
— no filter [ip ip-filter-id] [ipv6 ipv6-filter-id]
— no filter
— match-qinq-dot1p {top|bottom}
— no match-qinq-dot1p
— qos policy-id [shared-queuing]
— no qos policy-id
— 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
— managed-routes
— route {ip-prefix/length | ip-prefix netmask}
  [create]
— no route {ip-prefix/length | ip-prefix netmask}
— [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] 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-subscriber-profile-name
— no def-sub-profile
— multi-sub-sap subscriber-limit
— 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
— [no] shutdown
— [no] srrp srrp-id
— description description-string
— no description
— gw-mac mac-address
— no gw-mac
— keep-alive-interval interval
— no keep-alive-interval
— message-path sap-id
— no message-path
— [no] policy vrrp-policy-id
— priority priority
— no priority
— [no] shutdown
— tos-marking-state {trusted | untrusted}
— no tos-marking-state
— [no] private-retail-subnets
— [no] shutdown

```

IES Subscriber Management Configuration Commands

```

config
  — service
    — ies service-id [customer customer-id] [vpn vpn-id]
      — [no] interface ip-int-name
        — address ip-address/mask [netmask] [broadcast {all-ones | host-ones}]
        — no address
        — [no] allow-directed-broadcast
        — [no] arp-populate
        — arp-timeout seconds
        — no arp-timeout
        — cflowd [acl | interface]
        — no cflowd
        — description description-string
        — no description
        — dhcp
          — 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
            — action {replace | drop | keep}
            — no action
            — circuit-id [ascii-tuple | ifindex | sap-id]
            — no circuit-id
            — [no] remote-id
          — server server1 [server2...(up to 8 max)]
          — no server
          — [no] shutdown
          — [no] trusted
        — host-connectivity-verify [source {vrrp|interface}] [interval interval]
          [action {remove | alarm}]
        — icmp
          — [no] mask-reply
          — redirects [number seconds]
          — no redirects
          — ttl-expired number seconds]
          — no ttl-expired
          — unreachables [number seconds]
          — no unreachables
        — [no] local-proxy-arp
        — [no] loopback
        — [no] mac ieee-address
        — [no] proxy-arp-policy policy-name [policy-name...(up to 5 max)]
        — [no] remote-proxy-arp
        — [no] sap sap-id
      — [no] subscriber-interface ip-int-name
        — [no] address {ip-address/mask | ip-address netmask}
        — delayed-enable seconds [init-only]
        — no delayed-enable

```

- **description** *description-string*
- **no description**
- **dhcp**
 - **gi-address** *ip-address* [*src-ip-addr*]
 - **no gi-address**
- **[no] group-interface** *ip-int-name*
 - **arp-host**
 - **host-limit** *max-num-hosts*
 - **no host-limit**
 - **min-auth-interval** *min-auth-interval*
 - **no min-auth-interval**
 - **sap-host-limit** *max-num-hosts-sap*
 - **no sap-host-limit**
 - **[no] shutdown**
 - **[no] arp-populate**
 - **arp-timeout** *seconds*
 - **no arp-timeout**
 - **authentication-policy** *name*
 - **no authentication-policy**
 - **description** *description-string*
 - **no description**
 - **dhcp**
 - **client-applications** **dhcp**
 - **srrp-sync-data** **pppoe**
 - **srrp-sync-data** **dhcp** **pppoe**
 - **no srrp-sync-data**
 - **description** *description-string*
 - **no description**
 - **gi-address** *ip-address* [*src-ip-addr*]
 - **no gi-address**
 - **lease-populate** *nbr-of-leases*
 - **no lease-populate**
 - **[no] match-circuit-id**
 - **option**
 - **action** {**replace** | **drop** | **keep**}
 - **no action**
 - **circuit-id** [*ascii-tuple* | *ifindex* | *sap-id*]
 - **no circuit-id**
 - **[no] remote-id**
 - **[no] vendor-specific-option**
 - **[no] client-mac-address**
 - **[no] sap-id**
 - **[no] service-id**
 - **string** *text*
 - **no string**
 - **[no] system-id**
 - **server** *server1* [*server2...*(up to 8 max)]
 - **no server**
 - **[no] shutdown**
 - **[no] trusted**
 - **host-connectivity-verify** [*interval interval*] [**action** {**remove** | **alarm**}]
 - **icmp**
 - **[no] mask-reply**
 - **redirects** [*number seconds*]

- **no redirects**
- **ttl-expired** *number seconds*]
- **no ttl-expired**
- **unreachables** [*number seconds*]
- **no unreachables**
- **[no] local-proxy-arp**
- **[no] loopback**
- **[no] mac** *ieee-address*
- **[no] oper-up-while-empty**
- **[no] pppoe**
 - **description** *description-string*
 - **no description**
 - **pap-chap-user-db** *local-user-db-name*
 - **no pap-chap-user-db**
 - **pppoe-policy** *pppoe-policy-name*
 - **no pppoe-policy**
 - **sap-session-limit** *sap-session-limit*
 - **no sap-session-limit**
 - **session-limit** *session-limit*
 - **no session-limit**
 - **[no] shutdown**
- **[no] proxy-arp-policy** *policy-name* [*policy-name...*(up to 5 max)]
- **[no] remote-proxy-arp**
- **[no] sap** *sap-id*
 - **accounting-policy** *acct-policy-id*
 - **no accounting-policy** [*acct-policy-id*]
 - **anti-spoof** {**ip** | **ip-mac** | **nh-mac**}
 - **no anti-spoof**
 - **atm**
 - **egress**
 - **traffic-desc** *traffic-desc-profile-id*
 - **no traffic-desc**
 - **encapsulation** *atm-encap-type*
 - **ingress**
 - **traffic-desc** *traffic-desc-profile-id*
 - **oam**
 - **[no] alarm-cells**
 - **[no] periodic-loopback**
 - **[no] calling-station-id**
 - **[no] collect-stats**
 - **description** *description-string*
 - **no description**
 - **egress**
 - **agg-rate-limit** *agg-rate*
 - **no agg-rate-limit** *agg-rate*
 - **filter ip** *ip-filter-id*
 - **filter**
 - **no filter ip** *ip-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 | default}*
- **no mbs**
- **rate** *pir-rate [cir cir-rate]*
- **no rate**
- **scheduler-policy** *scheduler-policy-name*
- **no scheduler-policy**
- **host** *ip ip-address [mac ieee-address] [subscriber sub-ident-string] [sub-profile sub-profile-name] [sla-profile sla-profile-name]*
- **no host** *{[ip ip-address] [mac ieee-address]}*
- **no host** *all*
- **ingress**
 - **filter** *ip ip-filter-id*
 - **no filter**
 - **match-qinq-dot1p** *{top|bottom}*
 - **no match-qinq-dot1p**
 - **qos** *policy-id [shared-queuing]*
 - **no qos**
 - **scheduler-policy** *scheduler-policy-name*
 - **no scheduler-policy**
- **[no] multi-service-site** *customer-site-name*
- **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] sub-sla-mgmt**
 - **[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-subscriber-profile-name
— no def-sub-profile
— multi-sub-sap subscriber-limit
— 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
— [no] shutdown
— tos-marking-state {trusted | untrusted}
— no tos-marking-state
— [no] shutdown
— [no] srrp srrp-id
  — bfd-enable
  — description description-string
  — no description
  — gw-mac mac-address
  — no gw-mac
  — keep-alive-interval interval
  — no keep-alive-interval
  — message-path sap-id
  — no message-path
  — [no] policy vrrp-policy-id
  — priority priority
  — no priority
  — [no] shutdown

```

Redundant Interface Commands

```

config
  — service
    — ies
      — [no] redundant-interface ip-int-name
        — address {ip-address/mask | ip-address netmask} [remote-ip ip-address]
        — no address
        — [no] description description-string
        — [no] shutdown
        — [no] spoke-sdp sdp-id:vc-id
          — egress
            — filter [ip ip-filter-id]
            — vc-label ingress-vc-label
            — no vc-label [ingress-vc-label]
          — ingress
            — filter [ip ip-filter-id]
            — no filter
            — vc-label ingress-vc-label
            — no vc-label [ingress-vc-label]
        — [no] shutdown

config
  — service
    — vprn
      — [no] redundant-interface ip-int-name
        — address {ip-address/mask | ip-address netmask} [remote-ip ip-address]
        — no address
        — [no] description description-string
        — [no] shutdown
        — [no] spoke-sdp sdp-id:vc-id
          — egress
            — filter [ip ip-filter-id]
            — vc-label ingress-vc-label
            — no vc-label [ingress-vc-label]
          — ingress
            — filter [ip ip-filter-id]
            — no filter
            — vc-label ingress-vc-label
            — no vc-label [ingress-vc-label]
        — [no] shutdown

```

Show Commands

```

show
  — service
    — active-subscribers summary
    — active-subscribers [subscriber sub-ident-string] [sap-id sap-id] [sla-profile sla-profile-name] [detail | mirror]
    — active-subscribers hierarchy [subscriber sub-ident-string]
    — host-tracking [subscriber sub-ident-string] [detail | statistics | summary]
    — host-tracking groups [group group-ip-address] [detail | summary]
    — 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
— retailers
— wholesalers
— subscriber-hosts [sap sap-id] [ip ip-address[/mask]] [mac ieee-address] [sub-pro-
file sub-profile-name] [sla-profile sla-profile-name] [detail]
— gsmp
  — 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 (VPRN only)
— 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
— redundancy
  — multi-chassis all
  — multi-chassis mc-lag
  — multi-chassis sync
    — 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|statis-
tics]]
    — mc-ring global-statistics

show
— subscriber-mgmt
  — authentication [policy-name] [detail]
  — authentication coa-statistics
  — authentication [policy-name] statistics
  — explicit-subscriber-map
  — local-user-db local-user-db-name association [dhcp] [pppoe]
  — local-user-db local-user-db-name dhcp-all-hosts
  — local-user-db local-user-db-name dhcp-host dhcp-host-name

```

```

— local-user-db local-user-db-name dhcp-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
— msap-policy [msap-policy-name] [association]
— sla-profile [sla-profile-name] [association]
— 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]

show
— qos
— 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]
—   sap sap-id [scheduler scheduler-name] [ingress|egress] [detail]
—   subscriber sub-ident-string [scheduler scheduler-name] [ingress|egress][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]
—   sap sap-id [scheduler scheduler-name] [ingress|egress] [detail]
—   subscriber sub-ident-string [scheduler scheduler-name] [ingress|egress][detail]
— shared-queue shared-queue-policy-name [detail]

```

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
    — authentication [policy-name]
    — authentication coa-statistics
    — msap-policy msap-policy-name
    — 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
      — msap msap-id
      — msap-policy msap-policy-name
    — statistics
      — subscriber sub-ident-string
    — qos
      — scheduler-stats
        — subscriber sub-ident-string [scheduler scheduler-name] [ingress|egress]

clear
  — router
    — srrp
      — interface subscriber-interface [id srrp-id]
      — statistics interface subscriber-interface [id srrp-id]

```

Debug Commands

```
debug
— service
    — id service-id
        — arp
        — [no] arp-host
— 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] srrp
        — [no] events [interface ip-int-name]
        — [no] packets [interface ip-int-name]
```

Tools Commands

```
tools
— dump
    — redundancy
        — multi-chassis
            — mc-ring
            — srrp-sync-data [instance instance-id] [peer ip-address]
            — sync-database [peer ip-address] [port port-id | lag-id] [sync-tag sync-tag] [application {dhcps|igmp|igmp-snooping|mc-ring|srrp|subscriber-mgmt|mld-snooping}] [detail] [type {alarm-deleted|local-deleted}]
— perform
    — 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}
        — 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

Topics in this section:

- [Generic Commands on page 774](#)
- [ANCP and GSMP Commands on page 777](#)
- [RADIUS Policy Commands on page 785](#)
- [Multi-Chassis Redundancy Commands on page 833](#)
- [Subscriber Identification Policy Commands on page 848](#)
- [Subscriber Profile Commands on page 852](#)
- [SLA Profile Commands on page 871](#)
- [Subscriber Management Service Commands on page 895](#)
- [Managed SAP Policy Commands on page 905](#)
- [Subscriber Management Service Commands on page 928](#)
- [Category Map Commands on page 949](#)
- [Clear Commands on page 991](#)
- [Show Commands on page 961](#)
- [Tools Commands on page 994](#)
- [Debug Commands on page 998](#)
- [Monitor Commands on page 1002](#)

Generic Commands

description

Syntax	description <i>description-string</i> no description
Context	config>subscr-mgmt>authentication-policy config>subscr-mgmt>host-tracking 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>service>vppls>gsmp>group config>log>accounting-policy config>service>vpnr>redundant-interface 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>vpnr>subscriber-interface config>service>vpnr>sub-if>dhcp config>service>vpnr>subscriber-interface>group-interface config>service>vpnr>subscriber-interface>grp-if>sap config>service>vpnr>sub-if>grp-if>srrp config>service>vpnr>subscriber-interface>grp-if>dhcp config>service>vpnr>gsmp>group config>service>vpnr>gsmp>group>neighbor config>redundancy>multi-chassis>peer config>subscr-mgmt>cat-map
Description	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 context in the configuration file. The no form of this command removes any description string from the context.
Default	No description is associated with the configuration context.
Parameters	<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

Syntax	[no] shutdown
Context	<pre> config>subscr-mgmt>sub-ident-policy>primary config>subscr-mgmt>sub-ident-policy>secondary config>subscr-mgmt>sub-ident-policy>tertiary config>service>vpls>sap>sub-sla-mgmt config>service>vpls>gsmp config>service>vpls>gsmp>group config>service>vpls>gsmp>group>neighbor config>service>vprn>redundant-interface config>service>vprn>redundant-interface>spoke-sdp config>service>vprn>subscriber-interface config>service>vprn>subscriber-interface>group-interface config>service>vprn>subscriber-interface>grp-if>dhcp config>service>vprn>sub-if>grp-if>srrp config>service>ies>subscriber-interface config>service>ies>subscriber-interface>grp-if>dhcp config>service>ies>sub-if>grp-if>srrp config>service>ies>redundant-interface config>service>ies>sub-if>grp-if>arp-host config>service>vprn>gsmp>group>neighbor config>redundancy>multi-chassis>peer config>redundancy>multi-chassis>peer>mc-lag config>redundancy>multi-chassis>peer>sync </pre>
Description	The shutdown 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 no shutdown command. The shutdown 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. Unlike other commands and parameters where the default state is not indicated in the configuration file, shutdown and no shutdown are always indicated in system generated configuration files. Shutting down a subscriber interface will operationally shut down all child group interfaces and SAPs. Shutting down a group interface will operationally shut down all SAPs that are part of that group-interface. The no form of the command puts an entity into the administratively enabled state.
Default	no shutdown

subscriber-mgmt

Syntax	subscriber-mgmt
Context	config

Description 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.

All subscriber hosts belonging to the same subscriber are subject to the same hierarchical QoS (HQoS) processing. The HQoS processing is defined in the **sub-profile** (the subscriber profile). A sub-profile refers to an existing scheduler policy (configured in **the configure>qos>scheduler-policy** context) and offers the possibility to overrule the rate of individual schedulers within this policy.

Because all subscriber hosts use the same scheduler policy instance, they must all reside on the same complex.

ANCP and GSMP Commands

ancp

Syntax	ancp
Context	config>subscr-mgmt config>subscr-mgmt>sub-prof
Description	This command enables the context to configure Access Node Control Protocol (ANCP) parameters.

ancp-policy

Syntax	ancp-policy <i>name</i>
Context	config>subscr-mgmt>ancp
Description	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. ANCP polices control rates and subscribers based on port-up/port-down messages from the access node. When configured, the 7750 SR 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. 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.
Default	No policies are defined.
Parameters	<i>name</i> — Configures the ANCP policy name.

ancp-policy

Syntax	ancp-policy <i>name</i>
Context	config>subscr-mgmt>sub-prof>ancp
Description	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.
Default	No policies are defined.
Parameters	<i>name</i> — Specifies an existing ANCP policy name.

ingress

Syntax	ingress
Context	config>subscr-mgmt>ancp>ancp-policy
Description	This command configures ingress ANCP policy parameters.

rate-adjustment

Syntax	rate-adjustment <i>adjusted-percent</i> no rate-adjustment				
Context	config>subscr-mgmt>ancp>ancp-policy>ingress config>subscr-mgmt>ancp>ancp-policy>egress				
Description	This command configures a rate adjustment for the scheduler. The rate-adjustment command should be used when the rate returned by the DSLAM is calculated with different encapsulation than the 7750 SR. The node will adjust the rate by the percent specified as: $\text{DSLAM\_RATE} * \text{adjust-rate} / 100 \text{ — rate-reduction.}$ The no form of the command returns the default value.				
Default	none				
Parameters	<i>adjusted-percent</i> — Specifies a rate adjustment for the scheduler. <table> <tr> <td>Values</td><td>1 — 200</td></tr> <tr> <td>Default</td><td>100</td></tr> </table>	Values	1 — 200	Default	100
Values	1 — 200				
Default	100				

rate-reduction

Syntax	rate-reduction <i>kilobit-per-second</i> no rate-reduction
Context	config>subscr-mgmt>ancp>ancp-policy>ingress config>subscr-mgmt>ancp>ancp-policy>egress
Description	This command defines a constant rate reduction to the rate specified by the DSLAM. The rate-reduction 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: $\text{DSLAM\_RATE} * \text{adjust-rate} / 100 \text{ — rate-reduction}$
Default	none

rate-monitor

Syntax	rate-monitor <i>kilobit-per-second</i> [alarm] no rate-monitor
Context	config>subscr-mgmt>ancp>ancp-policy>ingress config>subscr-mgmt>ancp>ancp-policy>egress
Description	This command configures the rate monitor level.
Default	none
Parameters	<i>kilobit-per-second</i> — Specifies the rate, in kilobits, below which the system will generate an event. alarm — 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

Syntax	rate-modify { scheduler <i>scheduler-name</i> arbiter <i>arbiter-name</i> } no rate-modify
Context	config>subscr-mgmt>ancp>ancp-policy>ingress
Description	This command configures ingress rate modify scheduler parameters.
Default	none
Parameters	scheduler <i>scheduler-name</i> — Specifies a scheduler name. arbiter <i>arbiter-name</i> — Specifies an arbiter name

egress

Syntax	egress
Context	config>subscr-mgmt>ancp>ancp-policy
Description	This command configures egress ANCP policy parameters.

rate-modify

Syntax	rate-modify { scheduler <i>scheduler-name</i> arbiter <i>arbiter-name</i> } rate-modify agg-rate-limit no rate-modify
Context	config>subscr-mgmt>ancp>ancp-policy>egress
Description	This command configures egress rate modify scheduler parameters.

Triple Play Subscriber Management Configuration Commands

Default	none
Parameters	agg-rate-limit — specifies that the maximum total rate for all subscriber egress queues for each subscriber associated with the policy. scheduler <i>scheduler-name</i> — Specify a scheduler name. arbiter <i>arbiter-name</i> — Specifies an arbiter name

port-down

Syntax	[no] port-down
Context	config>subscr-mgmt>ancp>ancp-policy
Description	This command specifies the number of GSMP portdown messages received in this ANCP session.

disable-shcv

Syntax	[no] disable-shcv [alarm] [hold-time seconds]
Context	config>subscr-mgmt>ancp>ancp-policy>port-down
Description	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. If the optional alarm flag is used the node should send a SHCV alarm before suspending.
Default	no disable-shcv

ancp-static-map

Syntax	ancp-static-map
Context	config>subscr-mgmt>ancp
Description	This command enables the context to configure a static ANCP name map.
Default	ancp-static-map

entry

Syntax	entry key <i>ancp-string</i> customer <i>customer-id</i> multi-service-site <i>customer-site-name</i> ancp-policy <i>policy-name</i> entry key <i>ancp-string</i> sap <i>sap-id</i> ancp-policy <i>policy-name</i> no entry key <i>ancp-string</i> customer <i>customer-id</i> multi-service-site <i>customer-site-name</i> no entry key <i>ancp-string</i> sap <i>sap-id</i>
Context	config>subscr-mgmt>ancp>static-map
Description	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. 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.
Default	none
Parameters	key <i>ancp-string</i> — Specify the ASCII representation of the DSLAM circuit-id name. customer <i>customer-id</i> — Specify the associated existing customer name. multi-service-site <i>customer-site-name</i> — Specify the associated customer's configured MSS name. ancp-policy <i>policy-name</i> — Specify an existing ANCP policy name. sap <i>sap-id</i> — Specifies the physical port identifier portion of the SAP definition. See Common Service Commands on page 1020 for <i>sap-id</i> command syntax.

gsmp

Syntax	gsmp
Context	config>service>vpls config>service>vprn
Description	This command enables the context to configure GSMP connections maintained in this service.
Default	not enabled

group

Syntax	[no] group <i>name</i>
Context	config>service>vpls>gsmp config>service>vprn>gsmp
Description	This command specifies a GSMP name. A GSMP group name is unique only within the scope of the service in which it is defined.

ancp

Syntax	ancp
Context	config>service>vpls>gsmp>group config>service>vprn>gsmp>group
Description	This command configures ANCP parameters for this GSMP group.

dynamic-topology-discover

Syntax	[no] dynamic-topology-discover
Context	config>service>vpls>gsmp>group>ancp config>service>vprn>gsmp>group>ancp
Description	This command enables the ANCP dynamic topology discovery capability. The no form of this command disables the feature.

oam

Syntax	[no] oam
Context	config>service>vpls>gsmp>group>ancp config>service>vprn>gsmp>group>ancp
Description	This command specifies whether or not the GSMP ANCP OAM capability should be negotiated at startup of the GSMP connection. The no form of this command disables the feature.

hold-multiplier

Syntax	hold-multiplier <i>multiplier</i> no hold-multiplier
Context	config>service>vpls>gsmp>group config>service>vprn>gsmp>group
Description	This command configures the hold-multiplier for the GSMP connections in this group.
Parameters	<i>multiplier</i> — Specifies the GSMP hold multiplier value. Values 1 — 100

keepalive

Syntax	keepalive <i>seconds</i> no keepalive
Context	config>service>vpls>gsmp>group config>service>vprn>gsmp>group
Description	This command configures keepalive values for the GSMP connections in this group.
Parameters	<i>seconds</i> — Specifies the GSMP keepalive timer value in seconds. Values 1 — 25

neighbor

Syntax	[no] neighbor <i>ip-address</i>
Context	config>service>vpls>gsmp>group config>service>vprn>gsmp>group
Description	This command configures a GSMP ANCP neighbor.
Parameters	<i>ip-address</i> — Specifies the IP address of the GSMP ANCP neighbor.

local-address

Syntax	local-address <i>ip-address</i> no local-address
Context	config>service>vpls>gsmp>group>neighbor config>service>vprn>gsmp>group>neighbor
Description	This command configures the source ip-address used in the connection towards the neighbor. The local address is optional. If specified the node will accept connections only for that address in the service running ANCP. The address may be created after the reference but connections will not be accepted until it is created. If the local address is not used, the system accepts connections on any interface within the routing context.
Parameters	<i>ip-address</i> — Specifies the source IP address to be used in the connection toward the neighbor.

priority-marking

Syntax	priority-marking dscp <i>dscp-name</i> priority-marking prec <i>ip-prec-value</i> no priority-marking
Context	config>service>vpls>gsmp>group>neighbor config>service>vprn>gsmp>group>neighbor
Description	This command configures the type of priority marking to be used.
Parameters	dscp <i>dscp-name</i> — Specifies the DSCP code-point to be used. Values be, cp1, cp2, cp3, cp4, cp5, cp6, cp7, cs1, cp9, af11, cp11, af12, cp13, af13, cp15, cs2, cp17, af21, cp19, af22, cp21, af23, cp23, cs3, cp25, af31, cp27, af32, cp29, af33, cp31, cs4, cp33, af41, cp35, af42, cp37, af43, cp39, cs5, cp41, cp42, cp43, cp44, cp45, ef, cp47, nc1, cp49, cp50, cp51, cp52, cp53, cp54, cp55, nc2, cp57, cp58, cp59, cp60, cp61, cp62, cp63 prec <i>ip-prec-value</i> — Specifies the precedence value to be used. Values 0 — 7

RADIUS Policy Commands

authentication-policy

Syntax	authentication-policy <i>name</i> no authentication-policy
Context	config>subscr-mgmt
Description	This command creates the context to configure RADIUS server parameters for session authentication. The policies can be applied to an IES or VPRN interface, or a VPLS SAP. The no form of the command removes the RADIUS server configuration for session authentication. RADIUS servers can be configured for three different applications: 1. For authentication of dynamic Triple Play subscriber sessions, under config>subscr-mgmt>authentication-plcy 2. For 802.1x port authentication, under config>system>security>dot1x>radius-plcy 3. For CLI login users, under config>system>radius
Default	none
Parameters	<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.

radius-accounting-policy

Syntax	radius-accounting-policy <i>name</i> no radius-accounting-policy
Context	config>subscr-mgmt config>subscr-mgmt>sub-prof
Description	This command specifies a subscriber RADIUS based accounting policy.
Parameters	<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

Syntax	[no] accept-authorization-change
Context	config>subscr-mgmt

Triple Play Subscriber Management Configuration Commands

Description	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.
Default	no accept-authorization-change

host-accounting

Syntax	[no] host-accounting
Context	config>subscr-mgmt>acct-plcy
Description	This command enables per-host accounting. The no form of the command disables per-host accounting.

include-radius-attribute

Syntax	[no] include-radius-attribute
Context	config>subscr-mgmt>auth-plcy config>subscr-mgmt>acct-plcy
Description	This command enables the context to specify the RADIUS parameters that the system should include into RADIUS authentication-request messages.

use-std-acct-attributes

Syntax	[no] use-std-acct-attributes
Context	config>subscr-mgmt>radius-acct-policy>include-radius-attribute
Description	This command sets use of RADIUS standard define attributes for sending accounting information.
Default	Per default ALU VSAs are sent.

acct-authentic

Syntax	[no] acct-authentic
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 acct-authentic RADIUS attribute.

acct-delay-time

Syntax	[no] acct-delay-time
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 acct-delay-time RADIUS attribute.

called-station-id

Syntax	[no] called-station-id
Context	config>subscr-mgmt>auth-policy>include-radius-attribute config>subscr-mgmt>acct-plcy>include-radius-attribute
Description	This command includes called station id attributes. The no form of the command excludes called station id attributes.

calling-station-id

Syntax	calling-station-id {mac remote-id sap-id sap-string} no calling-station-id
Context	config>service>ies>if>sap config>service>ies>sub-if>grp-if>sap config>service>vpls>sap config>service>vprn>if>sap config>service>vprn>sub-if>grp-if>sap config>subscr-mgmt>auth-plcy>include-radius-attribute config>subscr-mgmt>acct-plcy>include>include-radius-attribute
Description	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 calling-station-id value is set at the SAP level, the calling-station-id attribute will not be sent.
Default	no calling-station-id
Parameters	mac — Specifies that the mac-address will be sent. remote-id — Specifies that the remote-id will be sent. sap-id — Specifies that the sap-id will be sent. sap-string — 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.

access-loop-options

Syntax	[no] access-loop-options
Context	config>subscr-mgmt>auth-plcy>include-radius-attribute
Description	This command enables the generation of the DSL Forum access loop characteristics RADIUS attributes.

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.

delegated-ipv6

Syntax	[no] delegated-ipv6
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 delegated-ipv6 RADIUS attribute.

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

dhcp-vendor-class-id

Syntax	[no] dhcp-vendor-class-id
Context	config>subscr-mgmt>auth-plcy>include-radius-attribute
Description	This command enables the generation of the dhcp-vendor-class-id RADIUS attribute.

framed-interface-id

Syntax	[no] framed-interface-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 framed-interface-id RADIUS attribute.

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

Syntax	[no] framed-ip-netmask
Context	config>subscr-mgmt>acct-plcy>include-radius-attribute
Description	This command enables the inclusion of the framed-ip-netmask attribute.

framed-ipv6-prefix

Syntax	[no] framed-ipv6-prefix
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 framed-ipv6-prefix RADIUS attribute.

ipv6-address

Syntax	[no] framed-ipv6-address
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 ipv6-address RADIUS attribute.

mac-address

Syntax	[no] mac-address config>subscr-mgmt>auth-policy>include-radius-attribute config>subscr-mgmt>acct-plcy>include-radius-attribute
Description	This command enables the generation of the client MAC address RADIUS attribute.

nas-identifier

Syntax	[no] nas-identifier
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-identifier RADIUS attribute.

nas-port

Syntax	[no] nas-port <i>bit-specification binary-spec</i>
Context	config>subscr-mgmt>acct-plcy>include-radius-attribute
Description	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 no form of this command disables your nas-port configuration.
Parameters	<i>binary spec</i> — <i>bit-specification binary-spec</i> <i>bit-specification</i> — 0 1 <i>bit-origin</i> <i>bit-origin</i> — * <i>number-of-bits origin</i> <i>number-of-bits</i> — 1—32 <i>origin</i> — o i s m p Values o: 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 1111 1111 1111 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 0/0/0/0/0/0.
Parameters	prefix-string <i>string</i> — Specifies that a user configurable string will be added to the RADIUS NAS port attribute, up to 8 characters in length. suffix <i>suffix-option</i> — Specifies the suffix type to be added to the RADIUS NAS port attribute.
Values	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 nas-port-type , the following will be sent: values: 32 (null-encap), 33 (dot1q), 34 (qinq), 15 (DHCP hosts). The nas-port-type can also be set as a specified value, with an integer from 0 to 255. The no form of the command reverts to the default.
Default	no nas-port-type
Parameters	0 — 255 — Specifies an enumerated integer that specifies the value that will be put in the RADIUS nas-port-type attribute.

nat-port-range

Syntax	[no] nat-port-range
Context	config>subscr-mgmt>acct-plcy>include-radius-attribute
Description	This command enables the generation of the nat-port-range attribute.
Default	no nat-port-range

pppoe-service-name

Syntax	[no] pppoe-service-name
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 pppoe-service-name RADIUS attribute.

remote-id

Syntax	[no] remote-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-remote-id for RADIUS.

tunnel-server-attrs

Syntax	[no] tunnel-server-attrs
Context	config>subscr-mgmt>auth-policy>include-radius-attribute
Description	This command includes tunner-server attribute.

sla-profile

Syntax	[no] sla-profile
Context	config>subscr-mgmt>acct-plcy>include-radius-attribute
Description	This command specifies that SLA profile attributes should be included into RADIUS accounting messages.

sub-profile

Syntax	[no] sub-profile
Context	config>subscr-mgmt>acct-plcy>include-radius-attribute
Description	This command specifies that subscriber profile attributes should be included into RADIUS accounting messages.

subscriber-id

Syntax	[no] subscriber-id
Context	config>subscr-mgmt>acct-plcy>include-radius-attribute
Description	This command specifies that subscriber ID attributes should be included into RADIUS accounting messages.

tunnel-server

Syntax	[no] tunnel-server
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 tunnel-server RADIUS attribute.

user-name

Syntax	[no] user-name
Context	config>subscr-mgmt>acct-plcy>include-radius-attribute
Description	This command enables the inclusion of the user-name attribute. The no form of the command disables the inclusion of the user-name attribute.
Default	no user-name

password

Syntax	password <i>password</i> [hash hash2] no password
Context	config>subscr-mgmt>auth-policy
Description	This command sets a password that is sent with user-name 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.
Default	none
Parameters	<i>password</i> — A text string containing the password. Allowed values are any string up to 10 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.

hash — 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. For security, all keys are stored in encrypted form in the configuration file with the **hash** parameter specified.

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.

ppp-user-name

Syntax	ppp-user-name append <i>domain-name</i> ppp-user-name default-domain <i>domain-name</i> ppp-user-name replace <i>domain-name</i> ppp-user-name strip no ppp-user-name
Context	config>subscr-mgmt>auth-plcy
Description	The command configures the PPP user name operation. Use this command to specify the operation to perform on the PAP/CHAP user name. The no form of the command reverts to the default.
Default	no ppp-user-name
Parameters	append <i>domain-name</i> — A string specified by tmnxSubAuthPlcyPppDomain, preceded with a '@', is appended to the PAP/CHAP user name. default-domain <i>domain-name</i> — The same action is performed as with appendDomain, but only if the PAP/CHAP user name does not already contain a domain name. replace <i>domain-name</i> — All characters after a '@' delimiter are replaced with the string specified by tmnxSubAuthPlcyPppDomain. strip — Any '@' character and all subsequent characters are removed from the PAP/CHAP user name.

pppoe-access-method

Syntax	pppoe-access-method {none padi pap-chap} no pppoe-access-method
Context	config>subscr-mgmt>auth-plcy
Description	This command indicates the authentication method used towards the RADIUS server in case the policy is used for PPPoE.
Parameters	none — Indicates that the client will be authenticated by the local user database defined under the group interface and not through RADIUS. padi — 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).

pap-chap — 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.

radius-authentication-server

Syntax	radius-accounting-server
Context	config>subscr-mgmt>acct-plcy
Description	This command creates the context for defining RADIUS accounting server attributes under a given session authentication policy.

access-algorithm

Syntax	access-algorithm {direct round-robin} no access-algorithm
Context	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 <i>local-user-db-name</i> no fallback-action
Context	config>subscr-mgmt>auth-plcy-srvr
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

router

Syntax	router <i>router-instance</i> router <i>service-name</i> 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	<i>router-instance</i> — Specifies the virtual router instance.						
Values	<table> <tr> <td>router-name:</td><td>Base, management</td></tr> <tr> <td>service-id:</td><td>1 — 2147483647</td></tr> <tr> <td>service-name:</td><td>Specifies the service name up to 64 characters in length.</td></tr> </table>	router-name:	Base, management	service-id:	1 — 2147483647	service-name:	Specifies the service name up to 64 characters in length.
router-name:	Base, management						
service-id:	1 — 2147483647						
service-name:	Specifies the service name up to 64 characters in length.						

retry

Syntax	retry <i>count</i> 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	<i>count</i> — The retry count.
Values	1 — 10

server

Syntax	server <i>server-index</i> address <i>ip-address</i> secret <i>key</i> [hash hash2] [port <i>port-num</i>] no server <i>index</i>
Context	config> subscr-mgmt>auth-policy>auth-plcy-srvr config>subscr-mgmt>acct-plcy-server
Description	This command adds a RADIUS server and configures the RADIUS server IP address, index, and key values. Up to five RADIUS servers can be configured at any one time. RADIUS servers are accessed in order from lowest to highest index for authentication 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.

The **no** form of the command removes the server from the configuration.

Default	No RADIUS servers are configured.
Parameters	<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. Values 1 — 5 address <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. secret <i>key</i> — The secret key to access the RADIUS server. This secret key must match the password on the RADIUS server. Values Up to 20 characters in length. hash — 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. For security, all keys are stored in encrypted form in the configuration file with the hash parameter specified. 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. <i>port</i> — Specifies the UDP port number on which to contact the RADIUS server for authentication. Values 1 — 65535

source-address

Syntax	source-address <i>ip-address</i> no source-address
Context	config>subscr-mgmt>auth-plcy-srvr config>subscr-mgmt>acct-plcy>server
Description	This command configures the source address of the RADIUS packet. The system IP address must be configured in order for the RADIUS client to work. See Configuring a System Interface in the 7750 SR OS Router Configuration Guide. Note that the system IP address must only be configured if the source-address is not specified. When the no source-address 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 no source-address was given. The no form of the command reverts to the default value.
Default	System IP address
Parameters	<i>ip-address</i> — The IP prefix for the IP match criterion in dotted decimal notation. Values 0.0.0.0 - 255.255.255.255

timeout

Syntax	timeout <i>seconds</i> no timeout
Context	config>subscr-mgmt>auth-plcy-srvr config>subscr-mgmt>acct-plcy>server This command configures the number of seconds the router waits for a response from a RADIUS server. The no form of the command reverts to the default value.
Default	3 seconds
Parameters	<i>seconds</i> — The number of seconds the router waits for a response from a RADIUS server, expressed as a decimal integer. Values 1 — 90

session-id-format

Syntax	session-id-format { description number } no session-id-format
Context	config>subscr-mgmt>acct-plcy
Description	This command specifies the format for the acct-session-id attribute used in RADIUS accounting requests.
Parameters	description — Specifies to use a string containing following information <subscriber>@<sap-id>@<SLA-profile>_<creation-time>. number — Specifies to use a unique number generated by the OS to identify a given session.

update-interval

Syntax	update-interval <i>minutes</i> no update-interval
Context	config>subscr-mgmt>acct-plcy
Description	This command specifies the interval at which accounting data of subscriber hosts will be updated.
Parameters	<i>minutes</i> — Specifies the interval, in minutes, at which accounting data of subscriber hosts will be updated. Values 5 — 1080

use-std-acct-attributes

Syntax	[no] use-std-acct-attributes
Context	config>subscr-mgmt>radius-acct-policy
Description	This command enables the use of standard accounting attributes. The no form of the command disables the use of standard accounting attributes.

re-authentication

Syntax	[no] re-authentication
Context	config>subscr-mgmt>auth-policy
Description	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). If any RADIUS attributes were received during initial setup (ESM strings, IP information), re-authentication always occurs. The no form of the command reverts to the default value.
Default	disabled

user-name-format

Syntax	user-name-format <i>format</i> [append <i>domain-name</i>] no user-name-format
Context	config>subscr-mgmt>auth-policy
Description	This command defines the format of the “user-name” field in the session authentication request sent to the RADIUS server. The no form of the command switches to the default format, mac .
Default	By default, the MAC Source Address of the DHCP DISCOVER message is used in the user-name field.
Parameters	<i>format</i> — Specifies the user name format in RADIUS message.
Values	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. circuit-id — If the system serves as a DHCP relay server which inserts option 82 information, 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.

tuple — The concatenation of MAC source address and circuit-ID are used in the user-name field.

ascii-converted-circuit-id — Identical to circuit-id, but the user name will be sent to the RADIUS server as a string of hex digits, 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.

append *domain-name* — When specified, the system will append a domain-name using @ as a separator between the user-name and the domain-name sections. The domain-name is defined based in the **local-user-db** entry the given host matches. If no domain-name is available, nothing is appended.

Filter Commands

filter

Syntax	filter
Context	configure
Description	This command manages the configuration of filters.

copy

Syntax	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]
Context	configure>filter
Description	This command copies filters and its entries.
Parameters	<i>src-filter-id</i> — Specifies the source filter ID. Values 1..65535 <i>src-entry-id</i> — Specifies the source entry ID. Values 1..65535 <i>dst-filter-id</i> — Specifies the destination filter ID. Values 1..65535 <i>dst-entry-id</i> — Specifies the destination entry ID. Values 1..65535 overwrite — Specifies an overwrite.

ip-filter

Syntax	ip-filter <i>filter-id</i> [create] no ip-filter <i>filter-id</i>
Context	configure>filter
Description	This command configures an IP filter.
Parameters	<i>filter-id</i> — Specifies the filter ID. Values 1..65535

ipv6-filter

Syntax	ipv6-filter <i>ipv6-filter-id</i> [create] no ipv6-filter <i>ipv6-filter-id</i>
Context	configure>filter
Description	This command configures an IPv6 filter.
Parameters	<i>filter-id</i> — Specifies the filter ID.
Values	1..65535

default-action

Syntax	default-action <i>drop forward</i>
Context	configure>filter>ip-filter configure>filter>ipv6-filter
Description	This command configures default-action for the IP or IPv6 filter.
Parameters	<i>drop forward</i> — This keyword specifies the filter action.

entry

Syntax	entry <i>entry-id</i> [time-range <i>time-range-name</i>] [create] no entry <i>entry-id</i>
Context	configure>filter>ip-filter configure>filter>ipv6-filter
Description	This command configures an IP or IPv6 filter entry.
Parameters	<i>entry-id</i> — Specifies the entry ID.
Values	1..65535
	<i>time-range-name</i> — Specifies the time range name.
Values	32 charas max

action

Syntax	action <i>drop forward</i> no action
Context	config>filter>ip-filter>entry config>filter>ipv6-filter>entry
Description	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 **dhcp**
no dhcp

Context config>filter>ip-filter>entry>match
 config>filter>ipv6-filter>entry>match

Description This command configures DSCP match condition.

dst-ip

Syntax	dst-ip no dst-ip
Context	config>filter>ip-filter>entry>match config>filter>ipv6-filter>entry>match
Description	This command configures the destination IP or IPv6 address match condition.

dst-port

Syntax	dst-port 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	icmp-code 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	icmp-type 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	src-ip 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 **src-port**
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

Syntax **tcp-ack**
no tcp-ack

Context config>filter>ip-filter>entry>match
config>filter>ipv6-filter>entry>match

Description This command configures the TCP ACK match condition.

tcp-syn

Syntax **tcp-syn**
no tcp-syn

Context config>filter>ip-filter>entry>match
config>filter>ipv6-filter>entry>match

Description This command configures the TCP SYN match condition.

group-inserted-entries

Syntax **group-inserted-entries application application location location**

Context config>filter>ip-filter
config>filter>ipv6-filter

Description This command groups auto-inserted entries.

Parameters *application* — Specifies the application.

Values radius | credit-control

location — Specifies the location.

Values top | bottom

renum

Syntax	renum <i>old-entry-id</i> <i>new-entry-id</i>
Context	config>filter>ip-filter config>filter>ipv6-filter
Description	This command renumbers an IP or IPv6 filter entry.
Parameters	<i>old-entry-id</i> — Specifies the old entry ID to be renumbered. Values 1..65535 <i>new-entry-id</i> — Specifies the new entry ID. Values 1..65535

scope

Syntax	scope <i>exclusive</i> <i>template</i> no scope
Context	config>filter>ip-filter config>filter>ipv6-filter
Description	This command configures the scope for the IP or IPv6 filter.
Parameters	<i>exclusive</i> <i>template</i> — Specifies the type of policy.

sub-insert-radius

Syntax	sub-insert-radius start-entry <i>entry-id</i> count <i>count</i> no sub-insert-radius
Context	config>filter>ip-filter config>filter>ipv6-filter
Description	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 no 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.
Default	per default insertion is disabled
Parameters	<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

Syntax	sub-insert-credit-control start-entry <i>entry-id</i> count <i>count</i> no sub-insert-credit-control
Context	config>filter>ip-filter config>filter>ipv6-filter
Description	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 no 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.
Default	per default insertion is disabled
Parameters	<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-wmark

Syntax	sub-insert-wmark [low <i>percentage</i>] [high <i>percentage</i>] no sub-insert-wmark
Context	config>filter>ip-filter config>filter>ipv6-filter
Description	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 no version of the command sets the default values for the respective thresholds.
Default	low - 90% high - 95%
Parameters	<i>percentage</i> — Defines in percentage the threshold used to raise an alarm. Values 1-100, integer

RADIUS Accounting Policy Custom Record Commands

custom-record

Syntax	[no] custom-record
Context	config>subscr-mgmt>acct-plcy
Description	This command enables the context to configure the layout and setting for a custom accounting record associated with this accounting policy. The no form of the command reverts the configured values to the defaults.

override-counter

Syntax	[no] override-counter <i>override-counter-id</i>
Context	config>log>acct-policy>cr
Description	This command enables the context to configure Application Assurance override counter parameters. The no form of the command removes the ID from the configuration.
Parameters	<i>override-counter-id</i> — Specifies the override counter ID. Values 1 — 8

e-counters

Syntax	e-counters [all] no e-counters
Context	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
Description	This command configures egress counter parameters for this custom record. The no form of the command
Parameters	all — Includes all counters.

i-counters

Syntax	i-counters [all] no i-counters
Context	config>log>acct-policy>cr>override-cntr config>log>acct-policy>cr>ref-override-cntr config>log>acct-policy>cr>ref-queue
Description	This command configures ingress counter parameters for this custom record. The no form of the command
Parameters	all — Includes all counters.

queue

Syntax	[no] queue <i>queue-id</i>
Context	config>log>acct-policy>cr
Description	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 no form of the command reverts to the default value
Parameters	<i>queue-id</i> — Specifies the queue-id for which counters will be collected in this custom record.

in-profile-octets-discarded-count

Syntax	[no] in-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 in-profile octets discarded count. The no form of the command excludes the in-profile octets discarded count.

in-profile-octets-forwarded-count

Syntax	[no] in-profile-octets-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 octets forwarded count. 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. 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. 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.

The **no** form of the command excludes the out of profile packets discarded count.

out-profile-octets-forwarded-count

Syntax	[no] out-profile-octets-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 out of profile octets forwarded count. The no form of the command excludes the out of profile octets forwarded count.

out-profile-packets-discarded-count

Syntax	[no] out-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 out of profile packets discarded count. The no form of the command excludes the out of profile packets discarded count.

out-profile-packets-forwarded-count

Syntax	[no] out-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 out of profile packets forwarded count. The no form of the command excludes the out of profile packets forwarded count. The no form of the command

all-octets-offered-count

Syntax	[no] all-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 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. 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

Syntax	[no] high-packets-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 packets discarded count. The no form of the command excludes the high packets discarded count.
Default	no high-packets-discarded-count

high-packets-offered-count

Syntax	[no] high-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 the high packets offered count. The no form of the command excludes the high packets offered count.
Default	no high-packets-offered -count

in-profile-octets-forwarded-count

Syntax	[no] in-profile-octets-forwarded-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 in profile octets forwarded count. The no form of the command excludes the in profile octets forwarded count.
Default	no in-profile-octets-forwarded-count

in-profile-packets-forwarded-count

Syntax	[no] in-profile-packets-forwarded-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 in profile packets forwarded count. The no form of the command excludes the in profile packets forwarded count.
Default	no in-profile-packets-forwarded-count

low-octets-discarded-count

Syntax	[no] low-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 low octets discarded count. The no form of the command excludes the low octets discarded count.
Default	no low-octets-discarded-count

low-packets-discarded-count

Syntax	[no] low-packets-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 low packets discarded count. The no form of the command excludes the low packets discarded count.
Default	no low-packets-discarded-count

low-octets-offered-count

Syntax	[no] low-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 low octets discarded count. The no form of the command excludes the low octets discarded count.

low-packets-offered-count

Syntax	[no] low-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 the low packets discarded count. The no form of the command excludes the low packets discarded count.

out-profile-octets-forwarded-count

Syntax	[no] out-profile-octets-forwarded-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 out of profile octets forwarded count. The no form of the command excludes the out of profile octets forwarded count.
Default	no out-profile-octets-forwarded-count

out-profile-packets-forwarded-count

Syntax	[no] out-profile-packets-forwarded-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 out of profile packets forwarded count. 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	ref-aa-specific-counter any no ref-aa-specific-counter
Context	config>log>acct-policy>cr
Description	This command The no form of the command

ref-override-counter

Syntax	ref-override-counter <i>ref-override-counter-id</i> ref-override-counter all no ref-override-counter
Context	config>log>acct-policy>cr
Description	This command configures a reference override counter. The no form of the command reverts to the default value.
Default	no ref-override-counter

ref-queue

Syntax	ref-queue <i>queue-id</i> ref-queue all no ref-queue
Context	config>log>acct-policy>cr
Description	This command configures a reference queue. The no form of the command reverts to the default value.
Default	no ref-queue

significant-change

Syntax	significant-change <i>delta</i> no significant-change
Context	config>log>acct-policy>cr
Description	This command configures the significant change required to generate the record.
Parameters	<i>delta</i> — Specifies the delta change (significant change) that is required for the custom record to be written to the xml file. Values 0 — 4294967295

BGP Peering Policy Commands

bgp-peering-policy

Syntax	bgp-peering-policy <i>policy-name</i> [create] no bgp-peering-policy <i>policy-name</i>
Context	config>subscr-mgmt
Description	This command configures the name of the BGP peering policy.
Parameters	<i>policy-name</i> — Specifies the BGP peer policy name up to 32 characters in length. create —

advertise-inactive

Syntax	[no] advertise-inactive
Context	config>subscr-mgmt>bgp-prng-plcy
Description	This command enables or disables the advertising of inactive BGP routers to other BGP peers. By default, BGP only advertises BGP routes to other BGP peers if a given BGP route is chosen by the route table manager as the most preferred route within the system and is active in the forwarding plane. This command allows system administrators to advertise a BGP route even though it is not the most preferred route within the system for a given destination.
Default	no advertise-inactive

aggregator-id-zero

Syntax	[no] aggregator-id-zero
Context	config>subscr-mgmt>bgp-prng-plcy
Description	This command is used to set the router ID in the BGP aggregator path attribute to zero when BGP aggregates routes. This prevents different routers within an AS from creating aggregate routes that contain different AS paths. When BGP is aggregating routes, it adds the aggregator path attribute to the BGP update messages. By default, BGP adds the AS number and router ID to the aggregator path attribute. When this command is enabled, BGP adds the router ID to the aggregator path attribute. The no form of the command used at the global level reverts to default where BGP adds the AS number and router ID to the aggregator path attribute.
Default	no aggregator-id-zero — BGP adds the AS number and router ID to the aggregator path attribute.

as-override

Syntax	[no] as-override
Context	config>subscr-mgmt>bgp-prng-plcy
Description	This command replaces all instances of the peer's AS number with the local AS number in a BGP route's AS_PATH. This command breaks BGP's loop detection mechanism. It should be used carefully.
Default	as-override is not enabled by default.

auth-keychain

Syntax	auth-keychain <i>name</i> no auth-keychain
Context	config>subscr-mgmt>bgp-prng-plcy
Description	This command configures the BGP authentication key for all peers. The keychain allows the rollover of authentication keys during the lifetime of a session.
Default	no auth-keychain
Parameters	<i>name</i> — Specifies the name of an existing keychain, up to 32 characters, to use for the specified TCP session or sessions.

authentication-key

Syntax	authentication-key [<i>authentication-key</i> <i>hash-key</i>] [hash hash2] no authentication-key
Context	config>subscr-mgmt>bgp-prng-plcy
Description	This command configures the BGP authentication key. 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. The no form of the command removes the authentication password from the configuration and effectively disables authentication.
Default	Authentication is disabled and the authentication password is empty.
Parameters	<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 (“ ”). <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 (“ ”).

This is useful when a user must configure the parameter, but, for security purposes, the actual unencrypted key value is not provided.

hash — 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. For security, all keys are stored in encrypted form in the configuration file with the **hash** parameter specified.

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.

cluster

Syntax	cluster <i>cluster-id</i> no cluster
Context	config>subscr-mgmt>bgp-prng-plcy
Description	This command configures the cluster ID for a route reflector server. Route reflectors are used to reduce the number of IBGP sessions required within an AS. Normally, all BGP speakers within an AS must have a BGP peering with every other BGP speaker in an AS. A route reflector and its clients form a cluster. Peers that are not part of the cluster are considered to be non-clients. When a route reflector receives a route, first it must select the best path from all the paths received. If the route was received from a non-client peer, then the route reflector sends the route to all clients in the cluster. If the route came from a client peer, the route reflector sends the route to all non-client peers and to all client peers except the originator. For redundancy, a cluster can have multiple route reflectors. Confederations can also be used to remove the full IBGP mesh requirement within an AS. The no form of the command deletes the cluster ID and effectively disables the Route Reflection for the given group.
Default	no cluster — No cluster ID is defined.
Parameters	<i>cluster-id</i> — The route reflector cluster ID is expressed in dot decimal notation. Values Any 32 bit number in dot decimal notation. (0.0.0.1 — 255.255.255.255)

connect-retry

Syntax	connect-retry <i>seconds</i> no connect-retry
Context	config>subscr-mgmt>bgp-prng-plcy
Description	This command configures the BGP connect retry timer value in seconds. When this timer expires, BGP tries to reconnect to the configured peer. The no form of the command used at the global level reverts to the default value.

Triple Play Subscriber Management Configuration Commands

Default	120 seconds
Parameters	<i>seconds</i> — The BGP Connect Retry timer value in seconds, expressed as a decimal integer.
Values	1 — 65535

damping

Syntax	[no] damping
Context	config>subscr-mgmt>bgp-prng-plcy
Description	This command enables BGP route damping for learned routes which are defined within the route policy. Use damping to reduce the number of update messages sent between BGP peers and reduce the load on peers without affecting the route convergence time for stable routes. Damping parameters are set via route policy definition. The no form of the command used at the global level disables route damping. When damping is enabled and the route policy does not specify a damping profile, the default damping profile is used. This profile is always present and consists of the following parameters: Half-life: 15 minutes Max-suppress: 60 minutes Suppress-threshold:3000 Reuse-threshold 750
Default	no damping — Learned route damping is disabled.

disable-4byte-asn

Syntax	[no] disable-4byte-asn
Context	config>subscr-mgmt>bgp-prng-plcy
Description	This command disables the use of 4-byte ASNs. It can be configured at all 3 level of the hierarchy so it can be specified down to the per peer basis. If this command is enabled 4-btye ASN support should not be negotiated with the associated remote peer(s). The no form of the command resets the behavior to the default which is to enable the use of 4-byte ASN.

disable-client-reflect

Syntax	[no] disable-client-reflect
Context	config>subscr-mgmt>bgp-prng-plcy

Description	This command disables the reflection of routes by the route reflector to the group or neighbor. This only disables the reflection of routes from other client peers. Routes learned from non-client peers are still reflected to all clients. The no form re-enables client reflection of routes.
Default	no disable-client-reflect — Client routes are reflected to all client peers.

disable-communities

Syntax	disable-communities [standard] [extended] no disable-communities
Context	config>subscr-mgmt>bgp-prng-plcy
Description	This command configures BGP to disable sending communities.
Parameters	standard — Specifies standard communities that existed before VPRNs or 2547. extended — Specifies BGP communities used were expanded after the concept of 2547 was introduced, to include handling the VRF target.

disable-fast-external-failover

Syntax	[no] disable-fast-external-failover
Context	config>subscr-mgmt>bgp-prng-plcy
Description	This command configures BGP fast external failover.

export

Syntax	export <i>policy</i> [<i>policy</i> ...] no export
Context	config>subscr-mgmt>bgp-prng-plcy
Description	This command specifies the export policies to be used to control routes advertised to BGP neighbors. When multiple policy names are specified, the policies are evaluated in the order they are specified. A maximum of five (5) policy names can be configured. The first policy that matches is applied. Note that if a non-existent route policy is applied to a VPRN instance, the CLI generates a warning message. This message is only generated at an interactive CLI session and the route policy association is made. No warning message is generated when a non-existent route policy is applied to a VPRN instance in a configuration file or when SNMP is used. The no form of this command removes all route policy names from the export list.
Default	no export — BGP advertises routes from other BGP routes but does not advertise any routes from other protocols unless directed by an export policy.

Parameters *policy* — A route policy statement name.

hold-time

Syntax **hold-time** *seconds*
no hold-time

Context config>subscr-mgmt>bgp-prng-plcy

Description This command configures the BGP hold time, expressed in seconds.

The BGP hold time specifies the maximum time BGP waits between successive messages (either keepalive or update) from its peer, before closing the connection.

Even though the router OS implementation allows setting the keepalive time separately, the configured keepalive timer is overridden by the hold-time value under the following circumstances:

1. If the specified hold-time is less than the configured keepalive time, then the operational keepalive time is set to a third of the hold-time; the configured keepalive time is not changed.
2. If the hold-time is set to zero, then the operational value of the keepalive time is set to zero; the configured keepalive time is not changed. This means that the connection with the peer is up permanently and no keepalive packets are sent to the peer.

The **no** form of the command used at the global level reverts to the default value.

Default 90 seconds

Parameters *seconds* — The hold-time, in seconds, expressed as a decimal integer. A value of 0 indicates the connection to the peer is up permanently.

Values 0, 3 — 65535

import

Syntax **import** *policy* [*policy*...]
no import

Context config>subscr-mgmt>bgp-prng-plcy

Description This command specifies the import policies to be used to control routes advertised to BGP neighbors. Route policies are configured in the **config>router>policy-options** context. When multiple policy names are specified, the policies are evaluated in the order they are specified. A maximum of five (5) policy names can be specified. The first policy that matches is applied.

The **no** form of this command removes all route policy names from the import list.

Default no import — BGP accepts all routes from configured BGP neighbors. Import policies can be used to limit or modify the routes accepted and their corresponding parameters and metrics.

Parameters *policy* — A route policy statement name.

keepalive

Syntax	keepalive <i>seconds</i> no keepalive
Context	config>subscr-mgmt>bgp-prng-plcy
Description	This command configures the BGP keepalive timer. A keepalive message is sent every time this timer expires. The keepalive value is generally one-third of the hold-time interval. Even though the OS implementation allows the keepalive value and the hold-time interval to be independently set, under the following circumstances, the configured keepalive value is overridden by the hold-time value: If the specified keepalive value is greater than the configured hold-time, then the specified value is ignored, and the keepalive is set to one third of the current hold-time value. If the specified hold-time interval is less than the configured keepalive value, then the keepalive value is reset to one third of the specified hold-time interval. If the hold-time interval is set to zero, then the configured value of the keepalive value is ignored. This means that the connection with the peer is up permanently and no keepalive packets are sent to the peer. The no form of the command used at the global level reverts to the default value..
Default	30 seconds
Parameters	<i>seconds</i> — The keepalive timer in seconds, expressed as a decimal integer. Values 0 — 21845

local-address

Syntax	local-address <i>ip-address</i> no local-address
Context	config>subscr-mgmt>bgp-prng-plcy
Description	Configures the local IP address used by the group or neighbor when communicating with BGP peers. Outgoing connections use the local-address as the source of the TCP connection when initiating connections with a peer. When a local address is not specified, the 7750 SR OS uses the system IP address when communicating with IBGP peers and uses the interface address for directly connected EBGp peers. This command is used at the neighbor level to revert to the value defined under the group level. The no form of the command removes the configured local-address for BGP. The no form of the command used at the group level reverts to the value defined at the global level. The no form of the command used at the neighbor level reverts to the value defined at the group level.
Default	no local-address — The router ID is used when communicating with IBGP peers and the interface address is used for directly connected EBGp peers. <i>ip-address</i> — The local address expressed in dotted decimal notation. Allowed values are a valid routable IP address on the router, either an interface or system IP address.

local-as

Syntax	local-as <i>as-number</i> [private] no local-as
Context	config>subscr-mgmt>bgp-prng-plcy
Description	This command configures a BGP virtual autonomous system (AS) number. In addition to the AS number configured for BGP in the <code>config>router>autonomous-system</code> context, a virtual (local) AS number is configured. The virtual AS number is added to the as-path message before the router's AS number makes the virtual AS the second AS in the as-path. This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). Thus, by specifying this at each neighbor level, it is possible to have a separate as-number per EBGp session. When a command is entered multiple times for the same AS, the last command entered is used in the configuration. The private attribute can be added or removed dynamically by reissuing the command. Changing the local AS at the global level in an active BGP instance causes the BGP instance to restart with the new local AS number. Changing the local AS at the global level in an active BGP instance causes BGP to re-establish the peer relationships with all peers in the group with the new local AS number. Changing the local AS at the neighbor level in an active BGP instance causes BGP to re-establish the peer relationship with the new local AS number. This is an optional command and can be used in the following circumstance: Provider router P is moved from AS1 to AS2. The customer router that is connected to P, however, is configured to belong to AS1. To avoid reconfiguring the customer router, the local-as value on router P can be set to AS1. Thus, router P adds AS1 to the as-path message for routes it advertises to the customer router. The no form of the command used at the global level will remove any virtual AS number configured. The no form of the command used at the group level reverts to the value defined at the global level. The no form of the command used at the neighbor level reverts to the value defined at the group level.
Default	no local-as
Parameters	<i>as-number</i> — The virtual autonomous system number, expressed as a decimal integer. Values 1 — 65535 private — Specifies the local-as is hidden in paths learned from the peering.

local-preference

Syntax	local-preference <i>local-preference</i> no local-preference
Context	config>subscr-mgmt>bgp-prng-plcy
Description	This command enables setting the BGP local-preference attribute in incoming routes if not specified and configures the default value for the attribute. This value is used if the BGP route arrives from a BGP peer without the local-preference integer set.

The specified value can be overridden by any value set via a route policy.

The **no** form of the command at the global level specifies that incoming routes with local-preference set are not overridden and routes arriving without local-preference set are interpreted as if the route had local-preference value of 100.

Default	no local-preference — Does not override the local-preference value set in arriving routes and analyze routes without local preference with value of 100.
Parameters	<i>local-preference</i> — The local preference value to be used as the override value, expressed as a decimal integer.
Values	0 — 4294967295

loop-detect

Syntax	loop-detect {drop-peer discard-route ignore-loop off} no loop-detect
Context	config>subscr-mgmt>bgp-prng-plcy
Description	This command configures how the BGP peer session handles loop detection in the AS path. Note that dynamic configuration changes of loop-detect are not recognized. The no form of the command used at the global level reverts to default, which is loop-detect ignore-loop .
Default	loop-detect ignore-loop
Parameters	drop-peer — Sends a notification to the remote peer and drops the session. discard-route — Discards routes received with loops in the AS path. ignore-loop — Ignores routes with loops in the AS path but maintains peering. off — Disables loop detection.

med-out

Syntax	med-out {number igp-cost} no med-out
Context	config>subscr-mgmt>bgp-prng-plcy
Description	This command enables advertising the Multi-Exit Discriminator (MED) and assigns the value used for the path attribute for the MED advertised to BGP peers if the MED is not already set. The specified value can be overridden by any value set via a route policy. The no form of the command used at the global level reverts to default where the MED is not advertised. no med-out
Parameters	<i>number</i> — The MED path attribute value, expressed as a decimal integer.

Values 0 — 4294967295

igp-cost — The MED is set to the IGP cost of the given IP prefix.

min-as-origination

Syntax	min-as-origination <i>seconds</i> no min-as-origination
Context	config>subscr-mgmt>bgp-prng-plcy
Description	This command configures the minimum interval, in seconds, at which a path attribute, originated by the local router, can be advertised to a peer. The no form of the command used at the global level reverts to default.
Default	15 seconds
Parameters	<i>seconds</i> — The minimum path attribute advertising interval in seconds, expressed as a decimal integer. Values 2 — 255

min-route-advertisement

Syntax	min-route-advertisement <i>seconds</i> no min-route-advertisement
Context	config>subscr-mgmt>bgp-prng-plcy
Description	This command configures the minimum interval, in seconds, at which a prefix can be advertised to a peer. The no form of the command reverts to default values.
Default	30 seconds
Parameters	<i>seconds</i> — The minimum route advertising interval, in seconds, expressed as a decimal integer. Values 1 — 255

multihop

Syntax	multihop <i>ttl-value</i> no multihop
Context	config>subscr-mgmt>bgp-prng-plcy
Description	This command configures the time to live (TTL) value entered in the IP header of packets sent to an EBGp peer multiple hops away.

This parameter is meaningful only when configuring EBGp peers. It is ignored if set for an IBGP peer.

The **no** form of the command is used to convey to the BGP instance that the EBGp peers are directly connected.

The **no** form of the command reverts to default values.

Default	1 — EBGp peers are directly connected. 64 — IBGP
Parameters	<i>ttl-value</i> — The TTL value, expressed as a decimal integer.
Values	1 — 255

next-hop-self

Syntax	[no] next-hop-self
Context	config>subscr-mgmt>bgp-prng-plcy
Description	This command configures the neighbor to always set the NEXTHop path attribute to its own physical interface when advertising to a peer. The no form of the command disables the command.
Default	no next-hop-self

passive

Syntax	[no] passive
Context	config>subscr-mgmt>bgp-prng-plcy
Description	This command enables the passive mode for the BGP neighbors. The no form of the command disables the passive mode.
Default	no passive

peer-as

Syntax	peer-as as-number no peer-as
Context	config>subscr-mgmt>bgp-prng-plcy
Description	This command configures the autonomous system number for the remote peer. The peer AS number must be configured for each configured peer. The no form of the command removes the <i>as-number</i> from the configuration.
Default	No AS numbers are defined.

Parameters *as-number* — Specifies the AS number for the remote peer.

Values 1 — 4294967295

preference

Syntax **[no] preference** *preference*

Context config>subscr-mgmt>bgp-prng-plcy

Description This command configures the route preference for routes learned from the configured peer(s).
The lower the preference the higher the chance of the route being the active route. The OS assigns BGP routes highest default preference compared to routes that are direct, static or learned via MPLS or OSPF.
The **no** form of the command used at the global level reverts to default value.

Default 170

Parameters *preference* — The route preference, expressed as a decimal integer.

Values 1 — 255

prefix-limit

Syntax **prefix-limit** *limit*
no prefix-limit

Context config>subscr-mgmt>bgp-prng-plcy

Description This command configures the maximum number of routes BGP can learn from a peer.
When the number of routes reaches 90% of this limit, an SNMP trap is sent. When the limit is exceeded, the BGP peering is dropped and disabled.
The **no** form of the command removes the **prefix-limit**.

Default no prefix-limit

Parameters *limit* — The number of routes that can be learned from a peer, expressed as a decimal integer.

Values 1 — 4294967295

remove-private

Syntax **[no] remove-private**

Context config>subscr-mgmt>bgp-prng-plcy

Description This command allows private AS numbers to be removed from the AS path before advertising them to BGP peers.

The OS software recognizes the set of AS numbers that are defined by IANA as private. These are AS numbers in the range 64512 through 65535, inclusive.

The **no** form of the command used at the global level reverts to default value.

Default **no remove-private** — Private AS numbers will be included in the AS path attribute.

type

Syntax **[no] type {internal | external}**

Context config>subscr-mgmt>bgp-prng-plcy

Description This command designates the BGP peer as type internal or external.

The type of **internal** indicates the peer is an IBGP peer while the type of external indicates that the peer is an EBGP peer.

By default, the OS derives the type of neighbor based on the local AS specified. If the local AS specified is the same as the AS of the router, the peer is considered **internal**. If the local AS is different, then the peer is considered **external**.

The **no** form of the command used at the group level reverts to the default value.

Default **no type** — Type of neighbor is derived on the local AS specified.

Parameters **internal** — Configures the peer as internal.
external — Configures the peer as external.

ttl-security

Syntax **ttl-security min-ttl-value**
no ttl-security

Context config>subscr-mgmt>bgp-prng-plcy

Description Configure TTL security parameters for incoming packets.

Parameters *min-ttl-value* — Specify the minimum TTL value for an incoming BGP packet.

Values 1 — 255

diameter-policy

Syntax **[no] diameter-policy diameter-policy-name**

Context config>sub-mgmt

Triple Play Subscriber Management Configuration Commands

Description	This command configures a diameter policy. The no form of the command removes the policy name from the configuration.
Default	no policy is created
Parameters	<i>diameter-policy-name</i> — Specifies a diameter policy name up to 32 characters in length.

Multi-Chassis Redundancy Commands

redundancy

Syntax	redundancy
Context	config
Description	This command allows the user to perform redundancy operations.
Parameters	force-switchover — Forces a switchover to the standby CPM card Values now keyword - switch to standby CPM) NOTE: 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)? synchronize — Synchronizes the secondary CPM. Values <boot-env config> : keywords

synchronize

Syntax	synchronize {boot-env config}
Context	config>redundancy
Description	This command performs a synchronization of the standby CPM's images and/or config files to the active CPM. Either the boot-env or config parameter must be specified. In the config>redundancy context, this command performs an automatically triggered standby CPM synchronization. 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 CPMs have identical operational parameters. This includes the saved configuration, CPM and IOM images. 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. 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. 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).
Default	enabled

Triple Play Subscriber Management Configuration Commands

Parameters	boot-env — Synchronizes all files required for the boot process (loader, BOF, images, and configuration files).
	config — Synchronize only the primary, secondary, and tertiary configuration files.
	Default config

multi-chassis

Syntax	multi-chassis
Context	config>redundancy
Description	This command enables the context to configure multi-chassis parameters.

peer

Syntax	[no] peer <i>ip-address</i>
Context	config>redundancy>multi-chassis
Description	This command configures a multi-chassis redundancy peer.
Parameters	<i>ip-address</i> — Specifies a peer IP address. Multicast address are not allowed.

authentication-key

Syntax	authentication-key [<i>authentication-key</i> <i>hash-key</i>] [hash hash2] no authentication-key
Context	config>redundancy>multi-chassis>peer
Description	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.
Parameters	<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 (“ ”).
	hash — 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.
	hash2 — 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-lag

Syntax	[no] mc-lag
Context	config>redundancy>multi-chassis>peer>mc-lag
Description	This command enables the context to configure multi-chassis LAG operations and related parameters. The no 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

Syntax	hold-on-neighbor-failure <i>multiplier</i> no hold-on-neighbor-failure
Context	config>redundancy>multi-chassis>peer>mc-lag
Description	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. The no form of this command sets this parameter to default value.
Default	3
Parameters	<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. Values 2 — 25

keep-alive-interval

Syntax	keep-alive-interval <i>interval</i> no keep-alive-interval
Context	config>redundancy>multi-chassis>peer>mc-lag
Description	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. The no form of this command sets the interval to default value
Default	1s (10 hundreds of milliseconds means interval value of 10)

Parameters *interval* — The time interval expressed in deci-seconds

Values 5 — 500

lag

Syntax **lag** *lag-id* **lacp-key** *admin-key* **system-id** *system-id* [**remote-lag** *lag-id*] **system-priority** *system-priority*
no lag *lag-id*

Context config>redundancy>multi-chassis>peer>mc-lag

Description 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.

The same **lacp-key**, **system-id**, and **system-priority** 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 (**lacp-key**, **system-id**, **system-priority**) must be configured the same on both nodes of the same redundant pair.

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.

Default none

Parameters *lag-id* — The LAG identifier, expressed as a decimal integer. Specifying the *lag-id* allows the mismatch between lag-id on redundant-pair. If no **lag-id** is specified it is assumed that neighbor system uses the same *lag-id* 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.).

Values 1 — 200

lacp-key *admin-key* — 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.

Values 1 — 65535

system-id *system-id* — Specifies a 6 byte value expressed in the same notation as MAC address

Values xx:xx:xx:xx:xx:xx - xx [00..FF]

remote-lag *lag-id* — Specifies the LAG ID on the remote system.

Values 1 — 200

system-priority *system-priority* — Specifies the system priority to be used in the context of the MC-LAG. The partner system will consider all ports using the same **lacp-key**, **system-id**, and **system-priority** as part of the same LAG.

Values 1 — 65535

source-address

Syntax	source-address <i>ip-address</i> no source-address
Context	config>redundancy>multi-chassis>peer
Description	This command specifies the source address used to communicate with the multi-chassis peer.
Parameters	<i>ip-address</i> — Specifies the source address used to communicate with the multi-chassis peer.

sync

Syntax	[no] sync
Context	config>redundancy>multi-chassis>peer
Description	This command enables the context to configure synchronization parameters.

igmp

Syntax	[no] igmp
Context	config>redundancy>multi-chassis>peer>sync
Description	This command specifies whether IGMP protocol information should be synchronized with the multi-chassis peer.
Default	no igmp

igmp-snooping

Syntax	[no] igmp-snooping
Context	config>redundancy>multi-chassis>peer>sync
Description	This command specifies whether IGMP snooping information should be synchronized with the multi-chassis peer.
Default	no igmp-snooping

local-dhcp-server

Syntax	[no] local-dhcp-server
Context	config>redundancy>multi-chassis>peer>sync
Description	This command synchronizes DHCP server information.

mc-ring

Syntax	[no] mc-ring
Context	config>redundancy>multi-chassis>peer>sync
Description	This command synchronizes mc-ring information.

mld-snooping

Syntax	[no] mld-snooping
Context	config>redundancy>multi-chassis>peer>sync
Description	This command synchronizes MLD snooping information.

port

Syntax	port <i>[port-id lag-id]</i> [sync-tag sync-tag] [create] no port <i>[port-id lag-id]</i>
Context	config>redundancy>multi-chassis>peer>sync
Description	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.
Parameters	<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. sync-tag sync-tag — Specifies a synchronization tag to be used while synchronizing this port with the multi-chassis peer.

range

Syntax	range <i>encap-range</i> sync-tag <i>sync-tag</i> no range <i>encap-range</i>						
Context	config>redundancy>multi-chassis>peer>sync>port						
Description	This command configures a range of encapsulation values.						
Parameters	<i>encap-range</i> — Specifies a range of encapsulation values on a port to be synchronized with a multi-chassis peer. <table><tr><td>Values</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> sync-tag <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.	Values	Dot1Q	<i>start-vlan-end-vlan</i>		QinQ	<i>Q1.start-vlan-Q1.end-vlan</i>
Values	Dot1Q	<i>start-vlan-end-vlan</i>					
	QinQ	<i>Q1.start-vlan-Q1.end-vlan</i>					

srrp

Syntax	[no] srrp
Context	config>redundancy>multi-chassis>peer>sync
Description	This command specifies whether subscriber routed redundancy protocol (SRRP) information should be synchronized with the multi-chassis peer.
Default	no srrp

sub-host-trk

Syntax	[no] sub-host-trk
Context	config>redundancy>multi-chassis>peer>sync
Description	This command synchronizes subscriber host tracking information.

sub-mgmt

Syntax	[no] sub-mgmt
Context	config>redundancy>multi-chassis>peer>sync
Description	This command specifies whether subscriber management information should be synchronized with the multi-chassis peer.
Default	no sub-mgmt

mc-ring

Syntax	mc-ring
Context	config>redundancy>multi-chassis>peer
Description	This command enables the context to configure the multi-chassis ring parameters.
Default	mc-ring

ring

Syntax	[no] ring sync-tag [create]
Context	config>redundancy>multi-chassis>peer>mcr
Description	This command configures a multi-chassis ring. The no form of the command removes the sync-tag from the configuration.

Default none

l3-ring

Syntax [no] l3-ring name [create]

Context config>redundancy>multi-chassis>peer>mcr

Parameters This command configures a layer 3 multi-chassis ring.

in-band-control-path

Syntax in-band-control-path

Context config>redundancy>multi-chassis>peer>mcr>ring
config>redundancy>multi-chassis>peer>mc>l3-ring

Description This command enables the context to configure control path parameters.

Default none

debounce

Syntax [no] debounce

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 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 <i>ip-int-name</i> 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	<i>ip-int-name</i> — Specifies an interface name up to 32 characters in length.

max-debounce-time

Syntax	max-debounce-time <i>max-debounce-time</i> 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	<i>max-debounce-time</i> — Specifies the maximum debounce time on the transition of the operational state of the inband control connection.
Values	5 — 200 seconds

service-id

Syntax	service-id <i>service-id</i> no service-id
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 service ID of the SAP used for the Ring-Node Connectivity Verification of this ring node.
Parameters	<i>service-id</i> — [Specifies an existing service ID or service name.
Values	service-id: 1 — 214748364 svc-name: A string up to 64 characters in length.

path-b

Syntax	[no] path-b
Context	config>redundancy>multi-chassis>peer>mcr>ring

Triple Play Subscriber Management Configuration Commands

Description	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.
Default	If not specified, the default is an empty set.

range

Syntax	[no] range <i>vlan-range</i>
Context	config>redundancy>multi-chassis>peer>mcr>ring>path-b config>redundancy>multi-chassis>peer>mcr>ring>path-excl
Description	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 with VLAN ID i, with i in [0..4094]. Setting the value to the empty string is equivalent to setting it to 512 zeroes.

ring-node

Syntax	[no] ring-node <i>ring-node-name</i>
Context	config>redundancy>mc>peer>mcr>ring
Description	This command specifies the unique name of a multi-chassis ring access node.

path-excl

Syntax	[no] path-excl
Context	config>redundancy>multi-chassis>peer>mcr>ring
Description	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.
Default	If not specified, the default is an empty set.

connectivity-verify

Syntax	connectivity-verify
Context	config>redundancy>multi-chassis>peer>mcr>ring config>redundancy>multi-chassis>peer>mc>l3-ring
Description	This command configures the node connectivity check.

interval

Syntax	interval <i>interval</i>
Context	config>redundancy>multi-chassis>peer>mcr>ring>>connectivity-verify config>redundancy>multi-chassis>peer>mc>l3-ring>connectivity-verify
Description	This command specifies the polling interval of the ring-node connectivity verification of this ring node.
Parameters	<i>interval</i> — Specifies the polling interval of the ring-node connectivity verification of this ring node.
Values	1 — 6000

service-id

Syntax	service-id <i>service-id</i> no service-id
Context	config>redundancy>mc>peer>mcr>ring-node>connect-verify config>redundancy>multi-chassis>peer>mc>l3-ring>connectivity-verify
Description	This command specifies the service ID of the SAP used for ring-node connectivity verification of this ring node.
Parameters	<i>service-id</i> — Specifies the service ID or service name.
Values	service-id: 1 — 214748364 svc-name: A string up to 64 characters in length.

src-ip

Syntax	src-ip <i>ip-address</i> no src-ip
Context	config>redundancy>mc>peer>mcr>ring-node>connect-verify config>redundancy>multi-chassis>peer>mc>l3-ring>connectivity-verify
Description	This command specifies the source IP address used in ring-node connectivity verification of this ring node.
Parameters	<i>ip-address</i> — Specifies the source IP address used in ring-node connectivity verification of this ring node.

src-mac

Syntax	src-mac <i>ieee-address</i> no src-mac
Context	config>redundancy>mc>peer>mcr>ring-node>connect-verify config>redundancy>multi-chassis>peer>mc>l3-ring>connectivity-verify
Description	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.
Parameters	<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

Syntax	vlan [0..4094]
Context	config>redundancy>mc>peer>mcr>ring-node>connect-verify config>redundancy>mc>peer>mcr>l3ring>node>cv
Description	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

Syntax	[no] srrp-instance <i>srrp-id</i>
Context	config>redundancy>multi-chassis>peer>mc>l3-ring
Description	This command configures an SRRP instance for layer 3 ring.
Parameters	<i>srrp-id</i> — Specifies the SRRP ID of this SRRP instance. Values 1 — 4294967295

Log Commands

accounting-policy

Syntax	[no] accounting-policy <i>acct-policy-id</i>
Context	config>log
Description	This command creates an access or network accounting policy. An accounting policy defines the accounting records that are created.
Parameters	<i>policy-id</i> — The policy ID that uniquely identifies the accounting policy, expressed as a decimal integer.
Values	1 — 99

collection-interval

Syntax	collection-interval <i>minutes</i> no collection-interval
Context	config>log>accounting-policy
Description	This command configures the accounting collection interval.
Parameters	<i>minutes</i> — indicates how frequently, in minutes, the statistics are collected and written to their destination.
Values	5 — 120

default

Syntax	[no] default
Context	config>log>accounting-policy This command contains the description for the accounting policy to be used with all SAPs that do not have an accounting policy.

record

Syntax	record <i>record-name</i> no record
Context	config>log>accounting-policy

Triple Play Subscriber Management Configuration Commands

Description	This command configures the accounting record type to be forwarded to the configured accounting file.
Parameters	<i>record-name</i> — service-ingress-octets service-egress-octets service-ingress-packets service-egress-packets network-ingress-octets network-egress-octets network-ingress-packets network-egress-packets compact-service-ingress-octets combined-service-ingress combined-network-ing-egr-octets combined-service-ing-egr-octets complete-service-ingress-egress complete-sdp-ingress-egress combined-sdp-ingress-egress complete-subscriber-ingress-egress aa-protocol aa-application aa-app-group aa-subscriber-protocol aa-subscriber-application aa-subscriber-app-group

shutdown

Syntax	[no] shutdown
Context	config>log>accounting-policy
Description	This command administratively enables or disables the accounting policy.

to

Syntax	to file <i>log-file-id</i>
Context	config>log>accounting-policy
Description	This command configures the log destination.
Parameters	<i>log-file-id</i> — Specifies the log ID.
Values	1 — 99

event-control

Syntax	event-control <i>application-id</i> [<i>event-name</i> <i>event-number</i>] [generate [<i>severity-level</i>] [<i>throttle</i>] event-control <i>application-id</i> [<i>event-name</i> <i>event-number</i>] suppress no event-control <i>application</i> [<i>event-name</i> <i>event-number</i>]
Context	config>log
Description	This command is used to specify that a particular event or all events associated with an application is either generated or suppressed. Events are generated by an application and contain an event number and description explaining the cause of the event. Each event has a default designation which directs it to be generated or suppressed. Events are generated with a default severity level that can be modified by using the <i>severity-level</i> option.

Events that are suppressed by default are typically used for debugging purposes. Events are suppressed at the time the application requests the event's generation. No event log entry is generated regardless of the destination. While this feature can save processor resources, there may be a negative effect on the ability to troubleshoot problems if the logging entries are squelched. In reverse, indiscriminate application may cause excessive overhead.

The **no** form of the command reverts the parameters to the default setting for events for the application or a specific event within the application. The severity, generate, and suppress options will also be reset to the initial values.

Default	Each event has a default suppress or generate state. To display a list of all events and the current configuration use the event-control command.
Parameters	<i>application-id</i> — The application whose events are affected by this event control filter. Default None, this parameter must be explicitly specified. Values A valid application name. To display a list of valid application names, use the applications command. Valid applications are: aps, atm, bgp, cflowd, chassis, debug, dhcp, efm_oam, filter, gsmp, igmp, igmp_snooping, ip, isis, lag, ldp, logger, mc_redundancy, mirror, mpls, ntp, oam, ospf, pim, port, ppp, rip, route_policy, rsvp, security, snmp, stp, svcmgr, system, user, vrrp, vrrt <i>event-name event-number</i> — To generate, suppress, or revert to default for a single event, enter the specific number or event short name. If no event number or name is specified, the command applies to all events in the application. To display a list of all event short names use the event-control command. Default none Values A valid event name or event number. <i>generate severity-level</i> — This keyword indicates that the specified events will generate logging events. Default generate <i>severity-name</i> — An ASCII string representing the severity level to associate with the specified generated events Default The system assigned severity name Values One of: cleared, indeterminate, critical, major, minor, warning. <i>suppress</i> — This keyword indicates that the specified events will not be logged. If the suppress keyword is not specified then the events are generated by default. Default generate throttle — Specifies whether or not events of this type will be throttled. By default, event throttling is off for each specific event type. It must be explicitly enabled for each event type where throttling is desired. This makes backwards compatibility easier to manage.

Subscriber Identification Policy Commands

sub-ident-policy

Syntax	[no] sub-ident-policy <i>sub-ident-policy-name</i>
Context	config>subscr-mgmt
Description	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. The subscriber identification policy default subscriber profile cannot be defined with the subscriber profile name default. 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	<i>sub-ident-policy-name</i> — Specifies the name of the subscriber identification policy.

app-profile-map

Syntax	app-profile-map
Context	config>subscr-mgmt>sub-ident-pol
Description	This command enables the context to configure an application profile mapping.

entry

Syntax	entry key <i>app-profile-string</i> app-profile <i>app-profile-name</i> no entry key <i>app-profile-string</i>
Context	config>subscr-mgmt>sub-ident-pol>app-profile-map
Description	This command configures an application profile string. The no form of the command removes the values from the configuration.

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 <i>sla-profile-string</i> sla-profile <i>sla-profile-name</i> no entry key <i>sla-profile-string</i>
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.

Subscriber Profile Commands

sub-profile

Syntax	[no] sub-profile <i>subscriber-profile-name</i>
Context	config>subscr-mgmt
Description	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. Subscribers are either explicitly mapped to a subscriber profile template or are dynamically associated by one of various non-provisioned subscriber profile definitions. A subscriber host can be associated with a subscriber profile in the following ways, listed from lowest to highest precedence: 1. The subscriber profile named default. 2. The subscriber profile defined as the subscriber SAP default. 3. The subscriber profile found by the subscriber identification policy sub-profile-map. 4. The subscriber profile found by the subscriber identification policy explicit map. 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. Default Subscriber profile: 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. The default profile has no effect on existing active subscriber on the system as they exist due to higher precedence mappings. Attempting to delete any subscriber profile (including the profile named default) while in use by existing active subscribers will fail.
Parameters	<i>subscriber-profile-name</i> — Specify the name of the subscriber profile.
Values	32 characters maximum, default

accounting-policy

Syntax	accounting-policy <i>acct-policy-id</i> no accounting-policy
Context	config>subscr-mgmt>sub-prof
Description	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 config>log context. The no form of this command removes the accounting policy association .
Default	no accounting policy
Parameters	<i>acct-policy-id</i> — Enter the accounting <i>policy-id</i> as configured in the config>log>accounting-policy context. Values 1 — 99

collect-stats

Syntax	[no] collect-stats
Context	config>subscr-mgmt>sub-prof
Description	When enabled, the agent collects non-RADIUS accounting statistics. When the no collect-stats 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 collect-stats command is issued then the counters written to the billing file include all the traffic while the no collect-stats command was in effect.
Default	collect-stats

egress-scheduler-policy

Syntax	egress-scheduler-policy <i>scheduler-policy-name</i> no egress-scheduler-policy
Context	config>subscr-mgmt>sub-prof
Description	This command specifies a scheduler policy to associate to the subscriber profile. Scheduler policies are configured in the configure>qos>scheduler>policy 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.
Parameters	<i>scheduler-policy-name</i> — Specify the scheduler policy name.

egress-agg-rate-limit

Syntax	egress-agg-rate-limit { max <i>kilobits-per-second</i> } [queue-frame-based-accounting] no egress-agg-rate-limit
Context	config>subscr-mgmt>sub-prof
Description	The egress-agg-rate-limit command is used to 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 will also include the frame encapsulation overhead within the octet counters. 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 cause all 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 <i>kilobits-per-second</i>} — The max keyword and <i>kilobits-per-second</i> 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. <i>kilobits-per-second</i> — 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 <i>queue-frame-based-accounting</i> — 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.

scheduler

Syntax	scheduler <i>scheduler-name</i> rate <i>pir-rate</i> [cir <i>cir-rate</i>] no scheduler <i>scheduler-name</i>
Context	config>subscr-mgmt>sub-prof>egr-sched-pol config>subscr-mgmt>sub-prof>ingr-sched-pol
Description	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, except for schedulers created in tier 1).
Parameters	scheduler <i>scheduler-policy-name</i> — Specify an existing scheduler policy name. <i>pir-rate</i> — 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 queues adaptation-rule parameters and the actual hardware where the queue is provisioned. Values 1 — 100000000, max Default max <i>cir-rate</i> — 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

host-tracking-policy

Syntax	host-tracking-policy <i>policy-name</i> [create] no host-tracking-policy <i>policy-name</i>
Context	config>subscr-mgmt config>subscr-mgmt>sub-prof
Description	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.

Default disabled

egress-rate-modify

Syntax **egress-rate-modify agg-rate-limit**
egress-rate-modify scheduler *scheduler-name*
no egress-rate-modify

Context config>subscr-mgmt>trk-plcy

Description This command specifies the egress-rate modification that is to be applied.
agg-rate-limit — Specifies the egress rate limit.
scheduler *scheduler-name* — Specifies the scheduler name to use.

ingress-scheduler-policy

Syntax **ingress-scheduler-policy *scheduler-policy-name***
no ingress-scheduler-policy

Context config>subscr-mgmt>sub-prof

Description This command specifies a scheduler policy to associate to the subscriber profile. Scheduler policies are configured in the **configure>qos>scheduler>policy** 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.

Parameters *scheduler-policy-name* — Specify an existing scheduler policy name.

policer-control-policy

Syntax **policer-control-policy *policy-name* [create]**
no policer-control-policy

Context config>subscr-mgmt>sub-prof>ingress
config>subscr-mgmt>sub-prof>egress

Description This command is used to create, delete, or modify policer control policies. The **policer-control-policy** 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.

Default no policer-control-policy

Parameters	<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

Syntax	max-rate {<i>kilobits-per-second</i> max} no max-rate
Context	config>subscr-mgmt>sub-prof>ingress>policer-control-policy config>subscr-mgmt>sub-prof>egress>policer-control-policy
Description	The max-rate 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. 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. 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. The policer-control-policy root max-rate setting may be overridden on each SAP or sub-profile where the policy is applied.
Default	max
Parameters	<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. Values Integer 0 – 20,000,000

max — The **max** parameter is mutually exclusive with defining a **kilobits-per-second** value. When **max** is specified, the parent policer does not enforce a maximum rate on the aggregate throughput 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

Syntax	priority-mbs-thresholds
Context	config>subscr-mgmt>sub-prof>ingress>policer-control-policy config>subscr-mgmt>sub-prof>egress>policer-control-policy
Description	The priority-mbs-thresholds command contains the root arbiter parent policer's min-thresh-separation command and each priority level's mbs-contribution 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. The priority-mbs-thresholds CLI node always exists and does not need to be created.
Default	None.

min-thresh-separation

Syntax	min-thresh-separation <i>size</i> [bytes kilobytes] no min-thresh-separation
Context	config>subscr-mgmt>sub-prof>ingress>policer-control-policy>priority-mbs-thresholds config>subscr-mgmt>sub-prof>egress>policer-control-policy>priority-mbs-thresholds
Description	The min-thresh-separation 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. 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 mbs-contribution command's optional fixed keyword is not specified): • When a parent policer instance's priority level has less than two child policers associated, the shared-portion for the level will be zero. • 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 min-thresh-separation.

The second function the system uses the **min-thresh-separation** value for is determining the value per priority level for the fair-portion:

- 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.

Default	no min-thresh-separation
Parameters	<i>size</i> [bytes kilobytes] — The <i>size</i> parameter is required when executing the min-thresh-separation command. It is expressed as an integer and specifies the shared portion in bytes or kilobytes that is selected by the trailing bytes or kilobytes keywords. If both bytes and kilobytes are missing, kilobytes is the assumed value. Setting this value has no effect on parent policer instances where the min-thresh-separation value has been overridden. Values 0 – 4194304 Default none [bytes kilobytes] — The bytes keyword is optional and is mutually exclusive with the kilobytes keyword. When specified, <i>size</i> 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, <i>size</i> is interpreted as specifying the size of min-thresh-separation in kilobytes. Values bytes or kilobytes Default kilobytes

priority

Syntax	priority <i>level</i>
Context	config>subscr-mgmt>sub-prof>ingress>policer-control-policy>priority-mbs-thresholds config>subscr-mgmt>sub-prof>egress>policer-control-policy>priority-mbs-thresholds
Description	The priority level command contains the mbs-contribution configuration command for a given strict priority level. Eight levels are supported numbered 1 through 8 with 8 being the highest strict priority. 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.
Default	None.

mbs-contribution

Syntax	mbs-contribution <i>size</i> [bytes kilobytes] [fixed] no mbs-contribution
Context	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 — 4194304

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 <i>acct-policy-name</i> [duplicate <i>acct-policy-name</i>] no radius-accounting-policy
Context	config>subscr-mgmt>sub-prof
Description	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 collect-stats command.
Parameters	<i>acct-policy-name</i> — Specifies an existing RADIUS based accounting policy. duplicate <i>acct-policy-name</i> — Specifies the RADIUS accounting policy to be used to generate duplicate accounting information.

sla-profile-map

Syntax	sla-profile-map
Context	config>subscr-mgmt>sub-prof
Description	This command enables the context to configure SLA profile mapping.

entry

Syntax	entry key <i>sub-profile-string</i> sub-profile <i>sub-profile-name</i> no entry key <i>sub-profile-string</i>
Context	config>subscr-mgmt>sub-prof>sla-prof-map
Description	This command configures SLA profile string mappings.
Parameters	<i>sub-profile-string</i> — Specifies the subscriber profile string. Values 16 characters maximum <i>sub-profile-name</i> — Specifies the subscriber profile name. Values 32 characters maximum

use-direct-map-as-default

Syntax	[no] use-direct-map-as-default
Context	config>subscr-mgmt>sub-prof>sla-prof-map
Description	This command enables direct mapping of the SLA profile as default. The no form of the command disables direct mapping.

hsmda

Syntax	hsmda
Context	config>subscr-mgmt>sub-prof
Description	This command enables the context to configure egress and ingress HSM DA queue parameters.

egress-queues

Syntax	egress-queues
Context	config>subscr-mgmt>sub-prof>hsmda
Description	This command enables the context to configure HSM DA egress queue parameters.

ingress-queues

Syntax	ingress-queues
Context	config>subscr-mgmt>sub-prof>hsmda>egress-queues
Description	This command enables the context to configure HSM DA ingress queue parameters.

agg-rate-limit

Syntax	agg-rate-limit <i>agg-rate</i> no agg-rate-limit
Context	config>subscr-mgmt>sub-prof>hsmda>egress-queues
Description	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, an attempt to bind an egress-scheduler-policy to the sub-profile will fail. 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. If a port scheduler is not defined on the egress port, the queues are allowed to operate based on their own bandwidth parameters. The no form of the command removes the aggregate rate limit from the sub-profile.
Default	no agg-rate-limit

Triple Play Subscriber Management Configuration Commands

Parameters	<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.
Values	1 — 40000000, max Kbps

qos

Syntax	qos <i>policy-id</i> no qos
Context	config>subscr-mgmt>sub-prof>hsmda>egress-queues
Description	This command assigns a SAP egress QoS policy to the HSMDA egress queue.
Parameters	<i>policy-id</i> — Specifies the policy ID of an existing QoS SAP egress policy.
Values	1 — 65535

qos

Syntax	qos <i>policy-id</i> no qos
Context	config>subscr-mgmt>sub-prof>hsmda>ingress-queues
Description	This command assigns a SAP ingress QoS policy to the HSMDA ingress queue.
Parameters	<i>policy-id</i> — Specifies the policy ID of an existing QoS SAP egress policy.
Values	1 — 65535

queue-override

Syntax	[no] queue-override
Context	config>subscr-mgmt>sub-prof>hsmda>egress-queues config>subscr-mgmt>sub-prof>hsmda>ingress-queues
Description	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.

packet-byte-offset

Syntax	packet-byte-offset { add <i>add-bytes</i> subtract <i>sub-bytes</i> } no packet-byte-offset
Context	config>subscr-mgmt>sub-prof>hsmda>egress-queues>queue-over

```
config>subscr-mgmt>sub-prof>hsmda>ingress-queues>queue-over
```

Description

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.

The accounting functions affected include:

- Offered High Priority / In-Profile Octet Counter
- Offered Low Priority / Out-of-Profile Octet Counter
- Discarded High Priority / In-Profile Octet Counter
- Discarded Low Priority / Out-of-Profile Octet Counter
- Forwarded In-Profile Octet Counter
- Forwarded Out-of-Profile Octet Counter
- Peak Information Rate (PIR) Leaky Bucket Updates
- Committed Information Rate (CIR) Leaky Bucket Updates
- Queue Group Aggregate Rate Limit Leaky Bucket Updates

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.

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. Up to 31 bytes may be added to the packet and up to 32 bytes may be removed from the packet. 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.

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.

The packet-byte-offset value may be overridden at the queue-group level.

Parameters

add *add-bytes* — Indicates that the byte value should be added to the packet for queue and queue group level accounting functions. Either the **add** or **subtract** keyword must be specified. The corresponding byte value must be specified when executing the packet-byte-offset command. The **add** keyword is mutually exclusive with the **subtract** keyword.

Values 0 — 31

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.

Values 1 — 32

queue

Syntax	queue <i>queue-id</i> [create] no queue <i>queue-id</i>
Context	config>subscr-mgmt>sub-prof>hsmda>egress-queues>queue-over config>subscr-mgmt>sub-prof>hsmda>ingress-queues>queue-over
Description	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 queue <i>queue-id</i> mapping applies to all packets, regardless of the packets destination MAC address. 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. The no form of the command returns the HSM DA queue mapping for queue to the default mapping for the forwarding class.
Parameters	<i>queue-id</i> — Specifies the queue ID to override. Values 1 — 8 create — This keyword is mandatory while creating a new queue override.

rate

Syntax	rate <i>pir-rate</i> [cir <i>cir-rate</i>] no rate
Context	config>subscr-mgmt>sub-prof>hsmda>egress-queues>queue-over config>subscr-mgmt>sub-prof>hsmda>ingress-queues>queue-over
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).

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	<i>pir-rate</i> — 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 <i>cir-rate</i> — 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. The sum keyword specifies that the CIR be used as the summed CIR values of the children schedulers or queues. Values 0 — 100000000, max, sum Default 0

slope-policy

Syntax	slope-policy <i>hsmda-slope-policy-name</i> no slope-policy
Syntax	slope-policy <i>hsmda-slope-policy-name</i> no slope-policy
Context	config>subscr-mgmt>sub-prof>hsmda>egress-queues>queue-over config>subscr-mgmt>sub-prof>hsmda>ingress-queues>queue-over
Description	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. 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-

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.

SLA Profile Commands

sla-profile

Syntax	sla-profile <i>sla-profile-name</i>
Context	config>subscr-mgmt
Description	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: • The QoS-policies to use – The classification – The queues – The queue mapping • The IP filters to use 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.
Default	none
Parameters	<i>sla-profile-name</i> — Specifies the name of the SLA profile.

egress

Syntax	egress
Context	config>subscr-mgmt>sla-profile
Description	This command configures egress parameters for the SLA profile.

ingress

Syntax	ingress
Context	config>subscr-mgmt>sla-profile
Description	This command configures ingress parameters for the SLA profile.

host-limit

Syntax	host-limit <i>max-number-of-hosts</i> [remove-oldest] no host-limit
Context	config>subscr-mgmt>sla-profile
Description	This command specifies the maximum number of hosts for this SLA profile.
Parameters	<i>max-number-of-hosts</i> — Specifies the host limit for this SLA profile. Values 1 — 100 remove-oldest — When the keywords remove-oldest are specified, the oldest subscriber host will be removed when the host limit is reached.

ip-filter

Syntax	[no] ip-filter <i>filter-id</i>
Context	config>subscr-mgmt>sla-profile>egress config>subscr-mgmt>sla-profile>ingress
Description	This command configures an egress or ingress IP filter.
Parameters	<i>filter-id</i> — Specify an existing IP filter policy ID. Values 1 — 65535

SLA Profile QoS Commands

qos

Syntax	qos <i>sap-egress-policy-id</i> [force] no qos
Context	config>subscr-mgmt>sla-prof>egress
Description	This command specifies the egress QoS policy applicable to this SLA profile. The policy must already be defined in the configure>qos>sap-egress context.
Default	1
Parameters	<i>sap-egress-policy-id</i> — Specifies the egress policy to be applied to the egress SLA profile. Values 1 — 65535 force — Forces a policy change.

qos

Syntax	qos <i>policy-id</i> [shared-queuing multipoint-shared service-queuing] [force] no qos
Context	config>subscr-mgmt>sla-prof>ingress
Description	This command specifies the ingress QoS policy applicable to this SLA profile. The policy must already be defined in the configure>qos>sap-ingress context.
Default	qos 1
Parameters	<i>sap-ingress-policy-id</i> — Specifies the policy to be applied to the ingress SLA profile. Values 1 — 65535 shared-queuing — This keyword is mutually exclusive with the multipoint-shared and service-queuing 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. multipoint-shared — This keyword is mutually exclusive with the shared-queuing and service-queuing 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-queuing — This keyword is mutually exclusive with the **multipoint-shared** and **shared-queuing** keywords to state that service queueing is needed.

force — Forces a policy change.

queue

Syntax	[no] queue <i>queue-id</i>
Context	config>subscr-mgmt>sla-prof>egress>qos config>subscr-mgmt>sla-prof>ingress>qos
Description	This command configures the context to configure egress or ingress queue parameters. Parameters defined in the config>qos>sap-egress <i>policy-id</i> or the config>qos>sap-ingress <i>policy-id</i> context are overridden by parameters specified in the subscriber management SLA profile context. 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 sla-profile SAP egress and SAP ingress policy IDs). 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. The no form of the command removes the <i>queue-id</i> from the SLA profile.
Default	none
Parameters	<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.
	Default none

avg-frame-overhead

Syntax	avg-frame-overhead <i>percent</i> no avg-frame-overhead
Context	config>subscr-mgmt>sla-prof>egress>qos>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 10000×0.1 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 50×20 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 500×1.1 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 7500×1.1 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

cbs

Syntax **cbs** *size-in-kbytes*
no cbs

Context config>subscr-mgmt>sla-prof>egress>qos>queue
config>subscr-mgmt>sla-prof>ingress>qos>queue

Description 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.

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 size as configured in the QoS policy.

Default no cbs

Parameters *size-in-kbytes* — 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).

Values 0 — 131072 or default

high-prio-only

Syntax	high-prio-only <i>percent</i> no high-prio-only
Context	config>subscr-mgmt>sla-prof>egress>qos>queue config>subscr-mgmt>sla-prof>ingress>qos>queue
Description	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.
Default	no high-prio-only
Parameters	<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. Values 0 — 100 default

mbs

Syntax	mbs <i>size-in-kbytes</i> no mbs
Context	config>subscr-mgmt>sla-prof>egress>qos>queue
Description	This command configures the 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 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. 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 to the size as configured in the QoS policy.
Default	no mbs

Parameters	<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.
Values	0 — 131072 or default

mbs

Syntax	mbs size [bytes kilobytes] no mbs
Context	config>subscr-mgmt>sla-prof>ingress>qos>queue
Description	The Maximum Burst Size (MBS) command configures the explicit definition of the maximum amount of buffers allowed for a specific queue. 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. The sap-ingress context for mbs provides a mechanism for overriding the default maximum size for the queue. 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. If the CBS value is larger than the MBS value, an error will occur, preventing the MBS change. 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 the MBS size to the size as configured in the QoS policy.
Default	no mbs
Parameters	<i>size [bytes kilobytes]</i> — 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 kilobytes parameter. A value of 0 causes the queue to discard all packets. Values 0 — 131072 or default

rate

Syntax	rate <i>pir-rate</i> [cir <i>cir-rate</i>] no rate
Context	config>subscr-mgmt>sla-prof>egress>qos>queue
Description	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. 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 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. 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 <i>queue-id</i>. The no 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 (max, 0).
Default	no rate
Parameters	<i>pir-rate</i> — 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 queues adaptation-rule parameters and the actual hardware where the queue is provisioned. Values 1 — 100000000, max Default max <i>cir-rate</i> — 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

qos-marking-from-sap

Syntax	[no] qos-marking-from-sap
Context	configure>subscr-mgmt>sla-profile>egress

Triple Play Subscriber Management Configuration Commands

Description	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. The no form of the command reverts to the egress QoS marking defined in SAP-egress policy defined at sla-profile level.
Default	qos-marking-from-sap

report-rate

Syntax	report-rate agg-rate-limit report-rate scheduler <i>scheduler-name</i> no report-rate
Context	config>subscr-mgmt>sla-prof>ingress config>subscr-mgmt>sla-prof>egress
Description	This command configures the type of rate to report for L2TP.
Parameters	<i>scheduler-name</i> — Specifies the scheduler name.

rate

Syntax	rate <i>pir-rate</i> [<i>cir cir-rate</i>] no rate
Context	config>subscr-mgmt>sla-prof>ingress>qos>queue
Description	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. 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. 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. The rate 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>. The no 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 (max, 0).
Default	no rate

Parameters	<i>pir-rate</i> — 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 queues adaptation-rule parameters and the actual hardware where the queue is provisioned. Values 1 — 100000000, max Default max <i>cir-rate</i> — Specifies the cir parameter 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

policer

Syntax	policer <i>policer-id</i> [create] no policer <i>policer-id</i>
Context	config>subscr-mgmt>sla-prof>ingress>qos config>subscr-mgmt>sla-prof>egress>qos
Description	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. 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. 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. 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. 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.

The **no** 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.

Parameters	<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—32

cbs

Syntax	cbs {<i>size</i> [bytes kilobytes] default} no cbs
Context	config>subscr-mgmt>sla-prof>ingress>qos>policer config>subscr-mgmt>sla-prof>egress>qos>policer
Description	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. The policer's cbs size defined in the QoS policy may be overridden on an sla-profile or SAP where the policy is applied. The no form of this command returns the policer to its default CBS size.
Default	none
Parameters	<i>size</i> [bytes kilobytes] — The <i>size</i> parameter is required when specifying cbs and is expressed as an integer representing the required size in either bytes or kilobytes. The default is kilobytes. The optional byte and kilobyte keywords are mutually exclusive and are used to explicitly define whether size represents bytes or kilobytes. 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

mbs

Syntax	mbs { <i>size</i> [bytes kilobytes] default } no mbs
Context	config>subscr-mgmt>sla-prof>ingress>qos>policer config>subscr-mgmt>sla-prof>egress>qos>policer
Description	This command is used to configure the policer's PIR leaky bucket's high priority violate threshold. The high-prio-only 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. 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 high-prio-only is available for the higher priority traffic. The policer's mbs size defined in the QoS policy may be overridden on an sla-profile or SAP where the policy is applied. The no form of this command returns the policer to its default MBS size.
Default	None
Parameters	<i>size</i> [bytes kilobytes] — The <i>size</i> parameter is required when specifying mbs and is expressed as an integer representing the required size in either bytes or kilobytes. The default is kilobytes. The optional byte and kilobyte keywords are mutually exclusive and are used to explicitly define whether size represents bytes or kilobytes. 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

Syntax	packet-byte-offset { add <i>bytes</i> subtract <i>bytes</i> } no packet-byte-offset
Context	config>subscr-mgmt>sla-prof>ingress>qos>policer config>subscr-mgmt>sla-prof>egress>qos>policer
Description	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 packet-byte-offset 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. When child policers are adding to or subtracting from the size of each packet, the parent policer's min-thresh-separation value should also need to be modified by the same amount. The policer's packet-byte-offset defined in the QoS policy may be overridden on an sla-profile or SAP where the policy is applied. The no version of this command is used to remove per packet size modifications from the policer.
Parameters	add <i>bytes</i> — The add keyword is mutually exclusive to the subtract keyword. Either add or subtract must be specified. When add 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. Values 1 — 32 Default None subtract <i>bytes</i> — The subtract keyword is mutually exclusive to the add keyword. Either add or subtract must be specified. When 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. Values 0—31 Default None

rate

Syntax	rate { max kilobits-per-second } [cir { max kilobits-per-second }] no rate
Context	config>subscr-mgmt>sla-prof>ingress>qos>policer config>subscr-mgmt>sla-prof>egress>qos>policer
Description	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 its exceed (CIR) or violate (PIR) threshold. The **cbs**, **mbs**, and **high-prio-only** commands are used to configure the policer's PIR and CIR thresholds.

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.

When a packet is red neither the PIR or CIR bucket depths are incremented by the packet's 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.

The policer's **adaptation-rule** command settings are used by the system to convert the specified rates into hardware timers and decrement values for the policer's buckets.

By default, the policer's metering rate is **max** and the profiling rate is 0 Kbps (all packets out-of-profile).

The **rate** settings defined for the policer in the QoS policy may be overridden on an **sla-profile** or SAP where the policy is applied.

The **no** form of this command is used to restore the default metering and profiling rate to a policer.

- Parameters** {**max** | *kilobits-per-second*} — Specifying the keyword **max** or an explicit *kilobits-per-second* 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 *kilobits-per-second* 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.
- Values** **max** or 0—20,000,000
- cir** {**max** | *kilobits-per-second*} — The optional **cir** keyword is used to override the default CIR rate of the policer. Specifying the keyword max or an explicit *kilobits-per-second* 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 *kilobits-per-second* 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.
- Values** **max** or 0—20,000,000

stat-mode

Syntax	stat-mode {no-stats minimal offered-profile-no-cir offered-priority-no-cir offered-limited-profile-cir offered-profile-cir offered-priority-cir offered-total-cir} no stat mode
Context	config>subscr-mgmt>sla-prof>ingress>qos>policer config>subscr-mgmt>sla-prof>egress>qos>policer
Description	This command is used to configure the forwarding plane counters that allow offered, output and discard accounting to occur for the policer. An ingress policer has multiple types of offered packets (explicit in-profile, explicit out-of-profile, high priority or low priority) and each of these offered types is interacting with the policer's metering and profiling functions resulting in colored output packets (green, yellow and red). Due to the large number of policers, it is not economical to allocate counters in the forwarding plane for all possible offered packet types and output conditions. Many policers will not be configured with a CIR profiling rate and not all policers will receive explicitly profiled offered packets. The stat-mode command allows provisioning of the number of counters each policer requires and how the offered packet types and output conditions should be mapped to the counters. While a no-stats mode is supported which prevents any packet accounting, the use of the policer's parent command requires at the policer's stat-mode to be set at least to the minimal setting so that offered stats are available for the policer's Fair Information Rate (FIR) to be calculated. Once a policer has been made a child to a parent policer, the stat-mode cannot be changed to no-stats unless the policer parenting is first removed. Each time the policer's stat-mode is changed, any previous counter values are lost and any new counters are set to zero. Each mode uses a certain number of counters per policer instance that are allocated from the forwarding plane's policer counter resources. You can view the the total/allocated/free stats by using the tools dump system-resources command. If insufficient counters exist to implement a mode on any policer instance, the stat-mode change will fail and the previous mode will continue unaffected for all instances of the policer. The default stat-mode when a policer is created within the policy is minimal. The stat-mode setting defined for the policer in the QoS policy may be overridden on an sla-profile or SAP where the policy is applied. If insufficient policer counter resources exist to implement the override, the stat-mode override command will fail. The previous stat-mode setting active for the policer will continue to be used by the policer. The no form of this command attempts to return the policer's stat-mode setting to minimal. The command will fail if insufficient policer counter resources exist to implement minimal where the QoS policer is currently applied and has a forwarding class mapping.
Parameters	no-stats — Counter resource allocation:0 The policer does not have any forwarding plane counters allocated and cannot provide offered, discard and forward statistics. A policer using no-stats cannot be a child to a parent policer and the policer's parent command will fail. When collect-stats is enabled, the lack of counters causes the system to generate the following statistics: a. offered-in = 0

- b. offered-out = 0
- c. discard-in = 0
- d. discard-out = 0
- e. forward-in = 0
- f. forward-out = 0

Counter 0 indicates that the accounting statistic returns a value of zero.

minimal — Counter resource allocation: 1

The default **stat-mode** for a policer is **minimal**. The **minimal** mode allocates 1 forwarding plane offered counter and one traffic manager discard counter. The forwarding counter is derived by subtracting the discard counter from the offered counter. The counters do not differentiate possible offered types (profile or priority) and do not count green or yellow output. This does not prevent the policer from supporting different offered packet types and does not prevent the policer from supporting a CIR rate.

This counter mode is useful when only the most basic accounting information is required.

The counters are used in the following manner:

- 1. 'offered = profile in/out, priority high/low
- 2. 'discarded = Same as 1
- 3. 'forwarded = Derived from 1 - 2

When **collect-stats** is enabled, the counters are used by the system to generate the following statistics:

- a. offered-in = 1
- b. offered-out = 0
- c. discard-in = 2
- d. discard-out = 0
- e. forward-in = 3
- f. 'orward-out = 0

Counter 0 indicates that the accounting statistic returns a value of zero.

With **minimal** enabled as the policer **stat-mode**, the SAP offered stats for the policer returned via MIB query and CLI show commands will return the following values:

- i. offered-in = 1
- ii. offered-out = 0
- iii. offered-undefined = 0
- iv. offered-managed = 0 (IMPM managed packets are not redirected from the policer)

Counter 0 indicates that the SAP policer statistic returns a value of zero.

offered-profile-no-cir — Counter resource allocation: 2

The **offered-profile-no-cir** mode allocates two forwarding plane offered counters and two traffic manager discard counters.

The **offered-profile-no-cir** mode is most useful when the policer is receiving only in-profile and out-of-profile pre-marked (and trusted) packets. It is expected that in this instance a CIR rate will not be defined since all packet are already pre-marked. This mode does not prevent the policer from receiving un-trusted (color undefined) nor does it prevent the policer from being configured with a CIR rate.

The counters are used in the following manner:

1. offered-in = profile in
2. offered-out = profile out, priority high/low
3. dropped-in = Same as 1
4. dropped-out = Same as 2
5. forwarded-in = Derived from 1 - 3
6. forwarded-out = Derived from 2 - 4

When **collect-stats** is enabled, the counters are used by the system to generate the following statistics:

- a. offered-in = 1
- b. offered-out = 2
- c. discard-in = 3
- d. discard-out = 4
- e. forward-in = 5
- f. forward-out = 6

With **offered-profile-no-cir** enabled as the policer **stat-mode**, the SAP offered stats for the policer returned via MIB query and CLI show commands will return the following values:

- i. offered-in = 1
- ii. offered-out = 2
- iii. offered-undefined = 0
- iv. offered-managed = 0 (IMPM managed packets are not redirected from the policer)

Counter 0 indicates that the SAP policer statistic returns a value of zero.

offered-priority-no-cir — Counter resource allocation:2

The **offered-priority-no-cir** mode allocates two forwarding plane offered counters and two traffic manager discard counters.

The **offered-priority-no-cir** mode is most useful when the policer is receiving only un-trusted packets and the ingress priority high and priority low classification options are being used without a CIR profiling rate defined. This mode does not prevent the policer from receiving trusted packets that are pre-marked in-profile or out-of-profile nor does it prevent the policer from being configured with a CIR rate.

The counters are used in the following manner:

1. offered-high = profile in, priority high
2. offered-low = profile out, priority low

- 3. dropped-high = Same as 1
- 4. dropped-low = Same as 2
- 5. forwarded-high = Derived from 1 - 3
- 6. forwarded-low = Derived from 2 - 4

When **collect-stats** is enabled, the counters are used by the system to generate the following statistics:

- a. offered-high = 1
- b. offered-low = 2
- c. discard-high = 3
- d. discard-low = 4
- e. forward-high = 5
- f. forward-low = 6

With **offered-priority-no-cir** enabled as the policer **stat-mode**, the SAP offered stats for the policer returned via MIB query and CLI show commands will return the following values:

- i. offered-high = 1
- ii. offered-low = 2
- iii. offered-undefined = 0
- iv. offered-managed = 0 (IMPM managed packets are not redirected from the policer)

Counter 0 indicates that the SAP policer statistic returns a value of zero.

offered-limited-profile-cir — Counter resource allocation:3

The **offered-limited-profile-cir** mode allocates three forwarding plane offered counters and three traffic manager discard counters.

The **offered-limited-profile-cir** mode is most useful when the policer is receiving trusted out-of-profile (profile out but no profile in) traffic and un-trusted packets are being applied to a defined CIR profiling rate. This mode does not prevent the policer from receiving trusted in-profile packets.

The counters are used in the following manner:

- 1. offered-undefined-that-turned-green = profile in, priority high/low
- 2. offered-undefined-that-turned-yellow-or-red = priority high/low
- 3. offered-out-that-stayed-yellow-or-turned-red = profile out
- 4. dropped-undefined-that-turned-green = Same as 1
- 5. dropped-undefined-that-turned-yellow-or-red = Same as 2
- 6. dropped-out-that-turned-yellow-or-red = Same as 3
- 7. forwarded-undefined-that-turned-green = Derived from 1 - 4
- 8. forwarded-undefined-that-turned-yellow = Derived from 2 - 5
- 9. forwarded-out-that-turned-yellow = Derived from 3 - 6

When **collect-stats** is enabled, the counters are used by the system to generate the following statistics:

- a. offered-in = 0
- b. offered-out = 1 + 2 + 3
- c. discard-in = 0
- d. discard-out = 4 + 5 + 6
- e. forward-in = 7
- f. forward-out = 8 + 9

With **offered-limited-profile-cir** enabled as the policer **stat-mode**, the SAP offered stats for the policer returned via MIB query and CLI show commands will return the following values:

- i. offered-in = 0
- ii. offered-out = 3
- iii. offered-undefined = 1 + 2
- iv. offered-managed = 0 (IMPM managed packets are not redirected from the policer)

Counter 0 indicates that the SAP policer statistic returns a value of zero.

offered-profile-cir — Counter resource allocation:4

The **offered-profile-cir** mode allocates four forwarding plane offered counters and four traffic manager discard counters.

The **offered-profile-cir** mode is most useful when the policer is receiving trusted out-of-profile and in-profile traffic and is also receiving un-trusted packets that are being applied to a defined CIR profiling rate. This mode differs from **offered-limited-profile-cir** mode in that it expects both trusted in-profile and out-of-profile packets while still performing CIR profiling on packets with un-trusted markings. It is expected that in most cases where both trusted and un-trusted packets are received, the predominate case will not include trusted in-profile packets making the offered-limited-profile-cir accounting mode acceptable.

The counters are used in the following manner:

- 1. offered-in-that-stayed-green-or-turned-red = profile in
- 2. offered-undefined-that-turned-green = priority high/low
- 3. offered-undefined-that-turned-yellow-or-red = priority high/low
- 4. offered-out-that-stayed-yellow-or-turned-red = profile out
- 5. dropped-in-that-stayed-green-or-turned-red = Same as 1
- 6. dropped-undefined-that-turned-green = Same as 2
- 7. dropped-undefined-that-turned-yellow-or-red = Same as 3
- 8. dropped-out-that-turned-yellow-or-red = Same as 4
- 9. forwarded-in-that-stayed-green = Derived from 1 - 5
- 10. forwarded-undefined-that-turned-green = Derived from 2 - 6
- 11. forwarded-undefined-that-turned-yellow = Derived from 3 - 7

12. forwarded-out-that-turned-yellow = Derived from 4 - 8

When **collect-stats** is enabled, the counters are used by the system to generate the following statistics:

- a. offered-in = 1
- b. offered-out = 2 + 3 + 4
- c. discard-in = 5 + 6
- d. discard-out = 7 + 8
- e. forward-in = 9 + 10
- f. forward-out = 11 + 12

With **offered-profile-cir** enabled as the policer **stat-mode**, the SAP offered stats for the policer returned via MIB query and CLI show commands will return the following values:

- i. offered-high = 1
- ii. offered-low = 4
- iii. offered-undefined = 2 + 3
- iv. offered-managed = 0 (IMPM managed packets are not redirected from the policer)

Counter 0 indicates that the SAP policer statistic returns a value of zero.

offered-priority-cir — Counter resource allocation:4

The **offered-priority-cir** mode allocates four forwarding plane offered counters and four traffic manager discard counters.

The **offered-priority-cir** mode is most useful when the policer is receiving only un-trusted packets that are being classified as high priority or low priority and are being applied to a defined CIR profiling rate. This mode differs from **offered-profile-cir** mode in that it does not expect trusted in-profile and out-of-profile packets but does not exclude the ability of the policer to receive them.

The counters are used in the following manner:

- 1. offered-high-that-turned-green = profile in, priority high
- 2. offered-high-that-turned-yellow-or-red = profile in, priority high
- 3. offered-low-that-turned-green = profile out, priority low
- 4. offered-low-that-turned-yellow-or-red = profile out, priority low
- 5. dropped-high-that-turned-green = Same as 1
- 6. dropped-high-that-turned-yellow-or-red = Same as 2
- 7. dropped-low-that-turned-green = Same as 3
- 8. dropped-low-that-turned-yellow-or-red = Same as 4
- 9. forwarded-high-that-turned-green = Derived from 1 - 5
- 10. forwarded-high-that-turned-yellow = Derived from 2 - 6
- 11. forwarded-low-that-turned-green = Derived from 3 - 7

12. forwarded-low-that-turned-yellow = Derived from 4 - 8

When **collect-stats** is enabled, the counters are used by the system to generate the following statistics:

- a. offered-high = 1 + 2
- b. offered-low = 3 + 4
- c. discard-in = 5 + 7
- d. discard-out = 6 + 8
- e. forward-in = 9 + 11
- f. forward-out = 10 + 12

With **offered-priority-cir** enabled as the policer **stat-mode**, the SAP offered stats for the policer returned via MIB query and CLI show commands will return the following values:

- i. offered-high = 1 + 2
- ii. offered-low = 3 + 4
- iii. offered-undefined = 0
- iv. offered-managed = 0 (IMPM managed packets are not redirected from the policer)

Counter 0 indicates that the SAP policer statistic returns a value of zero.

offered-total-cir — Counter resource allocation:2

The **offered-total-cir** mode allocates two forwarding plane offered counters and two traffic manager discard counters.

The **offered-total-cir** mode is most useful when the policer is not receiving trusted in-profile or out-of-profile traffic and both high and low priority classifications are not being used on the un-trusted packets and the offered packets are being applied to a defined CIR profiling rate. This mode does not prevent the policer from receiving trusted in-profile or out-of-profile packets and does not prevent the use of priority high or low classifications on the un-trusted packets.

The counters are used in the following manner:

- 1. offered-that-turned-green = profile in/out, priority high/low
- 2. offered- that-turned-yellow-or-red = profile in/out, priority high/low
- 3. dropped-offered-that-turned-green = Same as 1
- 4. dropped-offered-that-turned-yellow-or-red = Same as 2
- 5. forwarded-offered-that-turned-green = Derived from 1 - 3
- 6. forwarded-offered-that-turned-yellow = Derived from 2 - 4

When **collect-stats** is enabled, the counters are used by the system to generate the following statistics:

- a. offered-in = 1 + 2
- b. offered-out = 0
- c. discard-in = 3
- d. discard-out = 4

e. forward-in = 5

f. forward-out = 6

Counter 0 indicates that the accounting statistic returns a value of zero.

With **offered-total-cir** enabled as the policer **stat-mode**, the SAP offered stats for the policer returned via MIB query and CLI show commands will return the following values:

i. offered-high = 1 + 2

ii. offered-low = 0

iii. offered-undefined = 0

iv. offered-managed = 0 (IMPM managed packets are not redirected from the policer)

Counter 0 indicates that the SAP policer statistic returns a value of zero.

Explicit Subscriber Mapping Commands

explicit-sub-map

Syntax	explicit-sub-map
Context	config>subscr-mgmt
Description	This command configures an explicit subscriber mapping

entry

Syntax	entry key <i>sub-ident-string</i> [sub-profile <i>sub-profile-name</i>] [alias <i>sub-alias-string</i>] [sla-profile <i>sla-profile-name</i>] no entry key <i>sub-profile-string</i>
Context	config>subscr-mgmt>explicit-sub-map
Description	This command configures a subscriber identification string.
Parameters	<i>sub-ident-string</i> — Specifies the profile string. Values 16 characters maximum <i>sub-profile-name</i> — Specifies an existing subscriber profile name. Values 32 characters maximum alias <i>sub-alias-string</i> — Specifies an alias for the subscriber identification string. sla-profile <i>sla-profile-name</i> — Specifies an existing SLA profile.

Subscriber Management Service Commands

SAP Subscriber Management Commands

sub-sla-mgmt

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

host-connectivity-verify

Syntax	host-connectivity-verify [interval <i>interval</i>] [action {remove alarm}]
Context	config>service>vpls>sap config>service>ies>if>sap config>service>ies>sub-if>grp-if
Description	This command enables subscriber host connectivity verification on a given SAP within a VPLS or IES service. 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.
Default	no host-connectivity-verify
Parameters	interval <i>interval</i> — The interval, in minutes, which specifies the time interval which all known sources should be verified. The actual rate is then dependent on the number of known hosts and interval. Values 1— 6000 Note that a zero value can be used by the SNMP agent to disable host-connectivity-verify. 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 and etc.). DHCP-RELEASE will be signaled to corresponding DHCP server. Static hosts will be never removed. The alarm keyword raises an alarm indicating that the host is disconnected.

def-sla-profile

Syntax	def-sla-profile <i>default-sla-profile-name</i> no def-sla-profile
Context	config>service>vpls>sap>sub-sla-mgmt config>service>ies>if>sap>sub-sla-mgmt config>service>ies>sub-if>grp-if>sap>sub-sla-mgmt
Description	This command specifies a default SLA profile for this SAP. 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 SAP configuration.
Default	no def-sla-profile
Parameters	<i>default-sla-profile-name</i> — Specifies a default SLA profile for this SAP.

def-sub-profile

Syntax	def-sub-profile <i>default-subscriber-profile-name</i>
Context	config>service>vpls>sap>sub-sla-mgmt config>service>ies>if>sap>sub-sla-mgmt config>service>ies>sub-if>grp-if>sap>sub-sla-mgmt
Description	This command specifies a default subscriber profile for this SAP. 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	<i>default-sub-profile</i> — Specifies a default subscriber profile for this SAP.

sub-ident-policy

Syntax	sub-ident-policy <i>sub-ident-policy-name</i>
Context	config>service>vpls>sap>sub-sla-mgmt config>service>ies>if>sap>sub-sla-mgmt config>service>ies>sub-if>grp-if>sap>sub-sla-mgmt
Description	This command associates a subscriber identification policy to this SAP. 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.

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.

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.

The **no** form of the command removes the default subscriber identification policy from the SAP configuration.

Default	no sub-ident-policy
Parameters	<i>sub-ident-policy-name</i> — Specifies a subscriber identification policy for this SAP.

multi-sub-sap

Syntax	multi-sub-sap <i>number-of-sub</i> no multi-sub-sap
Context	config>service>vpls>sap>sub-sla-mgmt config>service>ies>if>sap>sub-sla-mgmt config>service>ies>sub-if>grp-if>sap>sub-sla-mgmt
Description	This command defines the maximum number of subscribers (dynamic + static) that can be simultaneously active on this SAP. 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 refers back to the default setting.
Default	no multi-sub-sap
Parameters	<i>multi-sub-sap</i> — Specifies the maximum allowed.

single-sub-parameters

Syntax	single-sub-parameters
Context	config>service>vpls>sap>sub-sla-mgmt config>service>ies>sub-if>grp-if>sap>sub-sla-mgmt config>service>ies>if>sap>sub-sla-mgmt
Description	This command configure single subscriber SAP parameters.

non-sub-traffic

Syntax	non-sub-traffic sub-profile <i>sub-profile-name</i> sla-profile <i>sla-profile-name</i> [subscriber <i>sub-ident-string</i>] no non-sub-traffic
Context	config>service>vpls>sap>sub-sla-mgmt>single-sub config>service>ies>if>sap>sub-sla-mgmt>single-sub config>service>ies>sub-if>grp-if>sap>sub-sla-mgmt>single-sub
Description	This command configures traffic profiles for non-IP traffic such as PPPoE. The no form of the command removes any configured profile.
Default	no non-sub-traffic
Parameters	<i>sub-profile-name</i> — Identifies the subscriber profile name. Values 32 characters maximum <i>sla-profile-name</i> — 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 config>service>ies>if>sap>sub-sla-mgmt>single-sub config>service>ies>sub-if>grp-if>sap>sub-sla-mgmt>single-sub
Description	This command specifies whether only profiled traffic is applicable for this SAP. When enabled, all SAP queues will be deleted. The no form of the command reverts to the default setting.
Default	no profiled-traffic-only

srrp

Syntax	[no] srrp <i>srrp-id</i>
Context	config>service>vpn>sub-if>grp-if
Description	This command creates an SRRP instance on a group IP interface. An SRRP instance manages all subscriber subnets within the group interfaces subscriber IP interface or other subscriber IP interfaces that are associated through a wholesale/retail relationship. Only one unique SRRP instance can be configured per group interface. The no form of the command removes an SRRP instance from a group IP interface. Once removed, the group interface ignores ARP requests for the SRRP gateway IP addresses that may exist on subscriber subnets associated with the group IP interface. Then the group interface stops routing using

the redundant IP interface associated with the group IP interface and will stop routing with the SRRP gateway MAC address. Ingress packets destined to the SRRP gateway MAC will also be silently discarded. This is the same behavior as a group IP interface that is disabled (shutdown).

Default	no srrp
Parameters	<i>srrp-id</i> — Specifies a 32 bit instance ID that must be unique to the system. The instance ID must also match the instance ID used by the remote router that is participating in the same SRRP context. SRRP is intended to perform a function similar to VRRP where adjacent IP hosts within local subnets use a default gateway to access IP hosts on other subnets. Values 1 — 4294967295

gw-mac

Syntax	gw-mac <i>mac-address</i> no gw-mac
Context	config>service>vprn>sub-if>grp-if>srrp
Description	This command overrides the default SRRP gateway MAC address used by the SRRP instance. Unless specified, the system uses the same base MAC address for all SRRP instances with the last octet overridden by the lower 8 bits of the SRRP instance ID. . The same SRRP gateway MAC address should be in-use by both the local and remote routers participating in the same SRRP context. One reason to change the default SRRP gateway MAC address is if two SRRP instances sharing the same broadcast domain are using the same SRRP gateway MAC. The system will use the SRRP instance ID to separate the SRRP messages (by ignoring the messages that does not match the local instance ID), but a unique SRRP gateway MAC is essential to separate the routed packets for each gateway IP address. The no form of the command removes the explicit SRRP gateway MAC address from the SRRP instance. The SRRP gateway MAC address can only be changed or removed when the SRRP instance is shutdown.
Parameters	<i>mac-address</i> — Specifies a MAC address that is used to override the default SRRP base MAC address Values Any MAC address except all zeros, broadcast or multicast addresses. The offset is expressed in normal Ethernet MAC address notation. The defined gw-mac cannot be 00:00:00:00:00:00, ff:ff:ff:ff:ff:ff or any multicast address. If not specified, the system uses the default SRRP gateway MAC address with the last octet set to the 8 least significant bits of the SRRP instance ID.

keep-alive-interval

Syntax	keep-alive-interval <i>interval</i> no keep-alive-interval
Context	config>service>vprn>sub-if>grp-if>srrp

Description	This command defines the interval between SRRP advertisement messages sent when operating in the master state. The interval is also the basis for setting the master-down timer used to determine when the master is no longer sending. The system uses three times the keep-alive interval to set the timer. Every time an SRRP advertisement is seen that is better than the local priority, the timer is reset. If the timer expires, the SRRP instance assumes that a master does not exist and initiates the attempt to become master. When in backup state, the SRRP instance takes the keep-alive interval of the master as represented in the masters SRRP advertisement message. Once in master state, the SRRP instance uses its own configured keep-alive interval. The keep-alive-interval may be changed at anytime, but will have no effect until the SRRP instance is in the master state. The no form of the command restores the default interval.				
Parameters	<i>interval</i> — Specifies the interval, in milliseconds, between SRRP advertisement messages sent when operating in the master state. <table><tr><td>Values</td><td>1 — 100</td></tr><tr><td>Default</td><td>10 milliseconds</td></tr></table>	Values	1 — 100	Default	10 milliseconds
Values	1 — 100				
Default	10 milliseconds				

message-path

Syntax	message-path <i>sap-id</i> no message-path
Context	config>service>vprn>sub-if>grp-if>srrp
Description	This command defines a specific SAP for SRRP in-band messaging. A message-path SAP must be defined prior to activating the SRRP instance. The defined SAP must exist on the SRRP instances group IP interface for the command to succeed and cannot currently be associated with any dynamic or static subscriber hosts. Once a group IP interface SAP has been defined as the transmission path for SRRP Advertisement messages, it cannot be administratively shutdown, will not support static or dynamic subscriber hosts and cannot be removed from the group IP interface. The SRRP instance message-path command may be executed at anytime on the SRRP instance. Changing the message SAP will fail if a dynamic or static subscriber host is associated with the new SAP. Once successfully changed, the SRRP instance will immediately disable anti-spoof on the SAP and start sending SRRP Advertisement messages if the SRRP instance is activated. Changing the current SRRP message SAP on an active pair of routers should be done in the following manner: 1. Shutdown the backup SRRP instance. 2. Change the message SAP on the shutdown node. 3. Change the message SAP on the active master node. 4. Re-activate the shutdown SRRP instance. Shutting down the backup SRRP instance prevents the SRRP instances from becoming master due to temporarily using differing message path SAPs. If an MCS peering is operational between the redundant nodes and the SRRP instance has been associated with the peering, the designated message path SAP will be sent from each member.

The **no** form of the command can only be executed when the SRRP instance is shutdown. Executing no message-path allows the existing SAP to be used for subscriber management functions. A new message-path SAP must be defined prior to activating the SRRP instance.

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition. See [Common Service Commands on page 1020](#) for *sap-id* command syntax.

policy

Syntax **[no] policy** *vrrp-policy-id*

Context config>service>vprn>sub-if>grp-if>srrp

Description This command associates one or more VRRP policies with the SRRP instance. A VRRP policy is a collection of connectivity and verification tests used to manipulate the in-use priorities of VRRP and SRRP instances. A VRRP policy can test the link state of ports, ping IP hosts, discover the existence of routes in the routing table or the ability to reach L2 hosts. When one or more of these tests fail, the VRRP policy has the option of decrementing or setting an explicit value for the in-use priority of an SRRP instance.

More than one VRRP policy may be associated with an SRRP instance. When more than one VRRP policy is associated with an SRRP instance the delta decrement of the in-use priority is cumulative unless one or more test fail that have explicit priority values. When one or more explicit tests fail, the lowest priority value event takes effect for the SRRP instance. When the highest delta-in-use-limit is used to manage the lowest delta derived in-use priority for the SRRP instance.

VRRP policy associations may be added and removed at anytime. A maximum of two VRRP policies can be associated with a single SRRP instance.

The **no** form of the command removes the association with *vrrp-policy-id* from the SRRP instance.

Parameters *vrrp-policy-id* — Specifies one or more VRRP policies with the SRRP instance.

Values 1 — 9999

priority

Syntax **priority** *priority*
no priority

Context config>service>vprn>sub-if>grp-if>srrp

Description This command overrides the default base priority for the SRRP instance. The SRRP instance priority is advertised by the SRRP instance to its neighbor router and is compared to the priority received from the neighbor router. The router with the best (highest) priority enters the master state while the other router enters the backup state. If the priority of each router is the same, the router with the lowest source IP address in the SRRP advertisement message assumes the master state.

The base priority of an SRRP instance can be managed by VRRP policies. A VRRP policy defines a set of connectivity or verification tests which, when they fail, may lower an SRRP instances base priority (creating an in-use priority for the instance). Every time an SRRP instances in-use priority changes when in master state, it sends an SRRP advertisement message with the new priority. If the

dynamic priority drops to zero or receives an SRRP Advertisement message with a better priority, the SRRP instance transitions to the *becoming backup* state.

When the priority command is not specified, or the no priority command is executed, the system uses a default base priority of 100. The priority command may be executed at anytime.

The **no** form of the command restores the default base priority to the SRRP instance. If a VRRP policy is associated with the SRRP instance, it will use the default base priority as the basis for any modifications to the SRRP instances in-use priority.

Parameters	<i>priority</i> — Specifies a base priority for the SRRP instance to override the default.
Values	1 — 254
Default	100

tos-marking-state

Syntax	tos-marking-state {trusted untrusted} no tos-marking-state
Context	config>service>vprn>interface config>service>vprn>sub-if>grp-if
Description	This command is used to alter 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 VPRN and network IP interface as untrusted. 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. 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. 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 save config command will not store the default tos-marking-state trusted state for network IP interfaces unless the detail parameter is also specified. The no 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.
Default	trusted
Parameters	trusted — 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. untrusted — 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.

mac-da-hashing

Syntax	mac-da-hashing no mac-da-hashing
Context	config>service>vpls>sap>sub-sla-mgmt
Description	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 no form of the command reverts to the default setting.
Default	no mac-da-hashing

host

Syntax	host {[ip <i>ip-address</i> [mac <i>mac-address</i>]]} [subscriber-sap-id subscriber <i>sub-ident-string</i> [sub-profile <i>sub-profile-name</i> [sla-profile <i>sla-profile-name</i> [ancp-string <i>ancp-string</i>] [app-profile <i>app-profile-name</i>] [inter-dest-id <i>intermediate-destination-id</i>]]] no host {[ip <i>ip-address</i>] [mac <i>ieee-address</i>]} no host all
Context	config>service>vpls>sap config>service>ies>sub-if>grp-if>sap config>service>ies>if>sap config>service>vprn>sub-if>grp-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 <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.
Default	none
Parameters	ip <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. mac <i>mac-address</i> — Specify this optional parameter when defining a static host. The MAC address must be specified for anti-spoof mac and anti-spoof ip-mac. 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.

Managed SAP Policy Commands

msap-policy

Syntax	msap-policy <i>msap-policy-name</i> [create] no msap-policy <i>msap-policy-name</i>
Context	config>subscr-mgmt
Description	This command configures a managed SAP policy. Managed SAPs allow the use of policies and a SAP template for the creation of a SAP.
Default	none
Parameters	<i>msap-policy-name</i> — Specifies the managed SAP policy name.
Values	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.
	create — Keyword used to create the managed SAP policy. The create keyword requirement can be enabled/disabled in the environment>create context.

cpu-protection

Syntax	cpu-protection
Context	config>sys>security config>service>vprn>sub-if>grp-if>sap
Description	This command enables the context to configure CPU protection policies.

cpu-protection

Syntax	cpu-protection <i>policy-id</i> [mac-monitoring] no cpu-protection
Context	config>subscr-mgmt>msap-policy [mac-monitoring] config>service>ies>sub-if>grp-if>sap [mac-monitoring] config>service>vpls>sap [mac-monitoring] config>service>vprn>sub-if>grp-if>sap [mac-monitoring]
Description	This command assigns an existing CPU protection policy to the SAP or interface. CPU protection policies are configured in the config>sys>security>cpu-protection context. The no form of the command removes the policy ID from the SAP or interface configuration.
Default	none

Triple Play Subscriber Management Configuration Commands

- Parameters** *policy-id* — Specifies an existing CPU protection policy to assign to the SAP or interface.
 mac-monitoring — Specifies that the per-source rate limit be applied.

cpu-protection

- Syntax** **cpu-protection** *policy-id*
 no cpu-protection
- Context** config>router>if
 config>service>ies>if
 config>service>vprn>if
- Description** This command assigns an existing CPU protection policy to the SAP or interface.
CPU protection policies are configured in the **config>sys>security>cpu-protection** context.
The **no** form of the command removes the policy ID from the SAP or interface configuration.
- Default** none
- Parameters** *policy-id* — Specifies an existing CPU protection policy to assign to the SAP.

ies-vprn-only-sap-parameters

- Syntax** **ies-vprn-only-sap-parameters**
- Context** config>subscr-mgmt>msap-policy
- Description** This command configures Managed SAP IES and VPRN properties.

sub-sla-mgmt

- Syntax** **[no] sub-sla-mgmt**
- Context** config>subscr-mgmt>msap-policy
 config>service>vprn>sub-if>grp-if>sap>sub-sla-mgmt
 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>ies>sub-if>grp-if>sap>sub-sla-mgmt

```
config>service>vprn>sub-if>grp-if>sap>sub-sla-mgmt
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	<i>app-profile-name</i> — specifies an existing application profile to be mapped to the subscriber profile by default.

def-inter-dest-id

Syntax	def-inter-dest-id {string <i>string</i> use-top-q} no def-inter-dest-id
Context	config>subscr-mgmt>msap-policy>sub-sla-mgmt
Description	This command specifies a default destination string for all subscribers associated with the SAP. The command also accepts the use-top-q flag that automatically derives the string based on the top most delineating Dot1Q tag from the SAP's encapsulation. The no form of the command removes the default subscriber identification string from the configuration. no def-sub-id
Default	no def-inter-dest-id
Parameters	use-top-q — Derives the string based on the top most delineating Dot1Q tag from the SAP's encapsulation. string <i>string</i> — Specifies the subscriber identification applicable for a subscriber host.

def-sub-id

Syntax	def-sub-id use-sap-id def-sub-id string <i>sub-id</i> no def-sub-id
Context	config>subscr-mgmt>msap-policy>sub-sla-mgmt config>service>ies>sub-if>grp-if>sap>sub-sla-mgmt config>service>vprn>sub-if>grp-if>sap>sub-sla-mgmt config>service>vpls>sap>sub-sla-mgmt
Description	This command specifies a default subscriber identification string. The no form of the command removes the default subscriber identification string from the configuration.
Default	no def-sub-id
Parameters	use-sap-id — Specifies to use the SAP ID as the subscriber ID.

string *sub-id* — Specifies the subscriber identification applicable for a subscriber host.

def-sla-profile

Syntax	def-sla-profile <i>default-sla-profile-name</i> 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	<i>default-sla-profile-name</i> — Specifies a default SLA profile for an MSAP.

def-sub-profile

Syntax	def-sub-profile <i>default-subscriber-profile-name</i>
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	<i>default-sub-profile</i> — Specifies a default subscriber profile for this SAP.

multi-sub-sap

Syntax	multi-sub-sap [<i>limit limit</i>] 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 refers back to the default setting.
Default	no multi-sub-sap

Parameters *multi-sub-sap* — Specifies the maximum allowed.

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.
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	<i>sub-ident-policy-name</i> — Specifies the name of the subscriber identification policy.

vpls-only-sap-parameters

Syntax	vpls-only-sap-parameters
Context	config>subscr-mgmt>msap-policy
Description	This command enables the context to configure MSAP VPLS properties.

arp-host

Syntax	arp-host
Context	config>subscr-mgmt>msap-policy>vpls-only config>service>vpls>sap>arp-host config>service>ies>sub-if>grp-if config>service>vprn>sub-if>grp-if
Description	This command enables the context to configure ARP host parameters.

host-limit

Syntax	host-limit <i>max-num-hosts</i> no host-limit
Context	config>subscr-mgmt>msap-policy>vpls-only config>service>vpls>sap>arp-host config>service>ies>sub-if>grp-if config>service>vprn>sub-if>grp-if>arp-host

Description	This command configures the maximum number of ARP hosts.
Parameters	<i>max-num-hosts</i> — Specifies the maximum number of ARP hosts.
Values	1 — 32767

min-auth-interval

Syntax	min-auth-interval <i>min-auth-interval</i> no min-auth-interval
Context	config>subscr-mgmt>msap-policy>vpls-only config>service>vpls>sap>arp-host config>service>ies>sub-if>grp-if config>service>vprn>sub-if>grp-if>arp-host
Description	This command configures the minimum authentication interval.
Parameters	<i>min-auth-interval</i> — Specifies the minimum authentication interval.
Values	1 — 6000

sap-host-limit

Syntax	sap-host-limit <i>max-num-hosts-sap</i> no sap-host-limit
Context	config>service>ies>sub-if>grp-if config>service>vprn>sub-if>grp-if>arp-host
Description	This command configures the maximum number of ARP hosts per SAP.
Parameters	<i>max-num-hosts-sap</i> — Specifies the maximum number of ARP hosts per SAP allowed on this IES interface.
Values	1 — 32767

arp-reply-agent

Syntax	arp-reply-agent [sub-ident] no arp-reply-agent
Context	config>subscr-mgmt>msap-policy>vpls-only
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 an MSAP 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 MSAP.

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 using the **host** command. Dynamic hosts are enabled on the system by enabling the **lease-populate** command in the **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 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.

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 MSAP.

Default not enabled

Parameters **sub-ident** — 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.

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 MSAP 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 MSAP'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.

dhcp

Syntax **dhcp**

Context config>subscr-mgmt>msap-policy>vpls-only

Description This command enables the context to configure DHCP parameters.

lease-populate

Syntax	lease-populate [<i>nbt-of-entries</i>] no lease-populate
Context	config>subscr-mgmt>msap-policy>vpls-only
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

option

Syntax	[no] option
Context	config>subscr-mgmt>msap-policy>vpls-only>dhcp
Description	This command enables DHCP Option 82 (Relay Agent Information Option) parameters processing and enters the context for configuring Option 82 sub-options. The no form of this command returns the system to the default.
Default	no option

action

Syntax	action {replace drop keep} no action
Context	config>subscr-mgmt>msap-policy>vpls-only>dhcp>option
Description	This command configures the Relay Agent Information Option (Option 82) processing.

Triple Play Subscriber Management Configuration Commands

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

Default The default is to keep the existing information intact.

Parameters **replace** — 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).

drop — The DHCP packet is dropped if an Option 82 field is present, and a counter is incremented.

keep — 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

Syntax **circuit-id [ascii-tuple | vlan-ascii-tuple]**
no circuit-id

Context config>subscr-mgmt>msap-policy>vpls-only>dhcp>option

Description When enabled, the router sends an ASCII-encoded tuple in the **circuit-id** 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 **circuit-id** sub-option of the DHCP packet will be left empty.

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

Default circuit-id

Parameters **ascii-tuple** — Specifies that the ASCII-encoded concatenated tuple consisting of the access-node-identifier, service-id, and interface-name is used.

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 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.

vendor-specific-option

Syntax [**no**] **vendor-specific-option**

Context config>subscr-mgmt>msap-policy>vpls-only>dhcp>option

Description This command configures the Alcatel-Lucent vendor specific sub-option of the DHCP relay packet.

client-mac-address

Syntax	[no] client-mac-address
Context	config>subscr-mgmt>msap-policy>vpls-only>dhcp>option>vendor
Description	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.

sap-id

Syntax	[no] sap-id
Context	config>subscr-mgmt>msap-policy>vpls-only>dhcp>option>vendor
Description	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

Syntax	[no] service-id
Context	config>subscr-mgmt>msap-policy>vpls-only>dhcp>option>vendor
Description	This command enables the sending of the service 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 service ID in the Alcatel-Lucent vendor specific sub-option of the DHCP relay packet.

string

Syntax	[no] string <i>text</i>
Context	config>subscr-mgmt>msap-policy>vpls-only>dhcp>option>vendor
Description	This command specifies the string in the Alcatel-Lucent vendor specific sub-option of the DHCP relay packet. The no form of the command returns the default value.
Parameters	<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

Syntax	[no] system-id
Context	config>subscr-mgmt>msap-policy>vpls-only>dhcp>option>vendor
Description	This command specifies whether the system-id is encoded in the Alcatel-Lucent vendor specific sub-option of Option 82.

proxy-server

Syntax	proxy-server
Context	config>subscr-mgmt>msap-policy>vpls-only>dhcp
Description	This command configures the DHCP proxy server.

emulated-server

Syntax	emulated-server <i>ip-address</i> no emulated-server
Context	config>subscr-mgmt>msap-policy>vpls-only>dhcp>proxy
Description	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. The no form of this command reverts to the default setting. The local proxy server will not become operational without the emulated-server address being specified.
Parameters	<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).

lease-time

Syntax	lease-time [days <i>days</i>] [hrs <i>hours</i>] [min <i>minutes</i>] [sec <i>seconds</i>] [radius-override] no lease-time
Context	config>subscr-mgmt>msap-policy>vpls-only>dhcp>proxy
Description	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 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.
Default	7 days 0 hours 0 seconds

Parameters	radius-override — Specifies that the local-proxy-server will use the configured lease-time information to provide DHCP clients.
	days — Specifies the number of days that the given IP address is valid.
	Values 0 — 3650
	hours — Specifies the number of hours that the given IP address is valid.
	Values 0 — 23
	minutes — Specifies the number of minutes that the given IP address is valid.
	Values 0 — 59
	seconds — Specifies the number of seconds that the given IP address is valid.
	Values 0 — 59

egress

Syntax	egress
Context	config>subscr-mgmt>msap-policy>vpls-only
Description	This command configures egress policies for MSAPs.

multicast-group

Syntax	multicast-group <i>group-name</i> no multicast-group
Context	config>subscr-mgmt>msap-policy>vpls-only>egress
Description	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. Only SAPs defined on Ethernet access ports are allowed into an egress-multicast-group. 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. 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.

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.

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.

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.

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.

Parameters

group-name — 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.

Default None, each egress multicast group must be explicitly configured.

Values Up to 32 egress multicast groups may be created on the system.

igmp-snooping

Syntax	igmp-snooping
Context	config>subscr-mgmt>msap-policy>vpls-only
Description	This command enables the Internet Group Management Protocol (IGMP) snooping context.
Default	none

fast-leave

Syntax	[no] fast-leave
Context	config>subscr-mgmt>msap-policy>vpls-only>igmp-snp
Description	This command enables fast leave. When IGMP fast leave processing is enabled, the SR-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'). 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.
Default	no fast-leave

import

Syntax	import <i>policy-name</i> no import
Context	config>subscr-mgmt>msap-policy>vpls-only>igmp-snp
Description	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 no form of the command removes the policy association from the SAP or SDP.
Default	no import (No import policy is specified)
Parameters	<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

Syntax	last-member-query-interval <i>tenths-of-seconds</i> no last-member-query-interval
Context	config>subscr-mgmt>msap-policy>vpls-only>igmp-snp
Description	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.
Default	10
Parameters	<i>seconds</i> — Specifies the frequency, in tenths of seconds, at which query messages are sent. Values 1 — 50

max-num-groups

Syntax	max-num-groups <i>max-num-groups</i> no max-num-groups
Context	config>subscr-mgmt>msap-policy>vpls-only>igmp-snp
Description	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.
Default	no max-num-groups
Parameters	<i>max-num-groups</i> — 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 <i>level-id</i> bw <i>bandwidth</i> no level <i>level-id</i>
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	<i>level-id</i> — Specifies has an entry for each multicast CAC policy constraint level configured on a system. Values 1 — 8 <i>bandwidth</i> — Specifies the bandwidth in kilobits per second (kbps) for the level. Values 1 — 2147483647

number-down

Syntax	number-down <i>number-lag-port-down</i> level <i>level-id</i> no number-down <i>number-lag-port-down</i>
Context	config>subscr-mgmt>msap-policy>vpls-only>igmp-snp>mcac
Description	This command configures the number of ports down along with level for multicast CAC policy on an MSAP
Parameters	<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. level <i>level-id</i> — Specifies the amount of bandwidth available within a given bundle for MC traffic for a specified level.

policy

Syntax	policy <i>policy-name</i> no policy
Context	config>subscr-mgmt>msap-policy>vpls-only>igmp-snp>mcac
Description	This command configures the multicast CAC policy name.
Parameters	<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

Syntax	unconstrained-bw <i>bandwidth</i> mandatory-bw <i>mandatory-bw</i> no unconstrained-bw
Context	config>subscr-mgmt>msap-policy>vpls-only>igmp-snp>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	<i>bandwidth</i> — The bandwidth assigned for interface's MCAC policy traffic, in kilo-bits per second (kbps). Values 0 — 2147483647 mandatory-bw <i>mandatory-bw</i> — Specifies the bandwidth pre-reserved for all the mandatory channels on a given interface in kilo-bits per second (kbps). 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. If the value of <i>mandatory-bw</i> is equal to the value of <i>bandwidth</i>, 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 <i>mandatory-bw</i> should always be less than or equal to that of <i>bandwidth</i>. An attempt to set the value of <i>mandatory-bw</i> greater than that of <i>bandwidth</i>, will result in inconsistent value error. Values 0 — 2147483647

mvr

Syntax	mvr
Context	config>subscr-mgmt>msap-policy>vpls-only>igmp-snp
Description	This command enables the context to configure Multicast VPLS Registration (MVR) parameters.

from-vpls

Syntax	from-vpls <i>service-id</i> no from-vpls
Context	config>subscr-mgmt>msap-policy>vpls-only>igmp-snp>mvr
Description	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.
Default	no from-vpls
Parameters	<i>service-id</i> — Specifies the MVR VPLS from which multicast channels should be copied into an MSAP.
Values	<i>service-id</i> : 1 — 2147483647 <i>svc-name</i> : 64 characters maximum

query-interval

Syntax	query-interval <i>seconds</i> no query-interval
Context	config>subscr-mgmt>msap-policy>vpls-only>igmp-snp
Description	This command configures the IGMP query interval. If the send-queries 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.
Default	125
Parameters	<i>seconds</i> — The time interval, in seconds, that the router transmits general host-query messages.
Values	2 — 1024

query-response-interval

Syntax	query-response-interval <i>seconds</i>
Context	config>subscr-mgmt>msap-policy>vpls-only>igmp-snp
Description	This command configures the IGMP query response interval. If the send-queries 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.
Default	10
Parameters	<i>seconds</i> — 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 <i>robust-count</i> 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	<i>robust-count</i> — 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.

Triple Play Subscriber Management Configuration Commands

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 we have a v2 querier, we will never send out a v3 group or group-source specific query when a host wants to leave a certain group.

Default no send-queries

version

Syntax	version <i>version</i> no version
Context	config>subscr-mgmt>msap-policy>vpls-only>igmp-snp
Description	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. When the send-query command 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 gets dropped and a new “wrong version” counter is incremented. If the send-query command is not configured, the version 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.
Parameters	<i>version</i> — Specify the IGMP version.
Values	1, 2, 3

mac-da-hashing

[no] mac-da-hashing

Context	config>subscr-mgmt>msap-policy>vpls-only>igmp-snp
Description	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. This command is only meaningful if subscriber management is enabled and can be configured for a VPLS service.

split-horizon-group

Syntax	split-horizon-group <i>group-name</i>
Context	config>subscr-mgmt>msap-policy>vpls-only>igmp-snp
Description	This command specifies the name of the split horizon group to which the MSAP belongs.

default-msap-policy

Syntax	default-msap-policy <i>policy-name</i> no default-msap-policy
Context	config>service>vpls>sap
Description	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. The <i>policy-name</i> parameter is only valid for a SAP with the keywords capture-sap specified in the SAP's configuration. The capture-sap 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. 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.
Default	no default-msap-policy
Parameters	<i>policy-name</i> — /Specifies an existing default managed SAP policy.

eval-msap

Syntax	eval-msap { policy <i>msap-policy-name</i> msap <i>sap-id</i> }
Context	tools>perform>subscr-mgmt
Description	This command evaluates managed SAP policies.
Parameters	policy <i>msap-policy-name</i> — Specifies an existing MSAP policy. msap <i>sap-id</i> — Specifies an MSAP sap-id.
Values	[<i>port-id</i>][<i>lag-id</i>]:<i>qtag1</i> [<i>port-id</i>][<i>lag-id</i>]:<i>qtag1.qtag2</i>

Subscriber Management Service Commands

This section contains:

- [Layer 3 Subscriber Interfaces Commands on page 928](#)
 - [Layer 3 Subscriber Interfaces SAP Commands on page 934](#)
-

Layer 3 Subscriber Interfaces Commands

subscriber-interface

Syntax	[no] subscriber-interface <i>ip-int-name</i>
Context	config>service>ies config>service>vprn
Description	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. Use the no form of the command to remove the subscriber interface.
Default	no subscriber interfaces configured
Parameters	<i>ip-int-name</i> — Specifies the interface name of a subscriber interface. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

address

Syntax	[no] address { <i>ip-address/mask</i> <i>ip-address netmask</i> }
Context	config>service>ies>subscriber-interface config>service>vprn>subscriber-interface
Description	This command creates or removes an IP address, IP subnet or broadcast address format for the interface. Multiple IP addresses can be associated with a subscriber-interface 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. In the IES subscriber interface context, this command is used to assign one or more (16 maximum) host IP addresses and subnets. This differs from a normal IES interfaces where secondary command creates an additional subnet after the primary address is assigned. A user can then add or remove addresses without having to keep a primary address. Use the no form of this command to remove the IP address assignment from the IP interface.

Default	no IP address or subnet associations configured
Parameters	<i>ip-address</i> — The IP address of the IP interface. The <i>ip-address</i> 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). <i>/</i> — The forward slash is a parameter delimiter and separates the <i>ip-address</i> portion of the IP address from the mask that defines the scope of the local subnet. No spaces are allowed between the <i>ip-address</i>, the “/” and the <i>mask-length</i> parameter. If a forward slash is not immediately following the <i>ip-address</i>, a dotted decimal mask must follow the prefix. <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. <i>netmask</i> — The subnet mask in dotted decimal notation. Values 0.0.0.0 - 255.255.255.255

authentication-policy

Syntax	authentication-policy <i>name</i> no authentication-policy
Context	config>service>vprn>if config>service>vprn>sub-if>grp-if
Description	This command assigns an authentication policy to the interface. The no form of this command removes the policy name from the group interface configuration.
Default	no authentication-policy
Parameters	<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.

arp-populate

Syntax	[no] arp-populate
Context	config>service>vprn>if config>service>vprn>sub-if>subscriber-interface config>service>vprn>sub-if>grp-if
Description	This command enables populating static and dynamic hosts into the system ARP cache. When enabled, the host’s IP address and MAC address are placed in the system ARP cache as a managed entry. Static hosts must be defined on the interface using the host command. Dynamic hosts are enabled on the system through enabling lease-populate in the IP interface DHCP context. In the event

that both a static host and a dynamic host share the same IP and MAC address, the system's ARP cache retains the host information until both the static and dynamic information are removed. Both static and dynamic hosts override static ARP entries. Static ARP entries are marked as inactive when they conflict with static or dynamic hosts and will be repopulated once all static and dynamic host information for the IP address are removed. Since static ARP entries are not possible when static subscriber hosts are defined or when DHCP lease state table population is enabled, conflict between static ARP entries and the arp-populate function is not an issue.

The **arp-populate** command will fail if an existing static subscriber host on the SAP does not have both MAC and IP addresses specified.

Once **arp-populate** is enabled, creating a static subscriber host on the SAP without both an IP address and MAC address will fail.

arp-populate can only be enabled on VPRN interfaces supporting Ethernet encapsulation.

Use the **no** form of the command to disable ARP cache population functions for static and dynamic hosts on the interface. All static and dynamic host information in the systems ARP cache will be removed. Any existing static ARP entries previously inactive due to static or dynamic hosts will be populated in the system ARP cache.

When **arp-populate** 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 arp-populate enabled.

Default not enabled

arp-timeout

Syntax **arp-timeout** *seconds*
no arp-timeout

Context config>service>vprn>interface
config>service>vprn>sub-if>grp-if

Description 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 **arp-timeout** is set to a value of zero seconds, ARP aging is disabled.

The **no** form of this command restores **arp-timeout** to the default value.

Default 14400 seconds

Parameters *seconds* — 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.

Values 0 — 65535

client-applications

Syntax	client-applications dhcp client-applications pppoe client-applications dhcp pppoe no client-applications
Context	config>service>ies>sub-if>grp-if>dhcp
Description	This command enables the clients that will try to contact the DHCP server(s). The no form of the command removes the server client type from the configuration.
Parameters	dhcp — Specifies that the DHCP relay will forward requests to the DHCP server(s). pppoe — Specifies that PPPoE will attempt to request an IP address for a PPPoE client from the DHCP server(s)ly assigned to PPPoE node.

delayed-enable

Syntax	delayed-enable seconds [init-only] no delayed-enable
Context	config>service>ies>sub-if
Description	This command delays making interface operational by the specified number of seconds. In environments with many subscribers, it can take time to synchronize the subscriber state between peers when the subscriber-interface is enabled (perhaps, after a reboot). To ensure that the state has time to be synchronized, the delayed-enable timer can be specified. The optional parameter init-only can be added to use this timer only after a reboot.
Default	no delayed-enable
Parameters	seconds — Specifies the number of seconds to delay before the interface is operational. Values 1 — 1200 init-only — Delays the initialization of the subscriber-interface to give the rest of the system time to complete necessary tasks such as allowing routing protocols to converge and/or to allow MCS to sync the subscriber information. The delay only occurs immediately after a reboot.

group-interface

Syntax	[no] group-interface ip-int-name
Context	config>service>ies>subscriber-interface config>service>vprn>subscriber-interface
Description	This command creates or deletes a group interface. This interface is designed for triple-play services where multiple SAPs are part of the same subnet. A group interface may contain one or more SAPs. Use the no form of the command to remove the group interface.

Triple Play Subscriber Management Configuration Commands

Default	no group interfaces configured
Parameters	<i>ip-int-name</i> — Specifies the interface name of a group interface. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

match-circuit-id

Syntax	[no] match-circuit-id
Context	config>service>vprn>sub-if>grp-if>dhcp
Description	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. 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. When the match-circuit-id command is enabled this as 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.
Default	no match-circuit-id

mac

Syntax	mac <i>ieee-address</i> no mac
Context	config>service>ies>subscriber-interface>group-interface
Description	This command assigns a specific MAC address to a subscriber group interface. The no form of the command returns the MAC address of the group interface to the default value.
Default	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	<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.

oper-up-while-empty

Syntax [no] **oper-up-while-empty**

Context config>service>ies>sub-if>group-interface

Description This command allows the subscriber interface to treat this group interface to be operationally enabled without any active SAPs.

This command is typically used with MSAPs where advertising the subnet prior to having a MSAP dynamically created is needed.

Layer 3 Subscriber Interfaces SAP Commands

accounting-policy

Syntax	accounting-policy <i>acct-policy-id</i> no accounting-policy
Context	config>service>ies>sub-if>grp-if>sap config>service>vprn>if>sap config>service>vprn>if>spoke-sdp config>service>vprn>sub-if>grp-if>sap
Description	This command specifies the policy to use to collect accounting statistics on a subscriber profile. A maximum of one accounting policy can be associated with a profile at one time. The no form of this command removes the accounting policy association.
Default	no accounting policy
Parameters	<i>acct-policy-id</i> — Enter the accounting <i>policy-id</i> as configured in the config>log>accounting-policy context.
Values	1 — 99

collect-stats

Syntax	[no] collect-stats
Context	config>service>ies>sub-if>grp-if>sap config>service>vprn>sub-if>grp-if>sap
Description	When enabled, the agent collects non-RADIUS accounting statistics on a subscriber profile. When the no collect-stats 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 collect-stats command is issued then the counters written to the billing file include all the traffic while the no collect-stats command was in effect.
Default	collect-stats

anti-spoof

Syntax	anti-spoof {ip ip-mac nh-mac} no anti-spoof
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

Description	This command configures the anti-spoof type of the MSAP. 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 (ip, ip-mac) defines the anti-spoof filter type enforced by the SAP when anti-spoof filtering is enabled. The no form of the command reverts back to the default. Note that for IES and VPRN subscriber group interfaces, setting no antispoof will set the default anti-spoofing type which is ip-mac.
Default	no anti-spoof
Parameters	ip — 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 ip command will fail. Note that this parameter is not applicable in the config>subscr-mgmt>msap-policy context. ip-mac — 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 ip-mac command will fail. This is also true if the default anti-spoof filter type of the SAP is ip-mac and the default is not overridden. The anti-spoof type ip-mac command will also fail if the SAP does not support Ethernet encapsulation. nh-mac — Indicates that the ingress anti-spoof is based on the source MAC and egress anti-spoof is based on the nh-ip-address.

egress

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

filter

Syntax	<pre>filter ip ip-filter-id filter no filter [ip ip-filter-id] no filter</pre>
Context	<pre>config>service>ies>sub-if>grp-if>sap>egress config>service>ies>sub-if>grp-if>sap>ingress config>service>vprn>sub-if>grp-if>sap>egress config>service>vprn>sub-if>grp-if>sap>ingress</pre>

Description	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. The filter command is used to associate a filter policy with a specified <i>ip-filter-id</i> with an ingress or egress SAP. The filter policy 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 the 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. 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 the scope command within the filter definition to change the scope to local or global. The default scope of a filter is local.
Special Cases	IES — Only IP filters are supported on an IES IP interface, and the filters only apply to routed traffic.
Parameters	ip — Keyword indicating the filter policy is an IP filter. <i>ip-filter-id</i> — Specifies the ID for the IP filter policy. Allowed values are an integer in the range of 1 and 65535 that corresponds to a previously created IP filter policy in the configure>filter context.

qos

Syntax	qos <i>policy-id</i> no qos
Context	config>service>ies>sub-if>grp-if>sap>egress config>service>vprn>sub-if>grp-if>sap>egress config>service>vprn>sub-if>grp-if>sap>ingress
Description	Associates a Quality of Service (QoS) policy with an egress Service Access Point (SAP) or IP interface. QoS 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. The qos command is used to associate egress QoS policies. The qos command only allows egress policies on SAP or IP interface egress. Attempts to associate a QoS policy of the wrong type returns an error. Only one ingress 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. By default, no specific QoS policy is associated with the SAP or IP interface for egress, so the default QoS policy is used. The normal behavior is for queues to be created per destination.

The **no** form of this command removes the QoS policy association from the SAP or IP interface, and the QoS policy reverts to the default.

policy-id — The egress policy ID to associate with SAP or IP interface on egress. The policy ID must already exist.

Values 1 — 65535

qos

Syntax	qos <i>policy-id</i> [shared-queuing] no qos
Context	config>service>ies>sub-if>grp-if>sap>ingress
Description	Associates a Quality of Service (QoS) policy with an ingress Service Access Point (SAP) or IP interface. QoS ingress 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. This qos command is used to associate ingress QoS policies. The qos command only allows ingress policies to be associated on SAP or IP interface ingress. 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. By default, no specific QoS policy is associated with the SAP or IP interface for ingress so the default QoS policy is used. The normal behavior is for queues to be created per destination. Shared and multipoint shared change this behavior creating either unicast or unicast and mcast shared queues. The no form of this command removes the QoS policy association from the SAP or IP interface, and the QoS policy reverts to the default. <i>policy-id</i> — The ingress policy ID to associate with SAP or IP interface on ingress. The policy ID must already exist. Values 1 — 65535 shared-queuing — This keyword can only be specified on SAP ingress. Specify the ingress shared queue policy used by a 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.

scheduler-policy

Syntax	scheduler-policy <i>scheduler-policy-name</i> no scheduler-policy
Context	config>service>ies>sub-if>grp-if>sap>egress config>service>ies>sub-if>grp-if>sap>ingress config>service>vprn>sub-if>grp-if>sap>egress

config>service>vprn>sub-if>grp-if>sap>ingress

Description 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 **config>qos>scheduler-policy scheduler-policy-name** 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.

host

Syntax [no] host ip ip-address [mac ieee-address] [subscriber sub-ident-string] [sub-profile sub-profile-name] [sla-profile sla-profile-name]
no host {[ip ip-address] [mac ieee-address]}
no host all

Context config>service>ies>sub-if>grp-if>sap
config>service>ies>if>sap
config>service>ies>sub-if>grp-if>sap
config>service>vprn>sub-if>grp-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 filters and ARP cache population.

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 populate features disabled. When enabled, each feature has different requirements for static hosts.

anti-spoof — When enabled, this feature uses static and dynamic host information to populate entries into an anti-spoof filter table. The anti-spoof filter entries generated will be of the same type as specified in the anti-spoof type parameter. If the SAP anti-spoof filter is defined as **ip**, each static host definition must specify an IP address. If the SAP anti-spoof filter is defined as **ip-mac**, each static host definition must specify both an IP address and MAC address. If definition of a static

host is attempted without the appropriate addresses specified for the enabled anti-spoof filter, the static host definition will fail.

arp-populate — When enabled, this feature uses static and dynamic host information to populate entries in the system ARP cache.

Attempting to define a static subscriber host that conflicts with an existing DHCP lease state table entry will fail.

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 entry and/or ARP cache entry is also removed.

Default	none
Parameters	ip <i>ip-address</i> — 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-populate. Only one static host may be configured on the SAP with a given IP address. mac <i>mac-address</i> — Specify this optional parameter when defining a static host. The MAC address must be specified for anti-spoof ip-mac and arp-populate. Multiple static hosts may be configured with the same MAC address given that each definition is distinguished by a unique IP address. subscriber <i>sub-ident-string</i> — Specify this optional parameter 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 VPRN SAP arp-reply-agent to determine the proper handling of received ARP requests from subscribers. For VPRN SAPs with arp-reply-agent enabled with the optional <i>sub-ident</i> parameter, the static subscriber hosts <i>sub-ident-string</i> 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 VPRN destinations. If the static subscriber hosts <i>sub-ident</i> 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 <i>sub-ident</i> 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 <i>sub-profile-name</i> — Specify this optional parameter 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 <i>sla-profile-name</i> — Specify this optional parameter 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.

ingress

Syntax **ingress**

Context config>service>ies>sub-if>grp-if>sap
config>service>vprn>sub-if>grp-if>sap

Description This command enables the context to configure ingress SAP Quality of Service (QoS) policies and filter policies.

If no SAP ingress 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.

multi-service-site

Syntax **multi-service-site** *customer-site-name*
no multi-service-site *customer-site-name*

Context config>service>ies>sub-if>grp-if>sap
config>service>vprn>sub-if>grp-if>sap

Description This command creates a new customer site or edits an existing customer site with the *customer-site-name* parameter. A customer site is an anchor point to create an ingress and egress virtual scheduler hierarchy. When a site is created, it must be assigned to a chassis slot or port with the exception of the 7750 SR-1 in which the slot is set to 1. When scheduler policies are defined for ingress and egress, the scheduler names contained in each policy are created according to the parameters defined in the policy. Multi-service customer sites exist for the sole purpose of creating a virtual scheduler hierarchy and making it available to queues on multiple Service Access Points (SAPs).

The scheduler policy association with the customer site normally prevents the scheduler policy from being deleted until after the scheduler policy is removed from the customer site. The multi-service-site object will generate a log message indicating that the association was deleted due to scheduler policy removal.

When the multi-service customer site is created, an ingress and egress scheduler policy association does not exist. This does not prevent the site from being assigned to a chassis slot or prevent service SAP assignment. After the site has been created, the ingress and egress scheduler policy associations can be assigned or removed at anytime.

Default None — Each customer site must be explicitly created.

Parameters *customer-site-name* — Each customer site must have a unique name within the context of the customer. If *customer-site-name* already exists for the customer ID, the CLI context changes to that site name for the purpose of editing the site scheduler policies or assignment. Any modifications made to an existing site will affect all SAPs associated with the site. Changing a scheduler policy association may cause new schedulers to be created and existing queues on the SAPs to no longer be orphaned. Existing schedulers on the site may cease to exist, causing queues relying on that scheduler to be orphaned.

If the *customer-site-name* does not exist, it is assumed that an attempt is being made to create a site of that name in the customer ID context. The success of the command execution depends on the following:

- The maximum number of customer sites defined for the chassis slot has not been met.
- The *customer-site-name* is valid.
- The **create** keyword is included in the command line syntax (if the system requires it).

When the maximum number of customer sites has been exceeded a configuration error occurs, the command will not execute and the CLI context will not change.

If the *customer-site-name* is invalid, a syntax error occurs, the command will not execute and the CLI context will not change.

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.

ATM Commands

atm

Syntax	atm
Context	config>service>ies>sub-if>grp-if>sap config>service>vprn>if>sap config>service>vprn>sub-if>grp-if>sap
Description	This command enables access to the context to configure ATM-related attributes. This command can only be used when a given context (for example, a channel or SAP) supports ATM functionality such as: • Configuring ATM port or ATM port-related functionality on MDAs supporting ATM functionality • Configuring ATM-related configuration for ATM-based SAPs that exist on MDAs supporting ATM functionality. If ATM functionality is not supported for a given context, the command returns an error.

egress

Syntax	egress
Context	config>service>ies>sub-if>grp-if>sap>atm config>service>vprn>if>sap>atm config>service>vprn>sub-if>grp-if>sap>atm
Description	This command enables the context to configure egress ATM attributes for the SAP.

encapsulation

Syntax	encapsulation atm-encap-type
Context	config>service>ies>sub-if>grp-if>sap>atm config>service>vprn>if>sap>atm config>service>vprn>sub-if>grp-if>sap>atm
Description	This command configures RFC 2684, <i>Multiprotocol Encapsulation over ATM Adaptation Layer 5</i>, encapsulation for an ATM PVCC delimited SAP. This command specifies the data encapsulation for an ATM PVCC delimited SAP. The definition references RFC 2684 and to the ATM Forum LAN Emulation specification. Ingress traffic that does not match the configured encapsulation will be dropped.
Default	The encapsulation is driven by the services for which the SAP is configured. For IES service SAPs, the default is aal5snap-routed .

Parameters	<i>atm-encap-type</i> — Specify the encapsulation type.
Values	aal5snap-routed — Routed encapsulation for LLC encapsulated circuit (LLC/ SNAP precedes protocol datagram) as defined in RFC 2684. aal5mux-ip — Routed IP encapsulation for VC multiplexed circuit as defined in RFC 2684

ingress

Syntax	ingress
Context	config>service>ies>sub-if>grp-if>sap>atm config>service>vprn>if>sap>atm config>service>vprn>sub-if>grp-if>sap>atm
Description	This command configures ingress ATM attributes for the SAP.

traffic-desc

Syntax	traffic-desc <i>traffic-desc-profile-id</i> no traffic-desc
Context	config>service>ies>sub-if>grp-if>sap>atm>egress config>service>ies>sub-if>grp-if>sap>atm>ingress config>service>vprn>if>sap>atm>egress config>service>vprn>if>sap>atm>ingress config>service>vprn>sub-if>grp-if>sap>atm>egress config>service>vprn>sub-if>grp-if>sap>atm>ingress
Description	This command assigns an ATM traffic descriptor profile to a given context (for example, a SAP). When configured under the ingress context, the specified traffic descriptor profile defines the traffic contract in the forward direction. When configured under the egress context, the specified traffic descriptor profile defines the traffic contract in the backward direction. The no form of the command reverts the traffic descriptor to the default traffic descriptor profile.
Default	The default traffic descriptor (trafficDescProfileId. = 1) is associated with newly created PVCC-delimited SAPs.
Parameters	<i>traffic-desc-profile-id</i> — Specify a defined traffic descriptor profile (see the QoS atm-td-profile command).

oam

Syntax	oam
Context	config>service>ies>sub-if>grp-if>sap>atm config>service>vprn>interface >sap>atm config>service>vprn>sub-if>grp-if>sap>atm

- Description** This command enables the context to configure OAM functionality for a PVCC delimiting a SAP. The ATM-capable MDAs support F5 end-to-end OAM functionality (AIS, RDI, Loopback):
- ITU-T Recommendation I.610 - B-ISDN Operation and Maintenance Principles and Functions version 11/95
 - GR-1248-CORE - Generic Requirements for Operations of ATM Network Elements (NEs). Issue 3 June 1996
 - GR-1113-CORE - Bellcore, Asynchronous Transfer Mode (ATM) and ATM Adaptation Layer (AAL) Protocols Generic Requirements, Issue 1, July 1994

alarm-cells

- Syntax** `[no] alarm-cells`
- Context**
`config>service>ies>sub-if>grp-if>sap>atm>oam`
`config>service>vprn>if>sap>atm>oam`
`config>service>vprn>sub-if>grp-if>sap>atm>oam`
- Description** This command configures AIS/RDI fault management on a PVCC. Fault management allows PVCC termination to monitor and report the status of their connection by propagating fault information through the network and by driving PVCCs operational status.
- When alarm-cells functionality is enabled, PVCCs operational status is affected when a PVCC goes into AIS or RDI state because of an AIS/RDI processing (i.e. assuming nothing else affects PVCCs operational status, PVCC goes DOWN, when it enters a fault state and comes back UP, when it exits that fault state) and RDI cell are generated when PVCC is operationally DOWN. No OAM-specific SNMP trap is raised whenever an endpoint enters/exits an AIS or RDI states, however, if as result of an OAM state change, the PVCC changes operational status, then a trap is expected from an entity the PVCC is associated with (for example a SAP).
- The no command disables alarm-cells functionality for a PVCC. When alarm-cells functionality is disabled, PVCCs operational status is no longer affected by PVCCs OAM state changes due to AIS/RDI processing (Note that when alarm-cells is disabled, a PVCC will change operational status to UP, if it was DOWN because of the alarm-cell processing) and RDI cells are not generated as result of PVCC going into AIS or RDI state, however, PVCCs OAM status will record OAM faults as described above.
- Default** Enabled for PVCCs delimiting IES SAPs

periodic-loopback

- Syntax** `[no] periodic-loopback`
- Context**
`config>service>ies>sub-if>grp-if>sap>atm>oam`
`config>service>vprn>if>sap>atm>oam`
`config>service>vprn>sub-if>grp-if>sap>atm`
- Description** This command enables periodic OAM loopbacks on this SAP. This command is only configurable on IES and VPRN SAPs. When enabled, an ATM OAM loopback cell is transmitted every period as configured in the `config>system>atm>oam>loopback-period period` context.

If a response is not received and consecutive retry-down retries also result in failure, the endpoint will transition to an alarm indication signal/loss of clock state. Then, an ATM OAM loopback cell will be transmitted every period as configured in the `loopback-period` *period*. If a response is received for the periodic loopback and consecutive retry-up retries also each receive a response, the endpoint will transition back to the up state.

The **no** form of the command sets the value back to the default.

Default no periodic-loopback

Redundant Interface Commands

redundant-interface

Syntax	[no] redundant-interface <i>ip-int-name</i>
Context	config>service>ies config>service>vprn
Description	This command configures a redundant interface.
Parameters	<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.

address

Syntax	address { <i>ip-address/mask</i> <i>ip-address netmask</i> } [remote-ip <i>ip-address</i>] no address
Context	config>service>vprn>redundant-interface
Description	This command assigns an IP address mask or netmask and a remote IP address to the interface.
Parameters	<i>ip-address/mask</i> — Assigns an IP address/IP subnet format to the interface. <i>ip-address 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. Assigns an IP address netmask to the interface. remote-ip <i>ip-address</i> — Assigns a remote IP to the interface.

spoke-sdp

Syntax	[no] spoke-sdp <i>sdp-id</i>
Context	config>service>vprn
Description	This command binds a service to an existing Service Distribution Point (SDP). 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. 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 an SDP with a VPRN 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 7750 SRdevices 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 <i>sdp-id</i> is bound to a service.
Special Cases	VPRN — Several SDPs can be bound to a VPRN service. Each SDP must be destined to a different 7750 SR router. If two <i>sdp-id</i> bindings terminate on the same 7750 SR, an error occurs and the second SDP binding is rejected.
Parameters	<i>sdp-id</i> — The SDP identifier. Allowed values are integers in the range of 1 and 17407 for existing SDPs. <i>vc-id</i> — The virtual circuit identifier. Values 1 — 4294967295

egress

Syntax	egress
Context	config>service>vprn>red-if>spoke-sdp
Description	This command configures egress SDP parameters.

ingress

Syntax	ingress
Context	config>service>vprn>red-if>spoke-sdp
Description	This command configures ingress SDP parameters.

vc-label

Syntax	vc-label egress-vc-label no vc-label [egress-vc-label]
Context	config>service>vprn>red-if>spoke-sdp>egress
Description	This command configures the egress VC label.
Parameters	<i>vc-label</i> — A VC egress value that indicates a specific connection. Values 16 — 1048575

vc-label

Syntax	vc-label <i>ingress-vc-label</i> no vc-label [<i>ingress-vc-label</i>]
Context	config>service>vprn>red-if>spoke-sdp>ingress
Description	This command configures the ingress VC label.
Parameters	<i>vc-label</i> — A VC ingress value that indicates a specific connection.
Values	2048 — 18431

filter

Syntax	filter { <i>ip ip-filter-id</i> } no filter
Context	config>service>vprn>red-if>spoke-sdp>ingress config>service>vprn>red-if>spoke-sdp>egress
Description	This command associates an IP filter policy with an ingress or egress Service Access Point (SAP) or IP interface. An IP filter policy can be associated with spoke SDPs. Filter policies control the forwarding and dropping of packets based on IP or MAC matching criteria. The filter command is used to associate a filter policy with a specified ip-filter-id with an ingress or egress SAP. The ip-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.
Parameters	ip <i>ip-filter-id</i> — Specifies IP filter policy. The filter ID must already exist within the created IP filters.
Values	1 — 65535

Category Map Commands

category-map

Syntax	category-map <i>category-map-name</i> [create] no category-map <i>category-map-name</i>
Context	config>subscr-mgmt
Description	This command specifies the category map name.
Default	none
Parameters	<i>category-map-name</i> — Specifies the category map name up to 32 characters in length. create — Mandatory keyword when creating a new category map.

credit-control-policy

Syntax	credit-control-policy <i>policy-name</i> [create] no credit-control-policy <i>policy-name</i>
Context	config>subscr-mgmt
Description	This command creates, configures or deletes a credit control policy.
Parameters	<i>policy-name</i> — Specifies the policy name, 32 characters max.

credit-control-server

Syntax	credit-control-server <i>radius</i> no credit-control-server
Context	config>subscr-mgmt>credit-control-policy
Description	This command configures the credit control server.
Parameters	<i>radius</i> — Specifies the radius.

default-category-map

Syntax	default-category-map <i>category-map-name</i> no default-category-map
Context	config>subscr-mgmt>credit-control-policy
Description	This command configures the default category map.

Triple Play Subscriber Management Configuration Commands

Parameters *category-map-name* — Specifies the category map name, 32 chars max.

error-handling-action

Syntax **error-handling-action** {**continue** | **block**}
 no error-handling-action

Context config>subscr-mgmt>credit-control-policy

Description This command configures the error handling action for the policy.

out-of-credit-action

Syntax **out-of-credit-action** *action*
 no out-of-credit-action

Context config>subscr-mgmt>credit-control-policy

Description This command configures the action to be performed when out of credit is reached.

Parameters *action* — Specifies the action to be taken when out of credit is reached.

Values continue|disconnect-host|re-authenticate|block-category|change-service-level

activity-threshold

Syntax **activity-threshold** *kilobits-per-second*
 no activity-threshold

Context config>subscr-mgmt>cat-map

Description 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).

Default 0

Parameters *kilobits-per-second* — Specifies the activity threshold value in kilobits per second.

Values 1 — 100000000

category

Syntax **category** *category-name* [**create**]
 no category *category-name*

Context config>subscr-mgmt>cat-map

Description This command specifies the category name.

Default	none
Parameters	<i>category-name</i> — Specifies the category name up to 32 characters in length. create — Mandatory keyword when creating a new category.

credit-type-override

Syntax	credit-type-override {volume time} no credit-type-override
Context	config>subscr-mgmt>cat-map>category
Description	This command overrides the credit-type configured in the config>subscr-mgmt>cat-map context for the given category.
Default	no credit-type-override
Parameters	volume — If different than the value specified in the credit-type command, the value overrides the credit-type. time — If different than the value specified in the credit-type command, the value overrides the credit-type.

default-credit

Syntax	default-credit volume credits bytes kilobytes megabytes gigabytes default-credit time seconds no default-credit
Context	config>subscr-mgmt>cat-map>category
Description	This command configures the default credit used by this category.
Default	no default-credit
Parameters	volume <i>kilobytes</i> — specifies the unit, in kilobytes, in which the default value for the volume credit is expressed. Values 1024 — 4294967295 time <i>seconds</i> — Specifies the default value for the time credit, in seconds. Values 60 — 4294967295

exhausted-credit-service-level

Syntax	[no] exhausted-credit-service-level
Context	config>subscr-mgmt>cat-map>category

Triple Play Subscriber Management Configuration Commands

Description	This command enables the context to configure the exhausted credit service level
Default	exhausted-credit-service-level

egress-ip-filter-entries

Syntax	[no] egress-ip-filter-entries
Context	config>subscr-mgmt>cat-map>category>exh-lvl
Description	This command configures the egress IP filter entries.

egress-ipv6-filter-entries

Syntax	[no] egress-ipv6-filter-entries
Context	config>subscr-mgmt>cat-map>category>exh-lvl
Description	This command configures the egress IPv6 filter entries.

ingress-ip-filter-entries

Syntax	[no] ingress-ip-filter-entries
Context	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 <i>entry-id</i> [create]
Context	config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ip config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ipv6 config>subscr-mgmt>cat-map>category>exh-lvl>egr-ip config>subscr-mgmt>cat-map>category>exh-lvl>egr-ipv6
Description	This command configures the IP filter entry.
Parameters	<i>entry-id</i> — Specifies the entry ID.
Values	1..65535

action

Syntax	action drop action forward action http-redirect <i>url</i> no action
Context	config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ip>entry config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ipv6>entry config>subscr-mgmt>cat-map>category>exh-lvl>egr-ip>entry config>subscr-mgmt>cat-map>category>exh-lvl>egr-ipv6>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. <i>url</i> — Specifies the HTTP URL to forward the filter entry.

match

Syntax	match [protocol <i>protocol-id</i>] no match
Context	config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ip>entry config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ipv6>entry config>subscr-mgmt>cat-map>category>exh-lvl>egr-ip>entry config>subscr-mgmt>cat-map>category>exh-lvl>egr-ipv6>entry
Description	This command configures the match criteria for this IP filter entry.
Parameters	<i>protocol-id</i> — Specifies the protocol number accepted in DHB.
Values	0..255

dscp

Syntax	dscp <i>dscp-name</i> no dscp
Context	config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ip>entry>match config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ipv6>entry>match config>subscr-mgmt>cat-map>category>exh-lvl>egr-ip>entry>match config>subscr-mgmt>cat-map>category>exh-lvl>egr-ipv6>entry>match
Description	This command configures DSCP match conditions.
Parameters	<i>dscp-name</i> — Specifies the DSCP name. Values 32 chars max

dst-ip

Syntax	dst-ip { <i>ip-address/mask</i> <i>ip-address netmask</i> } no dst-ip
Context	config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ip>entry>match config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ipv6>entry>match
Description	This command configures the destination IP match condition.
Parameters	<i>ip-address/mask</i> — Specifies the IPv4 address and mask. Values ip-address a.b.c.d mask 0..32 <i>ipv6-address/prefix-length</i> — Specifies the IPv6 address and length. Values ipv6-address x:x:x:x:x:x:x:x (where x is [0..FFFFH]) x:x:x:x:x:x:d.d.d.d (where d is [0..255]D) <i>prefix-length</i> — Specifies the prefix length. Values 1..128 <i>netmask</i> — Specifies the mask, expressed as a dotted quad. Values a.b.c.d

dst-port

Syntax	dst-port { <i>lt gt eq</i> } <i>dst-port-number</i> dst-port range <i>start end</i> no dst-port
Context	config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ip>entry>match config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ipv6>entry>match

```
config>subscr-mgmt>cat-map>category>exh-lvl>egr-ip>entry>match
config>subscr-mgmt>cat-map>category>exh-lvl>egr-ipv6>entry>match
```

Description	This command configures the destination port match condition.
Parameters	<i>lt gt eq</i> — Specifies the operator. <i>dst-port-number</i> — Specifies the destination port number as a decimal hex or binary. Values 0..65535

fragment

Syntax	fragment {true false}
Context	<pre>config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ip>entry>match config>subscr-mgmt>cat-map>category>exh-lvl>egr-ip>entry>match</pre>
Description	This command configures the fragmentation match condition.
Parameters	true false — Sets/resets fragmentation check.

icmp-code

Syntax	icmp-code <i>icmp-code</i> no icmp-code
Context	<pre>config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ip>entry>match config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ipv6>entry>match config>subscr-mgmt>cat-map>category>exh-lvl>egr-ip>entry>match config>subscr-mgmt>cat-map>category>exh-lvl>egr-ipv6>entry>match</pre>
Description	This command configures the ICMP code match condition.
Parameters	<i>icmp-code</i> — Specifies the ICMP code numbers accepted in DHB. Values 0..255

icmp-type

Syntax	icmp-type <i>icmp-type</i> no icmp-type
Context	<pre>config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ip>entry>match config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ipv6>entry>match config>subscr-mgmt>cat-map>category>exh-lvl>egr-ip>entry>match config>subscr-mgmt>cat-map>category>exh-lvl>egr-ipv6>entry>match</pre>
Description	This command configures the ICMP type match condition.

Parameters *icmp-type* — Specifies the ICMP type numbers accepted in DHB.
Values 0..255

ip-option

Syntax **ip-option** *ip-option-value* [*ip-option-mask*]
no ip-option

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 IP option match condition.

Parameters *ip-option-value* — Specifies the IP option value as a decimal hex or binary.
Values 0..255
ip-option-mask — Specifies the IP option mask as a decimal hex or binary.
Values 0..255

multiple-option

Syntax **multiple-option** {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 multiple-option match condition.

Parameters true|false — Sets or resets the multiple option check.

option-present

Syntax **option-present** {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 option-present match condition.

Parameters true | false — Sets or resets the option present check.

src-ip

Syntax	src-ip { <i>ip-address/mask</i> <i>ip-address netmask</i> }		
	no src-ip		
Context	config>subscr-mgmt>cat-map>category>exh-lvl>egr-ip>entry>match config>subscr-mgmt>cat-map>category>exh-lvl>egr-ipv6>entry>match		
Description	This command configures the source IP match condition.		
Parameters	<i>ip-address/mask</i> — Specifies the IPv4 address and mask.		
	Values	ip-address	a.b.c.d
		mask	0 — 32
	<i>netmask</i> — Specifies the mask, expressed as a dotted quad.		
	Values	a.b.c.d	
	<i>ipv6-address/prefix-length</i> — Specifies the IPv6 address and length.		
	Values	ipv6-address	x:x:x:x:x:x:x (where x is [0..FFFFH]) x:x:x:x:x:d.d.d.d (where d is [0..255]D)
	<i>prefix-length</i> — Specifies the prefix length.		
	Values	1..128	

src-port

Syntax	src-port { <i>lt gt eq</i> } <i>src-port-number</i> src-port range <i>start end</i> no src-port				
Context	config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ip>entry>match config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ipv6>entry>match config>subscr-mgmt>cat-map>category>exh-lvl>egr-ip>entry>match config>subscr-mgmt>cat-map>category>exh-lvl>egr-ipv6>entry>match				
Description	This command configures the source port match condition.				
Parameters	<i>lt gt eq</i> — Specifies the operators. <i>src-port-number</i> — Specifies the source port number as a decimal hex or binary. <table> <tr> <td>Values</td><td>0..65535</td></tr> </table> <i>dst-port-number</i> — Specifies the destination port number as a decimal hex or binary. <table> <tr> <td>Values</td><td>0..65535</td></tr> </table>	Values	0..65535	Values	0..65535
Values	0..65535				
Values	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>ingr-ipv6>entry>match config>subscr-mgmt>cat-map>category>exh-lvl>egr-ip>entry>match config>subscr-mgmt>cat-map>category>exh-lvl>egr-ipv6>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>ingr-ipv6>entry>match config>subscr-mgmt>cat-map>category>exh-lvl>egr-ip>entry>match config>subscr-mgmt>cat-map>category>exh-lvl>egr-ipv6>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 <i>pir-rate</i> 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	<i>pir-rate</i> — 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.

queue

Syntax	queue <i>queue-id</i> { ingress-only egress-only ingress-egress } no queue <i>queue-id</i>
Context	config>subscr-mgmt>cat-map>category
Description	This command configures a queue in this category.
Default	none
Parameters	ingress-only — Specifies that ingress queues are defined in this category. egress-only — Specifies that egress queues are defined in this category. ingress-egress — Specifies that ingress and egress queues are defined in this category.

rating-group

Syntax	rating-group <i>rating-group-id</i> no rating-group
Context	config>subscr-mgmt>cat-map>category
Description	This command configures the rating group applicable for this category.
Default	no rating group
Parameters	<i>rating-group-id</i> — Specifies the rating group applicable for this category.

credit-exhaust-threshold

credit-exhaust-threshold *threshold-percentage*
no credit-exhaust-threshold

Context	config>subscr-mgmt>cat-map
Description	This command specifies the credit exhaust threshold taken into account to take action as depicted in Figure 52. The no form of the command reverts the configured value to the default.
Default	100
Parameters	<i>threshold-percentage</i> — Specifies the percent to use for the credit exhaust threshold. Values 50 — 100

credit-type

Syntax	credit-type { volume time } no credit-type
Context	config>subscr-mgmt>cat-map
Description	This command specifies whether volume or time based accounting is performed.
Default	volume
Parameters	volume — specifies volume-based accounting. time — Specifies time-based accounting.

Show Commands

authentication

Syntax	authentication [<i>policy-name</i>] [detail] authentication coa-statistics authentication [<i>policy-name</i>] statistics
Context	show>subscr-mgmt
Description	This command displays subscriber authentication policy information.
Parameters	<i>policy-name</i> — Specify the subscriber management authentication policy name to display the number of authenticated and rejected packets. detail — Displays detailed information. coa-statistics — Displays the overall statistics for incoming RADIUS Change of Authorization requests. The indented lines in the output are drop reasons, used to split out the 'requests dropped' counter for clarity. statistics — Displays a list of policies with basic statistics (without specifying a policy name) or the detailed (including per-server) statistics for the given policy (specifying a policy-name).
Output	Show Subscriber Management Authentication Policy Output — The following table describes SAP ingress show command output.

Label	Description
Policy-Name	The name of the authentication policy.
Subscriber pkts Authenticated	Indicates the number of authentication requests that were successfully authenticated by this RADIUS server.
Subscriber pkts Rejected	Indicates the number of authentication requests that were rejected by this RADIUS server.

Sample Output

```
A:ALA-49>config>subscr-mgmt# show subscriber-mgmt authentication AuthPolicy
=====
Subscriber management authentication statistics
=====
Policy Name                subscriber pkts  subscriber pkts
                          Authenticated    Rejected
-----
AuthPolicy                  0                0
-----
Number of authentication policies : 1
=====
A:ALA-49>config>subscr-mgmt#
```

```
A:ALA-49>show>subscr-mgmt# show subscriber-mgmt authentication AuthPolicy detail
=====
Subscriber management authentication statistics detailed
=====
Policy name                : AuthPolicy
subscriber packets authenticated : 0
subscriber packets rejected   : 0
-----
Number of authentication policies : 1
=====
A:ALA-49>show>subscr-mgmt#

*B:gash-I# show subscriber-mgmt authentication testPlcy association
=====
Authentication Policy testPlcy
-----
SAP Associations
-----
No associations found.
-----
Interface Associations
-----
Service-Id : 1000 (IES)

- If Name : dhcp-grp-itf
=====
*B:gash-I#

*B:gash-I# show subscriber-mgmt authentication
=====
Authentication Policies
=====
Name                Description
-----
testPlcy            VPRN Radius testing authentication
policy
-----
Number of Authentication Policies : 1
=====
*B:gash-I#

*B:gash-I# show subscriber-mgmt authentication testPlcy
=====
Authentication Policy testPlcy
=====
Description          : VPRN Radius testing authentication policy
Re-authentication    : No                               Username Format      : MAC
Address
Last Mgmt Change     : 05/21/2007 13:25:31
-----
Include Radius Attributes
-----
Remote Id            : No                               Circuit Id          : No
NAS Port Id          : No                               NAS Identifier      : No
-----
Radius Servers
```

```

-----
Router           : Base           Source Address   :
10.98.0.1
Access Algorithm : Direct         Retry           : 1

Timeout          : 5
-----
Index IP Address  Port  Oper State
-----
1      10.98.0.2   1812  Unknown
=====
*B:gash-I#

*B:gash-I# show subscriber-mgmt authentication statistics
=====
Authentication Policy Statistics
=====
Policy Name                subscriber pkts  subscriber pkts
                           Authenticated    Rejected
-----
testPlcy                   0              0
-----
Number of Authentication Policies : 1
=====
*B:gash-I#

*B:gash-I# show subscriber-mgmt authentication testPlcy statistics
=====
Authentication Policy Statistics
=====
Policy name                : testPlcy
subscriber packets authenticated : 1
subscriber packets rejected   : 0
-----
radius server      requests  requests  requests  requests  requests  requests
                  accepted  rejected  no reply  md5 failed pending send failed
-----
10.98.0.2          1          0          0          0          0          0
=====
*B:gash-I#

*B:gash-I# show subscriber-mgmt authentication coa-statistics
=====
Radius Change-of-Authorization Statistics
=====
Requests Received          : 0
Requests Accepted          : 0
Requests Rejected          : 0
Requests Dropped           : 0
No authentication policy found : 0
    Invalid message        : 0
    Out of resources        : 0
    Authentication failure  : 0
=====
*B:gash-I#

```

explicit-subscriber-map

Syntax	explicit-subscriber-map
Context	show>subscriber-mgmt
Description	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#

```

local-user-db

Syntax	local-user-db <i>local-user-db-name</i> association [dhcp] [pppoe] local-user-db <i>local-user-db-name</i> dhcp-all-hosts local-user-db <i>local-user-db-name</i> dhcp-host <i>dhcp-host-name</i> local-user-db <i>local-user-db-name</i> dhcp-unmatched-hosts local-user-db [<i>local-user-db-name</i>] local-user-db <i>local-user-db-name</i> pppoe-all-hosts local-user-db <i>local-user-db-name</i> pppoe-host <i>pppoe-host-name</i> local-user-db <i>local-user-db-name</i> 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#
```

```
*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 State	Matched objects

host001	Down	-

```
Number of DHCP Hosts : 1
=====
```

```
*A:ALA-48>show>subscr-mgmt#
```

```
*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
```

```
Address               : N/A
```

```
Identification Strings
```

```
Subscriber Id         : N/A
SLA Profile String    : N/A
Sub Profile String    : N/A
```

Show Commands

```
App Profile String : N/A
ANCP String       : N/A
Inter Destination Id: N/A
=====
*A:ALA-48>show>subscr-mgmt#

*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                      vprn1000
-----
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                      vprn1000
-----
No. of Server(s): 1
=====
*A:ALA-48>show>subscr-mgmt#
```

msap-policy

Syntax	msap-policy [<i>msap-policy-name</i> [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 [<i>sla-profile-name</i> [association]]
Context	show>subscriber-mgmt
Description	This command displays SLA profile information.
Parameters	<i>sla-profile-name</i> — Specifies an existing SLA profile name. association — Displays the information configured with the specified <i>sla-profile-name</i> .

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
```

Show Commands

```
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#

```

sub-ident-policy

- Syntax** **sub-ident-policy** [*sub-ident-policy-name* [**association**]]
sub-ident-policy *sub-ident-policy-name* **script** {**primary** | **secondary** | **tertiary**}
- Context** show>subscriber-mgmt
- Description** This command displays subscriber identification policy information.
- Parameters** *sub-ident-policy-name* — Specifies an existing subscriber identification policy name.
association — Displays information configured with the specified *sub-ident-policy-name*.
script {**primary** | **secondary** | **tertiary**} — 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

Syntax	sub-profile [<i>sub-profile-name</i> [association]]
Context	show>subscriber-mgmt
Description	This command displays subscriber profile information.
Parameters	<i>sub-profile-name</i> — Specifies an existing subscriber profile name. association — 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#
```

sap-egress

Syntax **sap-egress** [*policy-id*] [**association** | **detail**]

Context show>qos

Description This command displays SAP egress policy information.

Parameters *policy-id* — Displays information for the specified SAP egress policy.
association — Displays the information configured with the specified *sap-egress* policy.
detail — Displays detailed information.

Sample Output

```
A:Dut-A# show qos sap-egress
=====
```

```
Sap Egress Policies
=====
```

Show Commands

```

Policy-Id      Scope      Description
-----
1              Template  Default SAP egress QoS policy.
30             Template  1 video channel, 1 BE
31             Template  1 video EF, 2xvideo AF, 1 BE
80             Template  Limit outgoing 80
100            Template  User100
110            Template  User110
120            Template  User120
130            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
=====
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      def      0/1
2              0          max      def      def      2/1         limit_8000
              closest  closest  def      def      0/1
3              0          max      def      def      3/1         limit_8000
              closest  closest  def      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

Syntax	sap-ingress [<i>policy-id</i>] [association match-criteria detail]
Context	show>qos
Description	This command displays SAP ingress policy information.
Parameters	<i>policy-id</i> — Displays information for the specified SAP ingress policy. association — Displays the information configured with the specified <i>sap-ingress</i> policy. match-criteria — Displays information about the matching criteria. detail — Displays detailed information.

Sample Output

```
A:Dut-A# show qos sap-ingress
```

```
Sap Ingress Policies
```

Policy-Id	Scope	Description
1	Template	Default SAP ingress QoS policy.
80	Template	Dot1p mappings/service for service 80
90	Template	Dot1p mappings/service for service 90
100	Template	Dot1p mappings/service for service 100
110	Template	Dot1p mappings/service for service 110
120	Template	Dot1p mappings/service for service 120
130	Template	Dot1p mappings/service for service 130
140	Template	Dot1p mappings/service for service 140
901	Template	User90_1
902	Template	User90_2
903	Template	User90_3
904	Template	User90_4
905	Template	User90_5
1000	Template	Dot1p mappings/service for all services

```
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

Show Commands

```

CIR Rule  PIR Rule  MBS
-----
1   Prio    0         7000   def    def    1/1      serv80
      closest closest   def    def    0/1
2   Prio    0         3500   def    def    2/1      serv80
      closest closest   def    def    0/1
3   Prio    0         2000   def    def    3/1      serv80
      closest closest   def    def    0/1
11  Prio    0         max     def    def    1/1      None
      closest closest   def    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

Syntax	scheduler-hierarchy
Context	show>qos
Description	This command enables the context to display information about policies that use this scheduler.

customer

Syntax	customer <i>customer-id</i> site <i>customer-site-name</i> [scheduler <i>scheduler-name</i>] [ingress egress] [detail]
Context	show>qos>scheduler-hierarchy show>qos>scheduler-stats
Description	This command displays the scheduler hierarchy per customer multi-service-site.
Parameters	customer <i>customer-id</i> — Displays information for the specified customer ID. site <i>customer-site-name</i> — Displays information for the specified multi-service <i>customer-site-name</i> . scheduler <i>scheduler-name</i> — 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.

sap

Syntax	sap <i>sap-id</i> [scheduler <i>scheduler-name</i>] [ingress egress] [detail]
Context	show>qos>scheduler-hierarchy show>qos>scheduler-stats
Description	This command displays the scheduler stats per SAP.
Parameters	<i>sap-id</i> — Specifies the physical port identifier portion of the SAP definition. See Common Service Commands on page 1020 for <i>sap-id</i> command syntax. scheduler <i>scheduler-name</i> — 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 <i>sub-ident-string</i> [scheduler <i>scheduler-name</i>] [ingress egress][detail]
Context	show>qos>scheduler-hierarchy show>qos>scheduler-stats
Description	This command displays the scheduler stats per subscriber.
Parameters	subscriber <i>sub-ident-string</i> — Displays the subscriber profile name. scheduler <i>scheduler-name</i> — 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.

Sample Output

```

B:Dut-A# show qos scheduler-hierarchy subscriber alcatel_100
=====
Scheduler Hierarchy - Subscriber alcatel_100
=====
Root (Ing)
| slot(1)
|--(S) : root
| |
| | |--(S) : serv_all
| | |
| | | |--(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
| | | |
| | | |--(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->1/2/1:100->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->1/2/1:101->1
| | | |
| | | |--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:101->1 1/1
| | | |
| | | |--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:101->1 3/2
| | | |
| | | |--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:101->1 1/2
| | | |
| | | |--(S) : AccessIngress:Sub=6:1 100->1/2/1:101->2
| | | |
| | | |--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:101->2 1/1
| | | |
| | | |--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:101->2 3/2
| | | |
| | | |--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:101->2 1/2
| | | |
| | | |--(S) : AccessIngress:Sub=6:1 100->1/2/1:101->3
| | | |
| | | |--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:101->3 1/1
| | | |
| | | |--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:101->3 3/2
| | | |
| | | |--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:101->3 1/2
| | | |

```

Root (Egr)

| slot(1)

|--(S) : root

| |

| |--(S) : serv_all

| |

| |--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:101->1

| |

| |--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:101->2

| |

| |--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:101->3

| |

| |--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:100->1

| |

| |--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:100->2

| |

| |--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:100->3

| |

=====

B:Dut-A#

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
|

```

Show Commands

```
| | 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

- Syntax** scheduler-name scheduler-name
- Context** show>qos
- Description** This command displays information about the specified scheduler name.
- Parameters** scheduler-name — Displays information about the specified scheduler.

scheduler-policy

- Syntax** scheduler-policy [scheduler-policy-name] [association | sap-ingress policy-id | sap-egress policy-id]
- Context** show>qos
- Description** This command displays information about the specified scheduler policy.
- Parameters**
 - scheduler-policy-name — Displays information for the specified scheduler policy.
 - sap-ingress policy-id — Displays information for the ingress policy.
 - sap-egress policy-id — Displays information for the egress policy.
 - association — Displays the information currently configured with the specified scheduler-policy-name.

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.

```

Show Commands

```
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

Syntax	scheduler-stats
Context	show>qos
Description	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

```

Show Commands

```
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
6 100 100 10 50 10 FALSE
7 10 100 3 25 10 FALSE
```

Triple Play Service Delivery Architecture

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)

Show Commands

```
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: mssl
  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 1020](#) 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:active# show subscriber-mgmt ancp-string ANCP-String-5
=====
ANCP-String "ANCP-String-5"
=====
Type       : ANCP
State      : Up
Ancp Policy: N/A

```

Show Commands

```
I. Rate      : 100 kbps          E. Rate      : 100 kbps
Adj I. Rate: N/A              Adj E. Rate: N/A
Act I. Rate: N/A              Act E. Rate: N/A
Service Id   : 999 (VPRN)
Group        : operator1
Neighbor     : xxx.xxx.xxx.xxx:54137
```

```
=====
A:active#
```

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#
```

Clear Commands

authentication

Syntax	authentication [<i>policy-name</i>] authentication coa-statistics
Context	clear
Description	This command clears subscriber authentication data.
Parameters	<i>policy-name</i> — Clears the authentication policy name. The policy must be already configured. coa-statistics — Clears statistics for incoming RADIUS Change of Authorization requests.

msap-policy

Syntax	msap-policy <i>msap-policy-name</i>
Context	clear> subscriber-mgmt
Description	This command deletes managed SAPs created by the managed SAP policy.
Parameters	<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.

radius-accounting

Syntax	radius-accounting [<i>policy-name</i>]
Context	clear> subscriber-mgmt
Description	This command clears RADIUS accounting data for the specified policy.
Parameters	<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

Syntax	scheduler-stats
Context	clear>qos
Description	This command clears scheduler statistics.

subscriber

Syntax	subscriber <i>sub-ident-string</i> [scheduler <i>scheduler-name</i>] [ingress egress]
Context	clear>qos>scheduler-stats
Description	This command clears scheduler stats per subscriber.
Parameters	<i>sub-ident-string</i> — Clears information for the subscriber profile name. scheduler <i>scheduler-name</i> — Clears information for the specified scheduler-name. egress — Clears egress information for the subscriber. ingress — Clears ingress information for the subscriber.

srrp

Syntax	srrp
Context	clear>router
Description	This command enables the context to clear and reset SRRP virtual router instances.

interface

Syntax	interface <i>subscriber-interface</i> [id <i>srrp-id</i>]
Context	clear>router>srrp
Description	This command clears and resets SRRP interface instances.
Parameters	<i>subscriber-interface</i> — Specifies an existing subscriber interface name. Values 32 chars max id <i>srrp-id</i> — Specifies an existing SRRP ID. Values 1 — 4294967295

statistics

Syntax	statistics interface <i>subscriber-interface</i> [id <i>srrp-id</i>]
Context	clear>router>srrp
Description	This command clears statistics for SRRP instances.

Parameters *subscriber-interface* — Specifies an existing subscriber interface name.

Values 32 chars max

id *srrp-id* — Specifies an existing SRRP ID.

Values 1 — 4294967295

Tools Commands

tools

Syntax	tools
Context	<root>
Description	The context to enable useful tools for debugging purposes.
Default	none
Parameters	dump — Enables dump tools for the various protocols. perform — Enables tools to perform specific tasks.

perform

Syntax	perform
Context	tools
Description	This command enables the context to enable tools to perform specific tasks.
Default	none

subscriber-mgmt

Syntax	subscriber-mgmt
Context	tools>perform
Description	This command enables tools to control subscriber management.

edit-lease-state

Syntax	edit-lease-state sap <i>sap-id</i> ip <i>ip-address</i> [subscriber <i>sub-ident-string</i>] [sub-profile-string <i>sub-profile-string</i>] [sla-profile-string <i>sla-profile-string</i>] edit-lease-state svc-id <i>service-id</i> ip <i>ip-address</i> [subscriber <i>sub-ident-string</i>] [sub-profile-string <i>sub-profile-string</i>] [sla-profile-string <i>sla-profile-string</i>]
Context	tools>perform>subscr-mgmt
Parameters	sap <i>sap-id</i> — Specifies the physical port identifier portion of the SAP definition. See Common Service Commands on page 1020 for <i>sap-id</i> command syntax. ip <i>ip-address</i> — Modifies lease state information for the specified IP address.

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.

Values	<i>service-id:</i>	1 — 2147483647
	<i>svc-name:</i>	64 characters maximum

credit-reset

Syntax **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}

Context tools>perform>subscr-mgmt

Description This command resets the credit for an SLA-profile instance.

Parameters **sap *sap-id*** — Specifies the physical port identifier portion of the SAP definition. See [Common Service Commands on page 1020](#) for *sap-id* command syntax.

ip *ip-address* — Modifies lease state information for the specified IP address.

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.

Values	<i>service-id:</i>	1 — 2147483647
	<i>svc-name:</i>	64 characters maximum

eval-lease-state

Syntax **eval-lease-state [svc-id *service-id*] [sap *sap-id*] [subscriber *sub-ident-string*] [ip *ip-address*]**

Context tools>perform>subscr-mgmt

Description This command evaluates lease state information.

Parameters	svc-id <i>service-id</i> — Evaluates lease state information for the specified service. Values <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> sap <i>sap-id</i> — Evaluates lease state information for the specified SAP. <i>sap-id</i> — Specifies the physical port identifier portion of the SAP definition. See Common Service Commands on page 1020 for <i>sap-id</i> command syntax. subscriber <i>sub-ident-string</i> — Evaluates lease state information for the specified subscriber identification string. ip <i>ip-address</i> — Evaluates lease state information for the specified IP address.	<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				

re-ident-sub

Syntax	re-ident-sub <i>old-sub-ident-string</i> to <i>new-sub-ident-string</i>
Context	tools>perform>subscr-mgmt
Description	This command renames a subscriber identification string.
Parameters	<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

Syntax	redundancy
Context	tools>dump
Description	This command enables the context to dump redundancy parameters.

multi-chassis

Syntax	multi-chassis
Context	tools>dump>redundancy
Description	This command enables the context to dump multi-chassis parameters.

mc-ring

Syntax	mc-ring
Context	tools>dump>redundancy>multi-chassis

Description This command dumps multi-chassis ring data.

sync-database

Syntax **sync-database** [**peer** *ip-address*] [**port** *port-id* | *lag-id*] [**sync-tag** *sync-tag*] [**application** *application*] [**detail**] [**type** *type*]

Context tools>dump>redundancy>multi-chassis

Description This command dumps multi-chassis sync database information.

Parameters

- peer** *ip-address* — Dumps the specified address of the multi-chassis peer.
- port** *port-id* — Dumps the specified port ID of the multi-chassis peer.
- port** *lag-id* — Dumps the specified Link Aggregation Group (LAG) on this system.
- sync-tag** *sync-tag* — Dumps the synchronization tag used while synchronizing this port with the multi-chassis peer.
- application** — Dumps the specified application information that was synchronized with the multi-chassis peer.
- Values** dhcps, igmp, igmp-snooping, mc-ring, srrp, sub-mgmt, mld-snooping, all
- detail** — Displays detailed information.
- type** *type* — Filters by the specified entry type.
- Values** alarm-deleted, local-deleted

srrp-sync-data

Syntax **srrp-sync-database** [**instance** *instance-id*] [**peer** *ip-address*]

Context tools>dump>redundancy>multi-chassis

Description This command dumps multi-chassis SRRP sync database information.

Parameters

- instance-id* — [1..4294967295]
- ip-address* — Dumps the specified address (in the form of a.b.c.d).

Debug Commands

arp

Syntax	arp { all <i>ip-address</i> } arp interface [<i>ip-int-name</i> <i>ip-address</i>]
Context	clear>router
Description	This command clears all or specific ARP entries. The scope of ARP cache entries cleared depends on the command line option(s) specified.
Parameters	all — Clears all ARP cache entries. <i>ip-addr</i> — Clears the ARP cache entry for the specified IP address. interface <i>ip-int-name</i> — Clears all ARP cache entries for the interface with the specified name. interface <i>ip-addr</i> — Clears all ARP cache entries for the specified interface with the specified address.

arp-host

Syntax	[no] arp-host
Context	debug>service>id
Description	This command enables and configures ARP host debugging. The no form of the command disables ARP host debugging.

authentication

Syntax	authentication [policy <i>policy-name</i>] [mac-addr <i>ieee-address</i>] [circuit-id <i>circuit-id</i>]
Context	debug>subscr-mgmt
Description	This command debugs subscriber authentication.
Parameters	policy <i>policy-name</i> — Specify an existing subscriber management authentication policy name. mac-addr <i>ieee-address</i> — Specifies the 48-bit MAC address xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx. circuit-id <i>circuit-id</i> — Specify the circuit-id, up to 256 characters.

sub-ident-policy

Syntax	[no] sub-ident-policy <i>policy-name</i>
Context	debug>subscr-mgmt
Description	This command debugs subscriber identification policies.
Parameters	<i>policy-name</i> — Specifies the subscriber identification policy to debug.

script-compile-error

Syntax	[no] script-compile-error
Context	debug>subscr-mgmt>sub-ident-plcy
Description	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

Syntax	[no] script-export-variables
Context	debug>subscr-mgmt>sub-ident-plcy
Description	This command sends the result (the three output variables) of the Python script to the logger when the script ran successfully.

script-output

Syntax	[no] script-output
Context	debug>subscr-mgmt>sub-ident-plcy
Description	This command sends the output (such as from 'print' statements) of the Python script to the logger.

script-output-on-error

Syntax	[no] script-output-on-error
Context	debug>subscr-mgmt>sub-ident-plcy
Description	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

Syntax	[no] script-runtime-error
Context	debug>subscr-mgmt>sub-ident-plcy
Description	This command sends the traceback of the Python script failure to the logger.

script-all-info

Syntax	script-all-info
Context	debug>subscr-mgmt>sub-ident-plcy
Description	This command enables the script-compile-error, script-export-variables, script-output, script-output-on-error, and script-runtime-error functionalities.

srrp

Syntax	[no] srrp
Context	debug>router
Description	This command enables debugging for SRRP packets. The no form of the command disables debugging.

events

Syntax	[no] events [interface <i>ip-int-name</i>]
Context	debug>router>srrp
Description	This command enables debugging for SRRP packets. The no form of the command disables debugging.

packets

Syntax	[no] packets [interface <i>ip-int-name</i>]
Context	debug>router>srrp
Description	This command enables debugging for SRRP packets. The no form of the command disables debugging.

radius

Syntax	radius [detail] [hex] no radius
Context	debug
Description	This command enables debugging for RADIUS connections.
Parameters	detail — Displays detailed output. hex — Displays the packet dump in hex format.

Monitor Commands

subscriber

Syntax	subscriber <i>sub-ident-string</i> sap <i>sap-id</i> sla-profile <i>sla-profile-name</i> [base ingress-queue-id <i>ingress-queue-id</i> egress-queue-id <i>egress-queue-id</i>] [interval <i>seconds</i>] [repeat <i>repeat</i>] [absolute rate]
Context	monitor>service
Description	This command monitors statistics for a subscriber.
Parameters	sub-ident-string — Specifies an existing subscriber identification profile to monitor. sap <i>sap-id</i> — Specifies the physical port identifier portion of the SAP definition. See Common Service Commands on page 1020 for <i>sap-id</i> command syntax. sla-profile <i>sla-profile-name</i> — Specifies an existing SLA profile. interval <i>seconds</i> — Configures the interval for each display in seconds. Default 11 Values 11 — 60 repeat <i>repeat</i> — Configures how many times the command is repeated. Default 10 Values 1 — 999 absolute — When the absolute keyword is specified, the raw statistics are displayed, without processing. No calculations are performed on the delta or rate statistics. Default mode delta rate — When the rate keyword is specified, the rate-per-second for each statistic is displayed instead of the delta. base — Monitor base statistics. ingress-queue-id <i>ingress-queue-id</i> — Monitors statistics for this queue. Values 1 — 32 egress-queue-id <i>egress-queue-id</i> — Monitors statistics for this queue. Values 1 — 8

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
-----

```

Monitor Commands

```

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

Syntax	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] host [sap <i>sap-id</i>] [wholesaler <i>service-id</i>] [port <i>port-id</i>] no-inter-dest-id [detail] host summary host [detail] wholesaler <i>service-id</i> (VPRN only)
Context	show>service>id
Description	This command displays static host information configured on this service.
Parameters	sap <i>sap-id</i> — Displays SAP information for the specified SAP ID. See Common Service Commands on page 1020 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.

wholesaler *service-id* — The VPRN service ID of the wholesaler. When specified in this context, SAP, SDP, interface, IP address and MAC parameters are ignored.

Values	<i>service-id:</i>	1 — 2147483647
	<i>svc-name:</i>	64 characters maximum

Service Manager Scripting Language

In This Chapter

This section describes the service manager scripting language and provides example scripts.

The following topics are included:

- [Service Manager Scripting Language on page 1008](#)
 - [Python Changes on page 1009](#)
 - [Configuration on page 1009](#)
 - [Operator Debugging on page 1011](#)
 - [Python Scripts on page 1012](#)
 - [Limitations on page 1013](#)
 - [Tips and Tricks on page 1014](#)
 - [Sample Python Scripts on page 1015](#)

Service Manager Scripting Language

The service manager scripting language 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` — Access to DHCP packet and export variables back to service manager.
- `binascii` — Common ASCII decoding like `base64`.
- `re` — Regular expression support.
- `struct` — Parse and manipulate binary strings.

The `alc` package contains only one object, `DHCP`, and has the following members:

Table 10: DHCP Object Members

Name	Read	Write	Class
<code>htype</code>	X		integer
<code>hlen</code>	X		integer
<code>hops</code>	X		integer
<code>flags</code>	X		integer
<code>ciaddr</code>	X		integer
<code>yiaddr</code>	X		integer
<code>siaddr</code>	X		integer
<code>giaddr</code>	X		integer
<code>chaddr</code>	X		string
<code>sname</code>	X		string
<code>file</code>	X		string
<code>options</code>	X		TLV
<code>sub_ident</code>		X	string
<code>sub_profile_string</code>		X	string
<code>sla_profile_string</code>		X	string
<code>anccp_string</code>		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 example [us5.py on page 1012](#), 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.

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.
-

Configuration

As an example consider script [us5.py on page 1012](#) which sets the `sub_ident` variable based on the circuit ID of three different DSLAMs:

```
import re
import alc
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(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)
```

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/
parsefaill.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: parsefaill.py
  File "ftp://xxx.xxx.xx.xx/py/parsefaill.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\rdut-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 flavours:
#      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)
```

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.

Tips and Tricks

- Use xrange() instead of range().
- Avoid doing a lot of string operations. 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)
```

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. 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\0x0Bmysdl-video
```

The following variables are returned:

```
sub_ident: mysdl-video
sub_profile_string: mysdl
sla_profile_string: video
```


Example

This script parses the option 82 “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 flavours:
5.  #      FENT      string      "TLV1: ATM:3/0:100.33"
6.  #      GELT      octet-stream 0x01010000A0A0A0A00000000640022
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
```

Common CLI Command Descriptions

In This Chapter

This section provides information about common Command Line Interface (CLI) syntax and command usage.

Topics in this chapter include:

- [sap on page 1020](#)
- [port on page 1024](#)

Common Service Commands

sap

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

Description This command specifies the physical port identifier portion of the SAP definition.

Parameters *sap-id* — The *sap-id* can be configured in one of the following formats:

Type	Syntax	Example
port-id	<i>slot/mda/port[.channel]</i>	2/1/11 1/2/3.1
null	<i>[port-id bundle-id bpgrp-id lag-id aps-id]</i>	<i>port-id:</i> 1/1/3 <i>bundle-id:</i> bundle-ppp-1/1.1 <i>bpgrp-id:</i> bpgrp-ima-1 <i>lag-id:</i> lag-63 <i>aps-id:</i> aps-1
dot1q	<i>[port-id bundle-id bpgrp-id lag-id aps-id]:qtag1</i>	<i>port-id:qtag1:</i> 1/1/3:100 <i>bundle-id:</i> bundle-ppp-1/1.1 <i>bpgrp-id:</i> bpgrp-ima-1 <i>lag-id:qtag1:</i> lag-63:102 <i>aps-id:qtag1:</i> aps-1:27
qinq	<i>[port-id bundle-id bpgrp-id lag-id]:qtag1.qtag2</i>	<i>port-id:qtag1.qtag2:</i> 1/1/3:100.10 <i>bundle-id:</i> bundle-ppp-1/1.1 <i>bpgrp-id:</i> bpgrp-ima-1 <i>lag-id:qtag1.qtag2:</i> lag-10:
atm	<i>[port-id aps-id bundle-id bpgrp-id][:vpi/vci vpi vpi1.vpi2]</i>	<i>port-id:</i> 1/1/1 <i>aps-id:</i> aps-1 <i>bundle-id:</i> bundle-ima-1/1.1 bundle-ppp-1/1.1 <i>bpgrp-id:</i> bpgrp-ima-1 <i>vpi/vci:</i> 16/26 <i>vpi:</i> 16 <i>vpi1.vpi2:</i> 16.200
frame-relay	<i>[port-id aps-id]:dlci</i>	<i>port-id:</i> 1/1/1:100 <i>aps-id:</i> aps-1 <i>dlci:</i> 16
cisco-hdlc	<i>slot/mda/port.channel</i>	<i>port-id:</i> 1/1/3.1
Values: <i>sap-id:</i>		
	null	<i>[port-id bundle-id bpgrp-id lag-id aps-id]</i>
	dot1q	<i>[port-id bundle-id bpgrp-id lag-id aps-id]:qtag1</i>
	qinq	<i>[port-id bundle-id bpgrp-id lag-id]:qtag1.qtag2</i>
	atm	<i>[port-id aps-id][:vpi/vci vpi vpi1.vpi2]</i>

frame	[<i>port-id</i> <i>aps-id</i>]: <i>dlci</i>
cisco-hdlc	<i>slot/mda/port.channel</i>
cem	<i>slot/mda/port.channel</i>
ima-grp	[<i>bundle-id</i>]: <i>vpi/vci vpi vpi1.vpi2</i>
port-id	<i>slot/mda/port[.channel]</i>
bundle-id	<i>bundle-type-slot/mda.bundle-num</i>
	<i>bundle</i> keyword
	<i>type</i> ima, fr, ppp
	<i>bundle-num</i> 1 — 256
bpgrp-id	<i>bpgrp-type-bpgrp-num</i>
	<i>bpgrp</i> keyword
	<i>type</i> ima, ppp
	<i>bpgrp-num</i> 1 — 1280
aps-id	<i>aps-group-id[.channel]</i>
	<i>aps</i> keyword
	<i>group-id</i> 1 — 64
ccag-id	<i>ccag-id.path-id[cc-type]:cc-id</i>
	<i>ccag</i> keyword
	<i>id</i> 1 — 8
	<i>path-id</i> a, b
	<i>cc-type</i> .sap-net, .net-sap
	<i>cc-id</i> 0 — 4094
eth-tunnel	<i>eth-tunnel-tunnel-id</i>
	<i>tunnel-id</i> 1 — 1024
	<i>eth-tunnel-sap-id</i> 0 — 4094
lag-id	<i>lag-id</i>
	<i>lag</i> keyword
	<i>id</i> 1 — 200
qtag1	0 — 4094
qtag2	*, 0 — 4094
vpi	NNI: 0 — 4095
	UNI: 0 — 255
vci	1, 2, 5 — 65535
dlci	16 — 1022
ipsec-id	<i>ipsec-id.private</i> <i>public:tag</i>
	<i>ipsec</i> keyword
	<i>id</i> 1 — 4
	<i>tag</i> 0 — 4094

bundle-id — Specifies the multilink bundle to be associated with this IP interface. The **bundle** keyword must be entered at the beginning of the parameter.

The command syntax must be configured as follows:

bundle-id: **bundle-type-slot-id/mda-slot.bundle-num**
bundle-id value range: 1 — 256

For example:

```
*A:ALA-12>config# port bundle-ppp-5/1.1
*A:ALA-12>config>port# multilink-bundle
```

bpgrp-id — Specifies the bundle protection group ID to be associated with this IP interface. The **bpgrp** keyword must be entered at the beginning of the parameter.

The command syntax must be configured as follows:

```
bpgrp-id:                bpgrp-type-bpgrp-num
type:                    ima
bpgrp-num value range:  1 — 1280
```

For example:

```
*A:ALA-12>config# port bpgrp-ima-1
*A:ALA-12>config>service>vpls$ sap bpgrp-ima-1
```

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 TDM	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 TDM	BCP-Dot1q	0 — 4094	The SAP is identified by the 802.1Q tag on the channel.
SONET/SDH TDM	Frame Relay	16 — 991	The SAP is identified by the data link connection identifier (DLCI).
SONET/SDH ATM	ATM	vpi (NNI) 0 — 4095 vpi (UNI) 0 — 255 vci 1, 2, 5 — 65535 -	The SAP is identified by port or by PVPC or PVCC identifier (vpi, vpi/vci, or vpi range)

sap ipsec-id.private|public:tag — This parameter associates an IPsec group SAP with this interface. This is the public side for an IPsec tunnel. Tunnels referencing this IPsec group in the private side may be created if their local IP is in the subnet of the interface subnet and the routing context specified matches with the one of the interface.

This context will provide a SAP to the tunnel. The operator may associate an ingress and egress QoS policies as well as filters and virtual scheduling contexts. Internally this creates an Ethernet SAP that will be used to send and receive encrypted traffic to and from the MDA. Multiple tunnels can be associated with this SAP. The “tag” will be a dot1q value. The operator may see it as an identifier. The range is limited to 1 — 4095.

port

Syntax	port <i>port-id</i>																																														
Description	This command specifies a port identifier.																																														
Parameters	<i>port-id</i> — The <i>port-id</i> can be configured in one of the following formats:																																														
Values	<table><tr><td>port-id</td><td>slot/mda/port[.channel]</td></tr><tr><td>bundle-id</td><td>bundle-type-slot/mda.bundle-num</td></tr><tr><td></td><td>bundle keyword</td></tr><tr><td></td><td>type ima ppp</td></tr><tr><td></td><td>bundle-num 1 — 256</td></tr><tr><td>bpgrp-id</td><td>bpgrp-type-bpgrp-num</td></tr><tr><td></td><td>bpgrp keyword</td></tr><tr><td></td><td>type ima, ppp</td></tr><tr><td></td><td>bpgrp-num 1 — 256</td></tr><tr><td>aps-id</td><td>aps-group-id[.channel]</td></tr><tr><td></td><td>aps keyword</td></tr><tr><td></td><td>group-id 1 — 64</td></tr><tr><td>ccag-id</td><td>ccag-id.<path-id>[cc-type]</td></tr><tr><td></td><td>ccag keyword</td></tr><tr><td></td><td>id 1 — 8</td></tr><tr><td></td><td>path-id a, b</td></tr><tr><td></td><td>cc-type [.sap-net .net-sap]</td></tr><tr><td>eth-tunnel-id</td><td>- eth-tunnel-id</td></tr><tr><td></td><td>eth-tunnel keyword</td></tr><tr><td></td><td>id 1 — 1024</td></tr><tr><td>lag-id</td><td>lag-id</td></tr><tr><td></td><td>lag keyword</td></tr><tr><td></td><td>id 1 — 200</td></tr></table>	port-id	slot/mda/port[.channel]	bundle-id	bundle-type-slot/mda.bundle-num		bundle keyword		type ima ppp		bundle-num 1 — 256	bpgrp-id	bpgrp-type-bpgrp-num		bpgrp keyword		type ima, ppp		bpgrp-num 1 — 256	aps-id	aps-group-id[.channel]		aps keyword		group-id 1 — 64	ccag-id	ccag-id.<path-id>[cc-type]		ccag keyword		id 1 — 8		path-id a, b		cc-type [.sap-net .net-sap]	eth-tunnel-id	- eth-tunnel-id		eth-tunnel keyword		id 1 — 1024	lag-id	lag-id		lag keyword		id 1 — 200
port-id	slot/mda/port[.channel]																																														
bundle-id	bundle-type-slot/mda.bundle-num																																														
	bundle keyword																																														
	type ima ppp																																														
	bundle-num 1 — 256																																														
bpgrp-id	bpgrp-type-bpgrp-num																																														
	bpgrp keyword																																														
	type ima, ppp																																														
	bpgrp-num 1 — 256																																														
aps-id	aps-group-id[.channel]																																														
	aps keyword																																														
	group-id 1 — 64																																														
ccag-id	ccag-id.<path-id>[cc-type]																																														
	ccag keyword																																														
	id 1 — 8																																														
	path-id a, b																																														
	cc-type [.sap-net .net-sap]																																														
eth-tunnel-id	- eth-tunnel-id																																														
	eth-tunnel keyword																																														
	id 1 — 1024																																														
lag-id	lag-id																																														
	lag keyword																																														
	id 1 — 200																																														

Standards and Protocol Support

Standards Compliance

IEEE 802.1ab-REV/D3 Station and Media Access Control Connectivity Discovery
IEEE 802.1d Bridging
IEEE 802.1p/Q VLAN Tagging
IEEE 802.1s Multiple Spanning Tree
IEEE 802.1w Rapid Spanning Tree Protocol
IEEE 802.1x Port Based Network Access Control
IEEE 802.1ad Provider Bridges
IEEE 802.1ah Provider Backbone Bridges
IEEE 802.1ag Service Layer OAM
IEEE 802.3ah Ethernet in the First Mile
IEEE 802.1ak Multiple MAC Registration Protocol
IEEE 802.3 10BaseT
IEEE 802.3ad Link Aggregation
IEEE 802.3ae 10Gbps Ethernet
IEEE 802.3ah Ethernet OAM
IEEE 802.3u 100BaseTX
IEEE 802.3x Flow Control
IEEE 802.3z 1000BaseSX/LX
ITU-T Y.1731 OAM functions and mechanisms for Ethernet based networks
ITU-T G.8031 Ethernet linear protection switching

Protocol Support

OSPF

RFC 1765 OSPF Database Overflow
RFC 2328 OSPF Version 2
RFC 2370 Opaque LSA Support
RFC 2740 OSPF for IPv6 (OSPFv3)
draft-ietf-ospf-ospfv3-update-14.txt
RFC 3101 OSPF NSSA Option
RFC 3137 OSPF Stub Router Advertisement
RFC 3623 Graceful OSPF Restart — GR helper

RFC 3630 Traffic Engineering (TE) Extensions to OSPF Version 2
RFC 4203 for Shared Risk Link Group (SRLG) sub-TLV

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 2547bis BGP/MPLS VPNs
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 (previously RFC 1966 & 2796)
RFC 4724 Graceful Restart Mechanism for BGP — GR helper
RFC 4760 Multi-protocol Extensions for BGP
RFC 4893 BGP Support for Four-octet AS Number Space
RFC 5065 Confederations for BGP (obsoletes 3065)

IS-IS

RFC 1142 OSI IS-IS Intra-domain Routing Protocol (ISO 10589)
RFC 1195 Use of OSI IS-IS for routing in TCP/IP & dual environments

RFC 2763 Dynamic Hostname Exchange for IS-IS
RFC 2966 Domain-wide Prefix Distribution with Two-Level IS-IS
RFC 2973 IS-IS Mesh Groups
RFC 3373 Three-Way Handshake for Intermediate System to Intermediate System (IS-IS) Point-to-Point Adjacencies
RFC 3567 Intermediate System to Intermediate System (ISIS) Cryptographic Authentication
RFC 3719 Recommendations for Interoperable Networks using IS-IS
RFC 3784 Intermediate System to Intermediate System (IS-IS) Extensions for Traffic Engineering (TE)
RFC 3787 Recommendations for Interoperable IP Networks
RFC 3847 Restart Signaling for IS-IS — GR helper
RFC 4205 for Shared Risk Link Group (SRLG) TLV
draft-ietf-isis-igp-p2p-over-lan-05.txt

LDP

RFC 3036 LDP Specification
RFC 3037 LDP Applicability
RFC 3478 Graceful Restart Mechanism for LDP — GR helper
RFC 5283 LDP extension for Inter-Area LSP
draft-jork-ldp-igp-sync-03

IPSec

RFC 2401 Security Architecture for the Internet Protocol
RFC 2409 The Internet Key Exchange (IKE)
RFC 3706 IKE Dead Peer Detection
RFC 3947 Negotiation of NAT-Traversal in the IKE
RFC 3948 UDP Encapsulation of IPsec ESP Packets

draft-ietf-ipsec-isakmp-xauth-06.txt —
Extended Authentication within
ISAKMP/Oakley (XAUTH)
draft-ietf-ipsec-isakmp-modecfg-05.txt
— 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 2463 Internet Control Message
Protocol (ICMPv6) for the Internet
Protocol Version 6 Specification
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
RFC3590 Source Address Selection for
the Multicast Listener Discovery
(MLD) Protocol
RFC 3810 Multicast Listener Discovery
Version 2 (MLDv2) for IPv6
RFC 4007 IPv6 Scoped Address
Architecture
RFC 4193 Unique Local IPv6 Unicast
Addresses
RFC 4291 IPv6 Addressing Architecture
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
draft-ietf-isis-ipv6-05
draft-ietf-isis-wg-multi-topology-xx.txt

Multicast

RFC 1112 Host Extensions for IP
Multicasting (Snooping)
RFC 2236 Internet Group Management
Protocol, (Snooping)
RFC 3376 Internet Group Management
Protocol, Version 3 (Snooping)
RFC 2362 Protocol Independent
Multicast-Sparse Mode (PIMSM)
RFC 3618 Multicast Source Discovery
Protocol (MSDP)
RFC 3446 Anycast Rendezvous Point
(RP) mechanism using Protocol
Independent Multicast (PIM) and
Multicast Source Discovery
Protocol (MSDP)
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)
draft-ietf-pim-sm-bsr-06.txt
draft-roosen-vpn-mcast-08.txt
draft-ietf-mboned-msdp-mib-01.txt
draft-ietf-l3vpn-2547bis-mcast-07:
Multicast in MPLS/BGP IP VPNs
draft-ietf-l3vpn-2547bis-mcast-bgp-05:
BGP Encodings and Procedures for
Multicast in MPLS/BGP IP VPNs
RFC 3956: Embedding the Rendezvous
Point (RP) Address in an IPv6
Multicast Address

MPLS

RFC 3031 MPLS Architecture
RFC 3032 MPLS Label Stack
Encoding (REV3443)
RFC 4379 Detecting Multi-Protocol
Label Switched (MPLS) Data Plane
Failures
RFC 4182 Removing a Restriction on the
use of MPLS Explicit NULL

RFC 5332 MPLS Multicast
Encapsulations

RIP

RFC 1058 RIP Version 1
RFC 2082 RIP-2 MD5 Authentication
RFC 2453 RIP Version 2

RSVP-TE

RFC 2430 A Provider Architecture
DiffServ & TE
RFC 2702 Requirements for Traffic
Engineering over MPLS
RFC2747 RSVP Cryptographic
Authentication
RFC3097 RSVP Cryptographic
Authentication
RFC 3209 Extensions to RSVP for
Tunnels
RFC 3564 Requirements for Diff-Serv-
aware TE
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 4875 Extensions to Resource
Reservation Protocol - Traffic
Engineering (RSVP-TE) for Point-
to-Multipoint TE Label Switched
Paths (LSPs)
draft-ietf-mpls-soft-preemption-14
MPLS Traffic Engineering Soft
Preemption
draft-ietf-ccamp-mpls-graceful-
shutdown-06 Graceful Shutdown in
GMPLS Traffic Engineering
Networks
draft-ietf-mpls-p2mp-lsp-ping-06
Graceful Shutdown in GMPLS
Traffic Engineering Networks

DIFFERENTIATED SERVICES

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 3140 Per-Hop Behavior
Identification Codes

TCP/IP

RFC 768 UDP
RFC 1350 The TFTP Protocol (Rev.
RFC 791 IP
RFC 792 ICMP
RFC 793 TCP
RFC 826 ARP
RFC 854 Telnet
RFC 951 BootP (rev)
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
draft-ietf-bfd-mib-00.txt Bidirectional
Forwarding Detection Management
Information Base
draft-ietf-bfd-base-05.txt Bidirectional
Forwarding Detection
draft-ietf-bfd-v4v6-1hop-06.txt BFD
IPv4 and IPv6 (Single Hop)
draft-ietf-bfd-multihop-06.txt 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

VPLS

RFC 4762 Virtual Private LAN Services
Using LDP
draft-ietf-l2vpn-vpls-mcast-reqts-04
draft-ietf-l2vpn-signaling-08

PSEUDO-WIRE

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 (draft-ietf-pwe3-atm-
encap-10.txt)
RFC 4816 PWE3 ATM Transparent Cell
Transport Service (draft-ietf-pwe3-
cell-transport-04.txt)
RFC 4448 Encapsulation Methods for
Transport of Ethernet over MPLS
Networks (draft-ietf-pwe3-ethernet-
encap-11.txt)

Standards and Protocols

RFC 4619 Encapsulation Methods for Transport of Frame Relay over MPLS Networks (draft-ietf-pwe3-frame-relay-07.txt)

RFC 4446 IANA Allocations for PWE3

RFC 4447 Pseudowire Setup and Maintenance Using LDP (draft-ietf-pwe3-control-protocol-17.txt)

RFC 5085, Pseudowire Virtual Circuit Connectivity Verification (VCCV): A Control Channel for Pseudowires

draft-ietf-l2vpn-vpws-iw-oam-02.txt

draft-ietf-pwe3-oam-msg-map-05.txt

draft-ietf-l2vpn-arp-mediation-04.txt

draft-ietf-pwe3-ms-pw-arch-05.txt

draft-ietf-pwe3-segmented-pw-11.txt

draft-hart-pwe3-segmented-pw-vccv-02.txt

draft-muley-dutta-pwe3-redundancy-bit-02.txt

draft-muley-pwe3-redundancy-02.txt

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

draft-ietf-ancp-framework-01.txt

draft-ietf-ancp-protocol-00.txt

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

RADIUS

RFC 2865 Remote Authentication Dial In User Service

RFC 2866 RADIUS Accounting

SSH

draft-ietf-secsh-architecture.txt SSH Protocol Architecture

draft-ietf-secsh-userauth.txt SSH Authentication Protocol

draft-ietf-secsh-transport.txt SSH Transport Layer Protocol

draft-ietf-secsh-connection.txt SSH Connection Protocol

draft-ietf-secsh-newmodes.txt SSH Transport Layer Encryption Modes

TACACS+

draft-grant-tacacs-02.txt

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.

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 2012 TCP-MIB

RFC 2013 UDP-MIB

RFC 2096 IP-FORWARD-MIB

RFC 2138 RADIUS

RFC 2206 RSVP-MIB

RFC 2452 IPv6 Management Information Base for the Transmission Control Protocol

RFC 2454 IPv6 Management Information Base for the User Datagram Protocol

RFC 2465 Management Information Base for IPv6: Textual Conventions and General Group

RFC 2558 SONET-MIB

RFC 2571 SNMP-FRAMEWORK-MIB

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 2665 EtherLike-MIB

RFC 2819 RMON-MIB

RFC 2863 IF-MIB

RFC 2864 INVERTED-STACK-MIB

RFC 2987 VRRP-MIB

RFC 3014 NOTIFICATION-LOG-MIB

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 5101 - Specification of the IP Flow
 Information Export (IPFIX)
 Protocol for the Exchange of IP
 Traffic Flow Information
 draft-ietf-disman-alarm-mib-04.txt
 draft-ietf-ospf-mib-update-04.txt
 draft-ietf-mpls-lsr-mib-06.txt
 draft-ietf-mpls-te-mib-04.txt
 draft-ietf-mpls-ldp-mib-07.txt
 draft-ietf-isis-wg-mib-05.txt
 IANA-IFType-MIB
 IEEE8023-LAG-MIB

TIMETRA-NG-BGP-MIB.mib
 TIMETRA-OAM-TEST-MIB.mib
 TIMETRA-OSPF-NG-MIB.mib
 TIMETRA-OSPF-V3-MIB.mib
 TIMETRA-PIM-NG-MIB.mib
 TIMETRA-PORT-MIB.mib
 TIMETRA-PPP-MIB.mib
 TIMETRA-QOS-MIB.mib
 TIMETRA-RIP-MIB.mib
 TIMETRA-ROUTE-POLICY-MIB.mib
 TIMETRA-RSVP-MIB.mib
 TIMETRA-SECURITY-MIB.mib
 TIMETRA-SERV-MIB.mib
 TIMETRA-SUBSCRIBER-
 MGMTMIB.mib
 TIMETRA-SYSTEM-MIB.mib
 TIMETRA-TC-MIB.mib
 TIMETRA-VRRP-MIB.mib
 TIMETRA-VRTR-MIB.mib

Proprietary MIBs

TIMETRA-APS-MIB.mib
 TIMETRA-ATM-MIB.mib
 TIMETRA-BGP-MIB.mib
 TIMETRA-BSX-NG-MIB.mib
 TIMETRA-CAPABILITY-7750-
 V4v0.mib
 TIMETRA-CFLOWD-MIB.mib
 TIMETRA-CHASSIS-MIB.mib
 TIMETRA-CLEAR-MIB.mib
 TIMETRA-FILTER-MIB.mib
 TIMETRA-GLOBAL-MIB.mib
 TIMETRA-IGMP-MIB.mib
 TIMETRA-ISIS-MIB.mib
 TIMETRA-LAG-MIB.mib
 TIMETRA-LDP-MIB.mib
 TIMETRA-LOG-MIB.mib
 TIMETRA-MIRROR-MIB.mib
 TIMETRA-MPLS-MIB.mib

Index

A

- accounting
 - volume and time 705
- ANCP 643
- Anti-spoofing filters 466
 - filtering packets 467
 - types 467
- ARP 473
 - local proxy ARP 474
 - reply agent 473

D

- DHCP
 - CLI 159
 - configuring 322
 - Option 82 handling 324
 - relay 325
 - snooping 322
 - lease state table 309
 - Option 82 304
 - overview 302
 - relay 304
 - snooping 307
- distributed service edges 23
- dual homing 688
- Dynamic tables 637

E

- Enhanced Subscriber Management 626
 - Configuring 711
 - Applying the profiles and policies 721
 - Basic configurations 713
 - Dynamic tables 637
 - Entities 640
 - Explicit subscriber profile mapping 632
 - Hosts and subscribers 629
 - Instantiating a host 641
 - Layer 3 subscriber interfaces 662
 - Profiles
 - SLA Profile 652

- Subscriber and SLA of a host 649
- Subscriber Profile 650

QoS

- 656
- In different profiles 656
- Policy overrides 656
- RADIUS authentication 605
- Scripts 648
 - Python 648
- SLA profile 632
- Subscriber identification policy 630
- Subscriber identification string 630
- Subscriber profile 631
- Subscriber SAPs 628
- Enhanced Subscriber Management (ESM) 626
- Services
 - Epipe
 - overview 45
- ESM entities 640

G

- GSMP 646

L

- Local User Database 326

M

- managed SAP (M-SAP) 701
- Multicast
 - IGMP 525
 - PIM 527
- multicast management 704
- Multi-chassis
 - Overview 657
 - synchronization 657, 691

P

- PPPoE 424
- Python scripting language 648

R

RADIUS-based accounting 620

Routed CO 660

S

Security features

MAC pinning 469

MAC protection 469

redirect policies 472

Service

entities

SDPs

bindings 41

policies 34

Services 32

Alcatel Service Model 35

configuring

command reference 61, 333, 497, 557

entities

customers 36

SAPs 37

configuration considerations 39

encapsulation types 38

Ethernet 38

dot1q 38

null 38

qinq 38

SDPs 41

encapsulation

GRE 43

encapsulation types 43

keepalives 44

spoke and mesh 42

IES

IP interfaces 49

overview 48

types 33

VPLS

overview 46

VPRN

overview 50

SRRP 692

Static subscriber hosts 655

Subscriber Management 603

mode

enhanced 626

standard 626

subscriber mirroring 704

T

TPSDA 18

BSA 23

BSR 23

DSLAM 24

infrastructure 22

service differentiation 25

Triple Play 301, 423, 465, 523

V

Services

VPRN overview 50

W

Web portal redirect 475