Triple Play Service Configuration Commands

Global Commands

shutdown

Syntax [**no**] **shutdown**

Context
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>lcl-addr-assign
config>service>vprn>sub-if>grp-if>lcl-addr-assign
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

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** *description-string*
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 *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.

Service Commands

vpls

Syntax **vpls** *service-id* **customer** *customer-id* **vpn** *vpn-id* [**m-vpls**]
vpls *service-id*
no vpls *service-id*

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 *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.

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 *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.

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** *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.

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 *service-id*. 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.

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 router 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)

m-vpls — Specifies a managed VPLS.

service-name

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

Context config>service>epipe
config>service>ies
config>service>vpls
config>service>vprn

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 and 7450 ESS platforms.

All services are required to assign a service ID to initially create a service. However, either the service ID or the service name can be used o identify and reference a given service once it is initially created.

Parameters *service-name* — Specifies a unique service name to identify the service. Service names may not begin with an integer (0-9).

ies

Syntax **ies** *service-id* **customer** *customer-id* **vpn** *vpn-id*
ies *service-id*
no ies *service-id*

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 *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.

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 router 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** *service-id* [**customer** *customer-id*]
no vprn *service-id*

Context config>service

Description This command creates or edits a Virtu9al 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 attempting to edit a service with the incorrect *customer-id* results in an error.

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

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

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

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

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

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

Values 1 — 2147483647

sap

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

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

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

All SAPs must be explicitly created. If no SAPs are created within a service or on an IP interface, a SAP will not exist on that object.
Enter an existing SAP without the **create** keyword to edit SAP parameters. The SAP is owned by the service in which it was created.
A SAP can only be associated with a single service. A SAP can only be defined on a port that has been configured as an access port using the **config interface** *port-type port-id* **mode access** command. Channelized TDM ports are always access ports.
If a port is shutdown, all SAPs on that port become operationally down. When a service is shutdown, SAPs for the service are not displayed as operationally down although all traffic traversing the service will be discarded. The operational state of a SAP is relative to the operational state of the port on which the SAP is defined.

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

Default No SAPs are defined.

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition. See Common CLI Command Descriptions on page 1509 for command syntax.

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

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

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

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

create — Keyword used to create a SAP instance.

disable-aging

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

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

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

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

When **no disable-aging** is specified for a VPLS, it is possible to disable aging for specifc 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** *high-water-mark*

Context config>service>vpls

Description This command specifies the value to send logs and traps when the threshold is reached.

Default 95

Parameters *high-water-mark* — When to send logs and traps.

Values 0 — 100

fdb-table-low-wmark

Syntax [**no**] **fdb-table-low-wmark** *low-water-mark*

Context config>service>vpls

Description This command specifies the value to send logs and traps when the threshold is reached.

Default 90

Parameters *low-water-mark* — When to send logs and traps.

Values 0 — 100

fdb-table-size

Syntax **fdb-table-size** *table-size*
no fdb-table-size [*table-size*]

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 maxium FDB table size to default.

Default **250** — Forwarding table of 250 MAC entries.

Parameters *table-size* — 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** *seconds*
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 *seconds* — The aging time for local MACs expressed in seconds.

Values 60 — 86400

remote-age

Syntax **remote-age** *seconds*
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 *seconds* — 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** [*group-name*] [*residential-group*]

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

residential-group — 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 not 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** *sap-id* [**split-horizon-group** *group-name*] [**create**] [**capture-sap**]
no sap *sap-id*

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 *sap-id* — Specifies the physical port identifier portion of the SAP definition. See Common Service Commands on page 1510 for *sap-id* command syntax.

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

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

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

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

split-horizon-group *group-name* **—** Specifies an existing split horizon group name.

capture-sap — Keyword to create a capture SAP.

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 *policy-id* 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 acccounting policy reverts to the default.

Default Default accounting policy.

Parameters *acct-policy-id* — Enter the accounting *policy-id* as configured in the **config>log>accounting-policy** context.

Values 1 — 99

collect-stats

Syntax [**no**] **collect-stats**

Context config>service>vpls>sap
config>service>vpls>spoke-sdp
config>service>vpls>mesh-sdp
config>service>ies>if>sap

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

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 *acl — cflowd* configuration associated with a filter.

interface — cflowd 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

track-srrp

Syntax **track-srrp** *srrp-id*
no track-srrp

Context configure>service>vpls>sap>

Description This is a capture SAP level command. This command is important in PPPoE deployments with MSAPs. PPPoE operation requires that the MAC address learned by the client at the very beginning of the session negotiation phase remains unchanged for the lifetime of the session (RFC 2516). This command will ensure that that the virtual MAC address used during the PPPoE session negotiation phase on the capture SAP is the same virtual MAC address that is used by the SRRP on the group-interface on which the session is established. Therefore, it is mandated that the SRRP instance (and implicitly the group-interface) where the session belongs to is known in advance. If the group-interface name for the session is returned by the RADIUS, it must be ensured that this group-interface

is the one on which the tracked SRRP instance is configured. PPPoE sessions on the same capture SAP cannot be shared across multiple group-interfaces, but instead they all must belong to a single group-interface that is known in advance.

The same restrictions will apply to IPoE clients in MC Redundancy scenario if they are to be supported concurrently on the same capture SAP as PPPoE.

The supported capture SAP syntax is this:

sap <port-id>:X.* capture-sap

The capture SAP syntax that is NOT supported is this:

sap <port-id>:*.* capture-sap

Default None

Parameters *srrp-id* — Specify the SRRP instance number.

Values 1..4294967295

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** *atm-encap-type*

Context config>service>vpls>sap>atm

Description This command specifies the data encapsulation for an ATM PVCC delimited SAP. The definition references RFC 2684, *Multiprotocol Encapsulation over ATM AAL5,* 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 *atm-encap-type* — 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>vpls>sap>atm

Description This command enables the context to configure ingress ATM attributes for the SAP.

traffic-desc

Syntax **traffic-desc** *traffic-desc-profile-id*
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 *traffic-desc-profile-id* — 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** *name*
no authentication-policy

Context config>service>vpls>sap
config>service>ies>sub-if>grp-if
config>service>vprn>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 *name* — 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 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 *default-sla-profile-name* — 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** *default-subscriber-profile-name*

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 *default-sub-profile* — 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** *sub-ident-policy-name*

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 identifcation policy from the SAP configuration.

Default no sub-ident-policy

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

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

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

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

  If *sub-ident* is not enabled on the SAP arp-reply-agent, subscriber identification matching is not performed on ARP requests received on the SAP.
  ARP requests are never forwarded back to the same SAP or within the receiving SAP's Split Horizon Group.

profiled-traffic-only

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 profiled traffic refers to single subscriber traffic on a dedicated SAP (in the VLAN-per-subscriber model). When enabled, subscriber queues are instantiated through the QOS policy defined in the sla-profile and the associated SAP queues are deleted. This can increase subscriber scaling by reducing the number of queues instantiated per subscriber (in the VLAN-per-subscriber model). In order for this to be achieved, any configured multi-sub-sap limit must be removed (leaving the default of 1).

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

non-sub-traffic

Syntax **non-sub-traffic sub-profile** *sub-profile-name* **sla-profile** *sla-profile-name* [**subscriber** *sub-ident-string*]
no non-sub-traffic

Context config>service>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. It is used in conjunction with the profiled-traffic-only on single subscriber SAPs and creates a subscriber host which is used to forward non-IP traffic through the single subscriber SAP without the need for SAP queues.

The **no** form of the command removes removes the profiles and disables the feature.

Parameters **sub-profile** *sub-profile-name* **—** 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 *sla-profile-name* **—** 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 *sub-ident-string* — 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 *sub-ident* parameter, the static subscriber host's sub-ident-string is used to determine whether an ARP request received on the SAP is sourced from a host belonging to the same subscriber as the destination host. When both the destination and source hosts from the ARP request are known on the SAP and the subscriber identifications do not match, the ARP request may be forwarded to the rest of the destinations.

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

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

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

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

profiled-traffic-only

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 profiled traffic refers to single subscriber traffic on a dedicated SAP (in the VLAN-per-subscriber model). When enabled, subscriber queues are instantiated through the QOS policy defined in the sla-profile and the associated SAP queues are deleted. This can increase subscriber scaling by reducing the number of queues instantiated per subscriber (in the VLAN-per-subscriber model). In order for this to be achieved, any configured multi-sub-sap limit must be removed (leaving the default of 1).

The **no** form of the command 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** *vpls-id*
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 *vpls-id* — Specifies the MVR VPLS from which multicast channels should be copied into this SAP.

group

Syntax [**no**] **group** *grp-address*

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 *grp-address* — 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** *policy-name*
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 *policy-name* — 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 Enabling IGMP Group Membership Report Filtering on page 678.

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** *policy-name*
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 *policy-name* — 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** *tenths-of-seconds*
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.

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

Default 10

Parameters *seconds* — Specifies the frequency, in tenths of seconds, at which query messages are sent.

Values 1 — 50

max-num-groups

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

unconstrained-bw

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

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

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

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

Values 0 — 2147483647

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

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

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

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

Values 0 — 2147483647

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

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

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

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

Default no mrouter-port

mvr

Syntax **mvr**

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

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

query-interval

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

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

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

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

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

Default 125

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

Values 2 — 1024

query-src-ip

Syntax **query-src-ip** *ipv6-address*
no query-src-ip

Context config>service>vpls>mld-snooping
config>service>vpls>igmp-snooping

Description This command configures the IP source address used in IGMP or MLD queries.

query-response-interval

Syntax **query-response-interval** *seconds*

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 *seconds* — 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** *ip-address*
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 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 *ip-address* — 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** *sap-id*
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 *sap-id* — Specifies the SAP to which multicast channels should be copied. See Common Service Commands on page 1510 for *sap-id* 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.

diameter-auth-policy

Syntax **diameter-auth-policy** *name*
no diameter-auth-policy

Context config>service>vpls>sap

Description This command is used to configure the Diameter NASREQ application policy to use for authentication.

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

host-limit

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

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

Description This command configures the maximum number of ARP hosts.

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

Default 1

Parameters *max-num-hosts* — specifies the maximum number of ARP hosts allowed on this SAP. Note that the operational maximum value may be smaller due to equipped hardware dependencies.

Values 1 — 131071

min-auth-interval

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

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

Description This command configures the minimum authentication interval.

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

Default 15

Parameters *min-auth-interval* — Specifies the minimum authentical 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 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

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** *ip-address*] [**mac** *mac-address*]} [**subscriber-sap-id** | **subscriber** *sub-ident-string*] [**sub-profile** *sub-profile-name*] [**sla-profile** *sla-profile-name*] [**ancp-string** *ancp-string*] [**inter-dest-id** *intermediate-destination-id*] [**mac-linking** *ip-address*][**rip-policy** *rip-policy-name*]
no host {[**ip** *ip-address*] [**mac** *ieee-address*]}
no host all [**force**]

Context config>service>vpls>sap
config>service>ies>if>sap
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 *ip-address* and *mac-address* must match the host's exact IP and MAC addresses as defined when it was created. When a static host is removed from the SAP, the corresponding anti-spoof filter entry and/or FDB entry is also removed.

Default none

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

mac *mac-address* — Specify this optional parameter when defining a static host. The MAC address must be specified for **anti-spoof mac** 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 *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 the 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.

mac-linking

Syntax **mac-linking** *ip-address*
no mac-linking

Context config>service>ies>sub-if>grp-if>sap>static-host

Description This command associates this IPv6 host to the specified IPv4 host through the learned MAC address. A learned MAC from the IPv6 host will be associated with the IPv4 host and vice versa.

Default none

rip-policy

Syntax **rip-policy** *rip-policy-name*
no rip-policy

Context config>service>ies>sub-if>grp-if>sap>static-host

Description This command specifies the name of the RIP policy up to 32 characters inlength.

The no form of the command removes the policy name from the static-host configuration.

Default none

igmp-host-tracking

Syntax **igmp-host-tracking**

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

Description This command enables the context to configure IGMP host tracking parameters.

disable-router-alert-check

Syntax [**no**] **disable-router-alert-check**

Context config>service>vprn>sub-if>grp-if>sap>igmp-host-tracking

Description This command enables the IGMP router alert check option.

The **no** form of the command disables the router alert check.

expiry-time

Syntax **expiry-time** *expiry-time*
no expiry-time

Context config>service>vprn>sub-if>grp-if>sap>igmp-snooping

Description This command configures the time that the system continues to track inactive hosts.

The **no** form of the command removes the values from the configuration.

Default no expiry-time

Parameters *expiry-time* — Specifies the time, in seconds, that this system continues to track an inactive host.

Values 1 — 65535

import

Syntax **import** *policy-name*
no import

Context config>service>vprn>sub-if>grp-if>sap>igmp-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 *policy-name* — The routing policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters excluding double quotes. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes. Routing policies are configured in the config>router>policy-options context The router policy must be defined before it can be imported.

max-num-group

Syntax **max-num-groups** *max-num-groups*
no max-num-groups

Context config>service>vprn>sub-if>grp-if>sap>igmp-snooping

Description This command configures the maximum number of multicast groups allowed to be tracked.

The **no** form of the command removes the values from the configuration.

Default no max-num-groups

Parameters *max-num-groups* — Specifies the maximum number of multicast groups allowed to be tracked.

Values 1 — 196607

max-num-sources

Syntax **max-num-sources** *max-num-sources*
no max-num-sources

Context config>service>vprn>sub-if>grp-if>sap>igmp-snooping

Description This command configures the maximum number of multicast sources allowedto be tracked per group.

The no form of the command removes the value from the configuration.

Parameters *max-num-sources* — Specifies the maximum number of multicast sources allowedto be tracked per group.

Values 1 — 1000

max-num-grp-sources

Syntax **max-num-grp-sources** [1..32000]
no max-num-grp-sources

Context config>service>vprn>sub-if>grp-if>sap>igmp-snooping

Description This command configures the maximum number of group sources for which IGMP can have local receiver information based on received IGMP reports on this interface. When this configuration is changed dynamically to a value lower than currently accepted number of group sources, the group sources that are already accepted are not deleted. Only new group sources will not be allowed. When this object has a value of 0, there is no limit to the number of group sources.

The **no** form of the command removes the value from the configuration.

Default no max-num-grp-sources

Parameters **1..32000 —** Specifies the maximum number of multicast sources allowed to be tracked per group

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** *ip-filter-id*
filter ipv6 *ipv6-filter-id*
filter mac *mac-filter-id*
no filter [**ip** *ip-filter-id*] [**mac** *mac-filter-id*] [**ipv6** *ipv6-filter-id*]

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** *ip-filter-id* — Specifies IP filter policy. The filter ID must already exist within the created IP filters.

Values 1 — 65535

ipv6 *ipv6-filter-id* — 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

agg-rate

Syntax **[no] agg-rate**

Context config>service>vprn>subscriber-interface>group-if>sap>egress
config>service>ies>subscriber-interface>group-if>sap>egress

Description This command is used to control an HQoS aggregate rate limit. It is used in conjunction with the following parameter commands: **rate**, **limit-unused-bandwidth**, **and queue-frame-based-accounting**.

rate

Syntax **rate** {**max** | **rate**}
no rate

Context config>service>vprn>subscriber-interface>group-if>sap>egress>agg-rate
config>service>ies>subscriber-interface>group-if>sap>egress>agg-rate

Description This command defines the enforced aggregate rate for all queues associated with the agg-rate context. A rate must be specified for the agg-rate context to be considered to be active on the context's object (SAP, subscriber, VPORT etc.).

limit-unused-bandwidth

Syntax [**no**] **limit-unused-bandwidth**

Context config>service>vprn>subscriber-interface>group-if>sap>egress>agg-rate
config>service>ies>subscriber-interface>group-if>sap>egress>agg-rate

Description This command is used to enable (or disable) aggregate rate overrun protection on the agg-rate context.

queue-frame-based-accounting

Syntax [**no**] **queue-frame-based-accounting**

Context config>service>vprn>subscriber-interface>group-if>sap>egress>agg-rate

Description This command is used to enabled (or disable) frame based accounting on all queues associated with the agg-rate context. Only supported on Ethernet ports. Not supported on HSMDA Ethernet ports.

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** *ip-filter-id* **—** Specifies IP filter policy. The filter ID must already exist within the created IP filters.

Values 1 — 65535

ipv6 *ipv6-filter-id* **—** 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** *group-name*
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 *group-name* — The *group-name* 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** *policy-id*
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 *policy-id* 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** *queue-id*

Context config>service>ies>if>sap>egress>queue-override
config>service>ies>if>sap>ingress>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 *queue-id* — 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>if>sap>ingress>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** *percent*
no avg-frame-overhead

Context config>service>ies>if>sap>egress>queue-override
config>service>ies>if>sap>ingress>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 x 0.1 or 1000 octets.

  For egress Ethernet queues, the frame encapsulation overhead is calculated by multiplying the number of offered-packets for the queue by 20 bytes. If a queue was offered 50 packets then the frame encapsulation overhead would be 50 x 20 or 1000 octets.
- Frame based offered-load — The frame based offered-load is calculated by adding the offered-load to the frame encapsulation overhead. If the offered-load is 10000 octets and the encapsulation overhead was 1000 octets, the frame based offered-load would equal 11000 octets.
- Packet to frame factor — The packet to frame factor is calculated by dividing the frame encapsulation overhead by the queue's offered-load (packet based). If the frame encapsulation overhead is 1000 octets and the offered-load is 10000 octets then the packet to frame factor would be 1000 / 10000 or 0.1. When in use, the avg-frame-overhead will be the same as the packet to frame factor making this calculation unnecessary.
- Frame based CIR — The frame based CIR is calculated by multiplying the packet to frame factor with the queue's configured CIR and then adding that result to that CIR. If the queue CIR is set at 500 octets and the packet to frame factor equals 0.1, the frame based CIR would be 500 x 1.1 or 550 octets.
- Frame based within-cir offered-load — The frame based within-cir offered-load is the portion of the frame based offered-load considered to be within the frame-based CIR. The frame based within-cir offered-load is the lesser of the frame based offered-load and the frame based CIR. If the frame based offered-load equaled 11000 octets and the frame based CIR equaled 550 octets, the frame based within-cir offered-load would be limited to 550 octets. If the frame based offered-load equaled 450 octets and the frame based CIR equaled 550 octets, the frame based within-cir offered-load would equal 450 octets (or the entire frame based offered-load).

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

- Frame based PIR — The frame based PIR is calculated by multiplying the packet to frame factor with the queue's configured PIR and then adding the result to that PIR. If the queue PIR is set to 7500 octets and the packet to frame factor equals 0.1, the frame based PIR would be 7500 x 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 *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>service>ies>if>sap>egress>queue-override>queue
config>service>ies>if>sap>ingress>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 *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 — 1048576 or default

high-prio-only

Syntax **high-prio-only** *percent*
no high-prio-only

Context config>service>ies>if>sap>egress>queue-override>queue
config>service>ies>if>sap>ingress>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 *percent* — The *percent* parameter is the percentage reserved for high priority traffic on the queue. If a value of 10KBytes is desired, enter the value 10. A value of 0 specifies that none of the MBS of the queue will be reserved for high priority traffic. This does not affect RED slope operation for packets attempting to be queued.

Values 0 — 100 | default

mbs

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

Context config>service>ies>if>sap>egress>queue-override>queue
config>service>ies>if>sap>ingress>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 *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 — 1073741824 or default

rate

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

Context config>service>ies>if>sap>egress>queue-override>queue
config>service>ies>if>sap>ingress>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 *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 *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 — 2000000000

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 — 2000000000, **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 opererationally down state) after all configured mesh SDPs are in opererationally 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** *table-size*
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 *table-size* — The maximum number of MAC entries in the FDB from this SAP.

Values 1 — 511999

multi-service-site

Syntax **multi-service-site** *customer-site-name*
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 *customer-site-name*. If the specified *customer-site-name* does not exist in the context of the service customer ID an error occurs and the command will not execute. If *customer-site-name* exists, the current and future defined queues on the SAP (ingress and egress) will attempt to use the scheduler hierarchies created within *customer-site-name* 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.

retail-service-id

Syntax **retail-svc-id** *service-id*
no retail-svc-id

Context config>service>vprn>sub-if>grp-if>sap>static-host

Description This command specifies the service id of the retailer IES/VPRN service to which the static IPv6 host belongs. A corresponding retailer subscriber interface must exist in the specified service.

Default no retai-svc-id

Parameters *service-id* — Specifies the retailer service id.

Values 1— 2148007978 or service name up to 64 characters in length

static-host

Syntax **static-host ip** *ip-prefix*[/*prefix-length*] [**mac** *ieee-address*] [**create**]
no static-host ip *ip-prefix*[/*prefix-length*] **mac** *ieee-address*
no static-host all [**force**]
no static-host ip *ip-prefix*[/*prefix-length*]

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-prefix[/prefix-length* **—** Specifies the IPv4 address, IPv6 address or the IPv6 prefix.

Values

ip-prefix[/prefix* : ipv4-prefix	a.b.c.d (host bits must be 0)
ipv4-prefix-le	[0..32]
ipv6-prefix	x:x:x:x:x:x:x:x (eight 16-bit pieces) x:x:x:x:x:x:d.d.d.d x - [0..FFFF]H d - [0..255]D
ipv6-prefix-le - [0..128]	

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
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** {*ip-prefix/length* | *ip-prefix netmask*} [**create**]
no route {*ip-prefix/length* | *ip-prefix netmask*} [**metric** *metric-value*]

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.

Parameters *ipv6-prefix/length* | *ipv6-prefix netmask* — This parameter associates an IPv6 managed route to the IPv6 static host. The IPv6 managed routes can overlap with the static host IPv6 address.

ipv4-prefix/length | *ipv6-prefix netmask* — This parameter associates an IPv4 managed route to the IPv4 static host.

Note: A maximum of 16 managed routes can be associated to a static host. IPv4 hosts can only have IPv4 managed routes and IPv6 hosts can only have IPv6 managed routes.

metric *metric-value* — A metric can be associated with the provisioned managed route.

sla-profile

Syntax **sla-profile** *sla-profile-name*
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 *sla-profile-name* — Specifies the SLA profile name.

sub-profile

Syntax **sub-profile** *sub-profile-name*
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 *sub-profile-name* — Specifies the sub-profile name.

subscriber

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

subscriber-sap-id

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

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

Syntax static-host-mgmt

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 common parameters for static hosts.

mac-learning-options

Syntax [**no**] **mac-learning-options**

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

Description This command configures additional methods by which the BNG will learn the subscriber host MAC.

data-triggered

Syntax [**no**] **data-triggered**

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

Description This command enables learning of MAC addresses from data packets.

The **no** form of the command disables learning of MAC addresses from data packets.

Default no data-triggered

single-mac

Syntax [**no**] **single-mac**

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

Description This command controls how the SAP will learn the static host MAC address. Enabling this command indicates that this particular SAP will only have one subscriber and will only have one MAC address for all hosts. With this parameter enabled, the subscriber's NS and RS source MAC address are used to automatically to populate the subscriber MAC address. To allow this auto-populate behavior, the subscriber's NS and RS source IP must be of type link local address.

Default no single-mac

static-mac

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

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

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

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

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

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

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

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

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

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

VPLS and IES SDP and SAP Commands

mesh-sdp

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

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

Description This command binds a VPLS service to an existing Service Distribution Point (SDP).

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

Note that this command creates a binding between a service and an SDP. The SDP has an operational state which determines the operational state of the SDP within the service. For example, if the SDP is administratively or operationally down, the SDP for the service will be down.

The SDP must already be defined in the **config>service>sdp** context in order to associate the SDP with an Epipe or VPLS service. If the **sdp** *sdp-id* is not already configured, an error message is generated. If the *sdp-id* does exist, a binding between that *sdp-id* and the service is created.

SDPs must be explicitly associated and bound to a service. If an SDP is not bound to a service, no far-end router devices can participate in the service.

The **no** form of this command removes the SDP binding from the service. The SDP configuration is not affected; only the binding of the SDP to a service. Once removed, no packets are forwarded to the far-end router.

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

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

Parameters *sdp-id* — The SDP identifier.

Values 1 — 17407

vc-id — The virtual circuit identifier. This value is used to validate the VC ID portion of each mesh SDP binding defined in the service. The default value of this object is equal to the service ID.

Values 1 — 4294967295

vc-type — This command overrides the default VC type signaled for the spoke or mesh binding to the far end of the SDP. The VC type is a 15 bit-quantity containing a value which represents the type of VC. The actual signaling of the VC type depends on the signaling parameter defined for the SDP. If signaling is disabled, the **vc-type** command can still be used to define the dot1q value expected by the far-end provider equipment. A change of the binding's VC type causes the binding to signal the new VC type to the far end when signaling is enabled.

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

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

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

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

spoke-sdp

Syntax **spoke-sdp** *sdp-id*[**:***vc-id*] [**vc-type** {**ether** | **vlan**}] [**split-horizon-group** *group-name*]
no spoke-sdp *sdp-id*[**:***vc-id*]

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 *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 service. Each SDP must use unique *vc-ids*. An error message is generated if two SDP bindings with identical *vc-ids* terminate on the same router. Split horizon groups can only be created in the scope of a VPLS service.

Parameters *sdp-id* — The SDP identifier.

Values 1 — 17407

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

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

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

spoke-sdp

Syntax **spoke-sdp** *sdp-id*[**:***vc-id*]
no spoke-sdp *sdp-id*[**:***vc-id*]

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 *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 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 *sdp-id* is bound to a service.

Special Cases **IES —** At most, only one *sdp-id* can be bound to an IES service.

Parameters *sdp-id* — The SDP identifier. Allowed values are integers in the range of 1 and 17407 for existing SDPs.

vc-id — The virtual circuit identifier.

Values 1 — 4294967295

egress

Syntax **egress**

Context config>service>vpls>mesh-sdp
config>service>vpls>spoke-sdp
config>service>ies>if>spoke-sdp

Description This command configures the egress SDP context.

ingress

Syntax **ingress**

Context config>service>vpls>mesh-sdp
config>service>vpls>spoke-sdp
config>service>ies>if>spoke-sdp

Description This command configures the ingress SDP context.

vc-label

Syntax [**no**] **vc-label** *egress-vc-label*

Context config>service>vpls>mesh-sdp>egress
config>service>vpls>spoke-sdp>egress
config>service>ies>if>spoke-sdp>egress

Description This command configures the egress VC label.

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

Values 16 — 1048575

vc-label

Syntax [**no**] **vc-label** *ingress-vc-label*

Context config>service>vpls>mesh-sdp>ingress
config>service>vpls>spoke-sdp>ingress
config>service>ies>if>spoke-sdp>ingress

Description This command configures the ingress VC label.

Parameters *vc-label* — A VC ingress value that indicates a specific connection.

Values 2048 — 18431

vlan-vc-tag

Syntax **vlan-vc-tag** *0..4094*
no vlan-vc-tag [*0..4094*]

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

Description This command specifies an explicit dot1q value used when encapsulating to the SDP far end. When signaling is enabled between the near and far end, the configured dot1q tag can be overridden by a received TLV specifying the dot1q value expected by the far end. This signaled value must be stored as the remote signaled dot1q value for the binding. The provisioned local dot1q tag must be stored as the administrative dot1q value for the binding.

When the dot1q tag is not defined, the default value of zero is stored as the administrative dot1q value. Setting the value to zero is equivalent to not specifying the value.

The **no** form of this command disables the command

Default no vlan-vc-tag

Parameters *0..4094* — Specifies a valid VLAN identifier to bind an 802.1Q VLAN tag ID.

avg-frame-overhead

Syntax **avg-frame-overhead** *percent*
no avg-frame-overhead

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

Description This command configures the average frame overhead to define the average percentage that the offered load to a queue will expand during the frame encapsulation process before sending traffic on-the-wire. While the avg-frame-overhead value may be defined on any queue, it is only used by the system for queues that egress a Sonet or SDH port or channel. Queues operating on egress Ethernet ports automatically calculate the frame encapsulation overhead based on a 20 byte per packet rule (8 bytes for preamble and 12 bytes for Inter-Frame Gap).

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

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

the avg-frame-overhead equals 10%, the frame encapsulation overhead would be 10000 x 0.1 or 1000 octets.

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

- Frame based offered-load — The frame based offered-load is calculated by adding the offered-load to the frame encapsulation overhead. If the offered-load is 10000 octets and the encapsulation overhead was 1000 octets, the frame based offered-load would equal 11000 octets.
- Packet to frame factor — The packet to frame factor is calculated by dividing the frame encapsulation overhead by the 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 x 1.1 or 550 octets.
- Frame based within-cir offered-load — The frame based within-cir offered-load is the portion of the frame based offered-load considered to be within the frame-based CIR. The frame based within-cir offered-load is the lesser of the frame based offered-load and the frame based CIR. If the frame based offered-load equaled 11000 octets and the frame based CIR equaled 550 octets, the frame based within-cir offered-load would be limited to 550 octets. If the frame based offered-load equaled 450 octets and the frame based CIR equaled 550 octets, the frame based within-cir offered-load would equal 450 octets (or the entire frame based offered-load).

  As a special case, when a queue or associated intermediate scheduler is configured with a CIR-weight equal to 0, the system automatically sets the 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 x 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 *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

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** *queue-id*

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 *queue-id* — 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** *size-in-kbytes*
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 *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 — 1048576 or default

high-prio-only

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

Values 0 — 100 | default

mbs

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

Context config>service>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 *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 — 1073741824 or default

mbs

Syntax **mbs** {*size-in-kbytes* | **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 enqueuing 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 — 1073741824 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 — 2000000000

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 — 2000000000, **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** *pir-rate* [**cir** *cir-rate*]
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 then 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 — 10000000, **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 7 defines the default behavior for Dot1P evaluation when the **match-qinq-dot1p** command is not executed.

Table 7: 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 8 defines the dot1p evaluation behavior when the top parameter is specified.

Table 8: 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 8: 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 9 defines the dot1p evaluation behavior when the bottom parameter is specified.

Table 9: 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 139) 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 specificially 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** *vc-id*

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 97. 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** *ieee-address*

Context config>service>vpls>mac-protect

Description This command specifies the 48-bit IEEE 802.3 MAC address.

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.

mac-subnet-length

Syntax **mac-subnet-length** *subnet-length*
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 *subnet-length* — Specifies the number of bits to be considered when performing MAC learning or MAC switching.

Values 24 — 48

mcr-default-gtw

Syntax **mcr-default-gtw**

Context config>service>vpls

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

ip

Syntax **ip** *address*
no ip

Context config>service>vpls>mcr-default-gtw

Description This command relates to a system configured for Dual Homing in L2-TPSDA. It defines the IP address used when the system sends out a gratuitous ARP on an active SAP after a ring heals or fails in order to attract traffic from subscribers on the ring with connectivity to that SAP.

Default no ip

Parameters *address* — Specifies the IP address in a.b.c.d. format.

mac

Syntax **mac** *ieee-address*
no mac

Context config>service>vpls>mcr-default-gtw

Description This command relates to a system configured for Dual Homing in L2-TPSDA. It defines the MAC address used when the system sends out a gratuitous ARP on an active SAP after a ring heals or fails in order to attract traffic from subscribers on the ring with connectivity to that SAP.

Default no mac

Parameters *ieee-address* — Specifies the address in xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx format (cannot be all zeros).

move-frequency

Syntax **move-frequency** *frequency*
no move-frequency

Context config>service>vpls>mac-move

Description This objects 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 *frequency* — Specifies the rate, in 5-second intervals for the maximum number of relearns.

Values 1 — 100

retry-timeout

Syntax **retry-timeout** *timeout*
no retry-timeout

Context config>service>vpls>mac-move

Description This objects 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 retry timeout is increased with the provisioned retry timeout in order to avoid thrashing. For example, when retry-timeout is set to 15, it increments (15,30,45,60...).

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

Default 10 (when mac-move is enabled)

Parameters *timeout* — 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** *high-water-mark*

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 *high-water-mark* — Specifies the multicast FIB high watermark as a percentage.

Values 1 — 100

Default 95%

mfib-table-low-wmark

Syntax [**no**] **mfib-table-low-wmark** *low-water-mark*

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 *low-water-mark* — Specifies the multicast FIB low watermark as a percentage.

Values 1 — 100

Default 90%

mfib-table-size

Syntax **mfib-table-size** *size*
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 *mfib-table-size* 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 maxium MFIB table size.

Default none

Parameters *size* — 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

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 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

Description This command enables the context to configure DHCP parameters.

gi-address

Syntax **gi-address** *ip-address* [*src-ip-addr*]
no gi-address

Context config>service>ies>if>dhcp
config>service>vprn>interface>dhcp
config>service>ies>subscriber-interface>grp-if>dhcp
config>service>ies>subscriber-interface>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 that this GI address is to be the source IP address for DHCP relay packets.

lease-populate

Syntax **lease-populate** [*nbr-of-leases*]
lease-populate [*nbr-of-leases*] **l2-header** [**mac** *ieee-address*]
no lease-populate

Context config>subscr-mgmt>msap-policy>vpls-only>dhcp
config>service>vpls>sap>dhcp
config>service>ies>interface>dhcp
config>service>vprn>interface>dhcp
config>service>ies>sub-if>grp-if>dhcp
config>service>vprn>sub-if>grp-if>dhcp
config>service>vprn>sub-if>dhcp

Description This command enables and disables dynamic host DHCPv4 lease state management for SAPs.

For VPLS, DHCP snooping must be explicitly enabled (using the **snoop** command) at all points where DHCP messages requiring snooping enter the VPLS instance (both from the DHCP server and from the subscribers). Lease state information is extracted from snooped DHCP ACK messages to populate lease state table entries for the SAP.

The optional number-of-entries parameter defines the number lease state table entries allowed.

- for this SAP in case of a VPLS service
- for this interface in case of an IES or VPRN interface
- for each SAP in case of an IES or VPRN group-interface
- for this interface in case of an IES or VPRN retail subscriber-interface

If number-of-entries is omitted, only a single entry is allowed. Once the maximum number of entries has been reached, subsequent lease state entries are not allowed and subsequent DHCP ACK messages are discarded.

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

- Populate a SAP based anti-spoof filter table to provide dynamic anti-spoof filtering. If the system is unable to populate the dynamic host information in the anti-spoof filter table on the SAP, the DHCP ACK message must be discarded without adding new lease state entry or updating an existing lease state entry.
- Populate the system's ARP cache based on the arp-populate configuration. Applicable to IES and VPRN interfaces or group-interfaces.
- Populate managed entries into a VPLS forwarding database. VPLS forwarding database population is an implicit feature that automatically places the dynamic host's MAC address into the VPLS FDB. When a dynamic host's MAC address is placed in the lease state table, it will automatically be populated into the VPLS forwarding database associated with the SAP on which the host is learned. The dynamic host MAC address will override any static MAC entries using the same MAC and prevent dynamic learning of the MAC on another interface. Existing static MAC entries with the same MAC address as the dynamic host are marked as inactive but not deleted. If all entries in the lease state table associated with the MAC address are removed, the static MAC may be populated. New static MAC definitions for the VPLS instance may be created while a dynamic host exists associated with the static MAC address
- Generate dynamic ARP replies if **arp-reply-agent** is enabled. Applicable to VPLS service SAPs

Default no lease-populate

Parameters *nbr-of-leases* — Specifies the number of DHCPv4 leases allowed. Note that the operational maximum value may be smaller due to equipped hardware dependencies.

Values 1 — 131071
1 — 500000 (retail subscriber interface)

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. This parameter is only applicable for group-interfaces.

mac — Specifies that the provisioned ieee-address will be used in the anti-spoofing entries for this SAP. The parameter may be changed mid-session. Existing sessions will not be re-programmed unless a tools perform command is issues for the lease. This parameter is only applicable for group-interfaces.

match-circuit-id

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 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 this part of the key is used to guarantee correctness in our lookup. This is really only needed when dealing with an IP aware DSLAM that proxies the client HW mac address.

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 `show>router>interface>detail`)

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** *string*]
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>vprn>if>dhcp
config>service>ies>subscriber-interface>grp-if>dhcp

Description According to RFC 3046, *DHCP Relay Agent Information Option*, 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** *egress-multicast-group-name*
no egress-multicast-group *group-name*

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 *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.

description

Syntax **description** *description-string*
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 *description-string* — 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** *destinations per pass*
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 *destinations per pass* — 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 *destinations per pass* 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 *destinations per pass* 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 *destinations per pass* 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** *ip-filter-id*]
egress-filter [**ipv6** *ipv6-filter-id*]
egress-filter [**mac** *mac-filter-id*]
no egress-filter [**ip** *ip-filter-id*] [**ipv6** *ipv6-filter-id*][**mac** *mac-filter-id*]

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** *ip-filter-id* — Specifies IP filter policy. The filter ID must already exist within the created IP filters.

Values 1 — 65535

ipv6 *ipv6-filter-id* — 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 then 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** *ip-int-name*
no interface *ip-int-name*

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 IP 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.

interface

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

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

address

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

Context config>service>ies>interface
config>service>vprn>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 .

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 adminstratively up state and an address is assigned, it becomes operationally up and the protocol interfaces and the MPLS LSPs associated with that IP interface will be reinitialized.

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

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

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

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

broadcast — The optional **broadcast** parameter overrides the default broadcast address used by the IP interface when sourcing IP broadcasts on the IP interface. If no broadcast format is specified for the IP address, the default value is **host-ones** which indictates 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 *ieee-address* — Specifies the 48-bit MAC address for the static ARP in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee, and ff are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

sap

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

Context config>service>ies>interface
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 router. Each SAP must be unique.

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

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

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

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 —** A SAP is defined within the context of an IP routed interface. Each IP interface is limited to a single SAP definition. Attempts to create a second SAP on an IP interface will fail and generate an error; the original SAP will not be affected.

secondary

Syntax **secondary** {*ip-address/mask* | *ip-address netmask*} [**broadcast all-ones** | **host-ones**] [**igp-inhibit**]
no secondary *ip-address*

Context config>service>ies>interface

This command assigns a secondary IP address/IP subnet/broadcast address format to the interface.

Default none

Parameters *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).

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.

netmask — 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 indictates 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.

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.

address

Syntax **address** {*ip-address/mask* | *ip-address netmask*} [**gw-ip-address** *ip-address*] [**track-srrp** *srrp-inst*] [**holdup-time** *msec*]
no address

Context configure>service>ies>sub-if
configure>service>vprn>sub-if

Description This command will configure the subscriber-interface address along with additional parameters related to multi-chassis redundancy.

Default none

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

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

mask — 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.

netmask — The subnet mask in dotted decimal notation.

Values 0.0.0.0 - 255.255.255.255

track-srrp *srrp-inst* **—** This command will enable the *subscriber-interface* route to track the SRRP state of the specified SRRP instance. The route will update its state attribute to reflect the *state* of SRRP instance:

•Master = srrp-master

•Any other = srrp-non-master

Routing policy can be applied towards the state attribute in order to customize the advertisement of the route. Only one SRRP instance can be tracked per subscriber-interface route. Tracked SRRP instance can be part of the Fate Sharing Group. This command can be enabled at any time.

holdup-time *msec* **—** Time to wait for the route before it accepts the new state attribute. This timer is used to prevent fluctuations in route advertisement caused by short lived SRRP instabilities, in the case that such condition arises.

Values *msec* [100…5000] msec

link-local-address

Syntax **link-local-address** *ipv6-address*
no link-local-address

Context configure>service>ies>sub-if>ipv6
configure>service>vprn>sub-if>ipv6

Description This command will configure the IPv6 Link Local address that will be used as a virtual SRRP IPv6 address by the Master SRRP node. This address will be sent in the Router Advertisements initiated by

the Master SRRP node. Clients will use this address as IPv6 default-gateway. Both SRRP nodes, Master and Backup, must be configured with the same Link Local address.

Default none

Parameters *ipv6-address* — IPv6 address in the form:

x:x:x:x:x:x:x:x:x:x:x:x:d.d.d.d
where:
x - [0..FF]
d - [0..255]

monitor-oper-group

Syntax **monitor-oper-group** *name* **priority-step** *step*
no monitor-oper-group

Context configure>service>ies>sub-if>grp-if>srrp
configure>service>vprn>sub-if>grp-if>srrp

Description This command will configure the association between the SRRP instance in a Fate Sharing Group and the operational-group that contains messaging SAPs. A state transition of a messaging SAP within an operational-group will trigger calculation of the priority for all SRRP instances that are associated with that operational-group.

Default none

Parameters *name* — Name of the operational-group that is tracking operational state of SRRP messaging SAPs

priority-step — The priority step for which the priority of an SRRP instance will be changed. If a messaging SAP within an operational-group transition to a non-UP state, the priority will be decreased by the step value. If the messaging SAP within the operational-group transition into the UP state, the priority of the SRRP instance will be increased by the step value.

step — The step for which the SRRP priority is modified expressed as a decimal integer.

Values 0 — 253

oper-group

Syntax **oper-group** *name*
no oper-group

Context configure>service>epipe>spoke-sdp
configure>service>ies>sub-if>grp-if>sap
configure>service>vprn>sub-if>grp-if>sap

Description This command will configure an epipe or a (messaging) SAP to be the member of the oper-group.

Epipe:

The epipe status can be monitored by the SRRP messaging SAP. The messaging SAP will transition into the DOWN state if the epipe is either DOWN or has a STANDBY status. In order for the

messaging SAP to assume the DOWN state, both RX and TX side of the PW must be shut. In other words, a PW in standby mode also must have the local TX disabled by the virtue of the 'slave' flag (standby-signaling-slave command under the spoke-sdp hierarchy). Without the TX disabled, the SAP monitoring the PW would not transition in the down state. The messaging SAP will be in the UP state if the epipe is in the UP state (Active staus).

(SRRP messaging) SAP:

The state of the messaging SAPs will be monitored by SRRP instances in a Fate Sharing Group. A state change of any of the messaging SAPs defined under the *group-interface* and within the oper-group will trigger recalculation of SRRP priority.

Default none

Parameters *name* — Specify name of the operational-group that contains the member epipe.

srrp-enabled-routing

Syntax **srrp-enabled-routing** [**hold-time** *decisec*]
no srrp-enabled-routing

Context configure>service>ies>sub-if>grp-if
configure>service>vprn>sub-if>grp-if

Description This command will enable SRRP state tracking by managed (IPv4 only) and subscriber-routes. Managed and subscriber-routes that are installed in the Route Table Manager (RTM) would be modified by the source (SRRP would update the route's *state* attribute - srrp-master, srrp-non-master) and this would trigger policy reevaluation with the corresponding action.

Default none

Parameters **hold-time —** Waiting period before which the route's state attribute is updated. The purpose of this command is to avoid propagation of quick successive SRRP state transitions into the routing.

decisec — Specify in deci seconds

Values 1-50

track-srrp

Syntax **track-srrp** *srrp-id*
no track-srrp

Context configure>service>vpls>sap>

Description This is a capture SAP level command. This command is important in PPPoE deployments with MSAPs. PPPoE operation requires that the MAC address learned by the client at the very beginning of the session negotiation phase remains unchanged for the lifetime of the session (RFC 2516). This command will ensure that that the virtual MAC address used during the PPPoE session negotiation phase on the capture SAP is the same virtual MAC address that is used by the SRRP on the group-interface on which the session is established. Therefore, it is mandated that the SRRP instance (and implicitly the group-interface) where the session belongs to is known in advance. If the group-

interface name for the session is returned by the RADIUS, it must be ensured that this group-interface is the one on which the tracked SRRP instance is configured. PPPoE sessions on the same capture SAP cannot be shared across multiple group-interfaces, but instead they all must belong to a single group-interface that is known in advance.

The same restrictions will apply to IPoE clients in MC Redundancy scenario if they are to be supported concurrently on the same capture SAP as PPPoE.

The supported capture SAP syntax is this:

sap <port-id>:X.* capture-sap

The capture SAP syntax that is NOT supported is this:

sap <port-id>:*.* capture-sap

Default none

Parameters *srrp-id* — Specify SRRP instance number.

Values 1..4294967295

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 *srrp-id* — 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** [*service-id*] **interface** *interface-name* **dst-ip** *ip-address*

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

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 *service-id* — 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 *interface-name* **—** Specifies the name of the interface running BFD.

dst-ip *ip-address* **—** Specifies the destination address of the interface running BFD.

gw-mac

Syntax **gw-mac** *mac-address*
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 *mac-address* — 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** *interval*
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 *interval* — 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** *sap-id*
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 *sap-id* — Specifies the physical port identifier portion of the SAP definition. See Common Service Commands on page 1510 for *sap-id* command syntax.

one-garp-per-sap

Syntax **[no] one-garp-per-sap**

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

Description This command is applicable to PPPoE only deployments in which there are multiple subnets under the subscriber-interface. In such case, if the switchover occurs, it will be sufficient to send a single Gratuitous ARP on every VLAN to update the Layer 2 forwarding path in the access aggregation network. This single gratuitous ARP will contain the IP address of the first gw-adddress found under the subscriber-interface address.

prefix

Syntax **prefix** *ipv6-address/prefix-length* **[pd] [wan-host] track-srrp** *srrp-instance* **[holdup-time** *milli-seconds*]
no prefix

Context configure>service>vprn>sub-if>ipv6>sub-pfx
configure>service>ies>sub-if>ipv6>sub-pfx

Description This command will configure the IPv6 subscriber-interface address along with additional parameters related to multi-chassis redundancy.

Default none

Parameters *ip-address* — The IP address of the IP interface. The *ip-address* portion of the **address** command specifies the IP host address that will be used by the IP interface within the subnet. This address

must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).

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

mask — 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.

netmask — The subnet mask in dotted decimal notation.

Values 0.0.0.0 - 255.255.255.255

track-srrp *srrp-inst* **—** This command will enable the subscriber-interface IPv6 route to track the SRRP state of the specified SRRP instance. The route will update its state attribute to reflect the state of SRRP instance:

Master = srrp-master

Any other = srrp-non-master

Routing policy can be applied towards the state attribute in order to customize the advertisement of the route. Only one SRRP instance can be tracked per subscriber-interface route. Tracked SRRP instance can be part of the Fate Sharing Group. This command can be enabled at any time.

holdup-time *msec* **—** Time to wait for the route before it accepts the new state attribute. This timer is used to prevent fluctuations in route advertisement caused by short lived SRRP instabilities, in the case that such condition arises.

msec [100…5000] msec

policy

Syntax [**no**] **policy** *vrrp-policy-id*

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 vrrp-policy-id from the SRRP instance.

Parameters *vrrp-policy-id* — Specifies one or more VRRP policies with the SRRP instance.

Values 1 — 9999

preempt

Syntax [**no**] **preempt**

Context configure>service>ies>sub-if>grp-if>srrp
configure>service>vprn>sub-if>grp-if>srrp

Description When preempt is enabled, a newly initiated SRRP instance can overrides an existing Master SRRP instance if its priority value is higher than the priority of the current Master.

If preempt is disabled, an SRRP instance only becomes Master if the master down timer expires before an SRRP advertisement message is received from the adjacent SRRP enabled node.

Default preempt

priority

Syntax **priority** *priority*
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 *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 *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 *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

host-connectivity-verify

Syntax **host-connectivity-verify** [**source** {**vrrp** | **interface**}] [**interval** *interval*] [**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 ***interval*** **—** 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.

ip-mtu

Syntax **ip-mtu** *octets*
no ip-mtu

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

Description This command specifies the maximum size of IP packets on this group-interface. Packets larger than this will be fragmented.

The ip-mtu applies to all IPoE host types (DHCP, ARP, static). For PPP/L2TP sessions, the ip-mtu is not taken into account for the mtu negotiation; the ppp-mtu in the ppp-policy should be used instead.

The **no** form of the command removes the octets value from the configuration.

Default none

Parameters *octets* — Specifies the largest frame size (in octets) that this interface can handle.

Values 512 — 9000

private-retail-subnets

Syntax [**no**] **private-retail-subnets**

Context config>service>vprn>sub-if

Description This command controls the export of retail subnets and prefixes to the wholesale forwarding service. When this attribute is configured, subnets and prefixes configured on the 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.

Configuring private retail subnets is not supported for IPoEv4 host management (DHCPv4, IPv4 static-host and ARP-host). If PPPoE sessions need to coexist with IPoEv4 hosts, then this attribute should not be configured on the retail subscriber 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. In this case, private retail subnets will be enabled by default and it will not be possible to deconfigure it.

unnumbered

Syntax **unnumbered** [*ip-int-name*|*ip-address*]
no unnumbered

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

Description This command can be configured only for subscriber-interfaces that do not have an IPv4 address explicitly configured and is therefore operationally in a DOWN state. By configuring this command, the subscriber-interface will borrow the IPv4 address from the referenced interface (directly or indirectly via IP address) that must be operationally UP and located in the same routing instance as the subscriber-interface. This will allow the subscriber-interface to becomes operationally UP and consequently allow forwarding of the subscriber traffic.

Such interface is referred as unnumbered interface, since it does not have explicitly configured a unique IP address. Subscriber-hosts under the unnumbered subscriber-interface are installed in the fib as /32 hosts.

Without this command the subscriber-interface is operationally DOWN and subscriber-host instantiation is not possible.

This command is mutually exclusive with the allow-unmatched-subnets command under the same CLI hierarchy.

The operation of IPv6 host is not affected by this command.

Default no unnumbered

Parameters *ip-int-name* — Specifies the interface name from which an IPv4 address will be borrowed.

ip-address — The IP address from an optionationally up interface that will be used for subscriber interface.

ipv6

Syntax **ipv6**

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

Description This command enables the context to enable IPv6 IPoE bridged mode.

ipoe-bridged-mode

Syntax [**no**] **ipoe-bridged-mode**

Context config>service>ies>sub-if>grp-if>ipv6
config>service>vprn>sub-if>grp-if>ipv6

Description This command enters the context to enable IPv6 IPoE bridged mode.

The **no** form of the command disables the IPv6 IPoE bridged mode.

allow-multiple-wan-addresses

Syntax [**no**] **allow-multiple-wan-addresses**

Context config>service>ies>sub-if>grp-if>ipv6
config>service>vprn>sub-if>grp-if>ipv6

Description This command enables host to have two WAN addresses, one from DHCP IA_NA and one from SLAAC assignment.

Default no allow-multiple-wan-addresses

nd

Syntax **nd**

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

Description This command enables the context to configure neighbor discovery parameters.

dad-snooping

Syntax [**no**] **dad-snooping**

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

Description This command allows the router to populate the neighbor discovery table through snooping subscribers' duplicate address detection messages.

Default no dad-snooping

neighbor-limit

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

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

Description This command configures the maximum number of neighbors learned.

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

router-advertisements

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

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

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

current-hop-limit

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

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

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

Default 64

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

Values 0 — 255

dns-options

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

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

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

include-dns

Syntax [**no**] **include-dns**

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

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

Default no include-dns

rdnss-lifetime

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

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

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

Default rdnss-lifetime 3600

Parameters *seconds* — Specifies the time in seconds.

Values 900 — 3600

infinite — The RDNSS address can be used permanently.

force-mcast

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

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

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

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

managed-configuration

Syntax [**no**] **managed-configuration**

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

Description This command sets or resets managed address configuration flag for this group-interface.

max-advertisement

Syntax **max-advertisement** *seconds*
no max-advertisement

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

Description This command specifies the maximum time allowed between sending unsolicited router advertisements from this interface.

Parameters *seconds* — Specifies the maximum advertisement interval in seconds for this group-interface.

Values 900 — 11800

min-advertisement

Syntax **min-advertisement** *seconds*
no min-advertisement

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

Description This command specifies the minimum time allowed between sending unsolicited router advertisements from this interface.

Parameters *seconds* — Specifies the minimum advertisement interval in seconds for this group-interface.

Values 900 — 1350

mtu

Syntax **mtu** *bytes*
no mtu

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

Description This command specifies the value to be placed in link MTU options sent by the router on this interface.

Parameters *bytes* — Sets the advertised MTU value in bytes for this group-interface.

Values 1280 — 9212

other-stateful-configuration

Syntax [**no**] **other-stateful-configuration**

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

Description This command sets and resets the other-stateful-configuration flag for this group-interface.

prefix-options

Syntax [**no**] **prefix-options**

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

Description This command enables the context to configure prefix options for this group-interface.

autonomous

Syntax [**no**] **autonomous**

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

Description This command enables or disables the option that determines whether or not the prefix can be used for stateless address autoconfiguration.

on-link

Syntax [**no**] **on-link**

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

Description This command specifies whether the prefix will be assigned to an interface on the specified link.

preferred-lifetime

Syntax **preferred-lifetime** *seconds*
preferred-lifetime infinite
no preferred-lifetime

reachable-time

Syntax **reachable-time** *milli-seconds*
no reachable-time

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

Description This command configures the value to be placed in the reachable time field in router advertisement messages sent from this interface.

Default 0

Parameters *milli-seconds* — Specifies the reachable time in milli-seconds for advertisements from this group-interface.

Values 0 — 3600000

retransmit-time

Syntax **retransmit-time** *milli-seconds*
no retransmit-time

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

Description This command configures the value to be placed in the retransmit timer field in router advertisements sent from this interface.

Default 0

Parameters *milli-seconds* — Specifies the retransmit time in milli-seconds for advertisement from this group-interface.

Values 0 — 1800000

router-lifetime

Syntax **router-lifetime** *seconds*
router-lifetime no-default-router
no router-lifetime

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

Description This command configures the value to be placed in the router lifetime field of router advertisements sent from this interface.

Default 4500

Parameters *seconds* — Specifies the router lifetime in seconds for this group-interface.

Values 2700 — 9000

no-default-router — Specifies that the router is not to be used as a default router.

router-solicit

Syntax **router-solicit**

Context config>service>ies>sub-if>grp-if>ipv6
config>service>vprn>sub-if>grp-if>ipv6
config>service>vprn>sub-if>ipv6
config>service>ies>sub-if>ipv6

Description This command enables the context to configure parameters used for router-solicit based authentication.

inactivity-timer

Syntax **inactivity-timer** [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]
no inactivity-timer

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

Description This command specifies the time before an inactive host is removed.

Default no interactive-timer

Parameters **infinite** — An idle host is never removed.

days *days* — An idle host is removed if idle within the number of specified days.

hrs *hours* — An idle host is removed if idle within the number of specified hours.

min *minutes* — An idle host is removed if idle within the number of specified minutes.

sec *seconds* — An idle host is removed if idle within the number of specified seconds.

min-auth-interval

Syntax **min-auth-interval** [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]
no min-auth-interval

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

Description This command specify the minimum interval between two consecutive authentication attempts from the same host.

Default no min-auth-interval

Parameters **days** *days* — The number of days that a user must wait for the next authentication attempt.

hrs *hours* — The number of hours that a user must wait for the next authentication attempt.

min *minutes* — The number of minutes that a user must wait for the next authentication attempt.

sec *seconds* — The number of seconds that a user must wait for the next authentication attempt.

user-db

Syntax [**no**] **user-db**

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

Description This command enables the use of the local-user-database for authentication.

Default no user-db

Parameters *local-user-db-name* — Specifies the name of the local-user-database to authenticate the router-solicit. The local-user-database can also return a static prefix or a pool name for address assignment.

shutdown

Syntax [**no**] **shutdown**

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

Description This command enables or disables SLAAC triggered host creation.

Default no shutdown

local-proxy-arp

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

Context config>service>ies>interface
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.

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 *number* and *seconds* 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. (*Default: redirects 100 10*)

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 *number* — The maximum number of ICMP redirect messages to send. This parameter must be specified with the *second* parameter.

Values 10 — 1000

seconds — The time frame in seconds used to limit the *number* of ICMP redirect messages that can be issued.

Values 1 — 60

remote-proxy-arp

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

Context config>service>ies>interface
config>service>vprn>interface
config>service>ies>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

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 will appear in the core routing ARP table. A static ARP can only be configured if it exists on the network attached to the IP interface.

If an entry for a particular IP address already exists and a new MAC address is configured for the IP address, the existing MAC address will be replaced with the new MAC address.

The **no** form of this command removes a static ARP entry.

Default none

Parameters *ip-address* — Specifies the IP address for the static ARP in IP address dotted decimal notation.

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

ttl-expired

Syntax **ttl-expired** *number seconds*
no ttl-expired

Context config>service>ies>if>icmp

Description This command configures the rate Internet Control Message Protocol. (ICMP) TTL expired messages are issued by the IP interface.
By default, generation of ICMP TTL expired messages is enabled at a maximum rate of 100 per 10 second time interval.

The **no** form of this command disables the limiting the rate of TTL expired messages on the router interface.

Default ttl-expired 100 10

Parameters *number* — The maximum number of ICMP TTL expired messages to send, expressed as a decimal integer. This parameter must be specified with the *seconds* parameter.

Values 10 — 1000

seconds — The time frame in seconds used to limit the *number* of ICMP TTL expired messages that can be issued, expressed as a decimal integer.

Values 1 — 60

unreachables

Syntax **unreachables** [*number seconds*]
no unreachables

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 *number* and *seconds* 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 *number* — The maximum number of ICMP unreachable messages to send. This parameter must be specified with the *seconds* parameter.

Values 10 — 1000

seconds — The time frame in seconds used to limit the *number* of ICMP unreachable messages that can be issued.

Values 1 — 60

Interface IPv6 Commands

ipv6

Syntax [**no**] **ipv6**

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

Description This command enables the context to configure IPv6 for an interface.

address

Syntax **address** *ipv6-address/prefix-length* [**eui-64**]
no address *ipv6-address/prefix-length*

Context config>service>ies>if>ipv6

Description This command assigns an IPv6 address to the IES interface.

Parameters *ipv6-address/prefix-length* — Specify the IPv6 address on the interface.

Values		
	ipv6-address	x:x:x:x:x:x:x:x (eight 16-bit pieces) x:x:x:x:x:x:d.d.d.d x [0 — FFFF]H d [0 — 255]D
	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.

default-dns

Syntax **default-dns** *ipv6-address* [**secondary** *ipv6-address*]
no default-dns

Context config>service>ies>sub-if>ipv6
config>service>vprn>sub-if>ipv6

Description Configure last resort IPv6 DNS addresses that can be used for name resolution by IPoEv6 hosts (IA_NA, IA_PD and SLAAC) and PPPoEv6 hosts (IA_NA, IA_PD and SLAAC).

Default no default-dns

Parameters *ipv6-addres* — s - IPv6 address of the primary DNS server.

secondary *ipv6-address* — IPv6 address of the secondary DNS server (optional).

address

Syntax **address** *ipv6-address/prefix-length* [**pd**] [**wan-host**] [**track-srrp** *srrp-instance*] [**holdup-time** *milli-seconds*]
no address *ipv6-address/prefix-length*

Context config>service>ies>sub-if>ipv6
config>service>vprn>sub-if>ipv6

Description This command assigns an IPv6 address/subnet to the subscriber interface. Note that SRRP is not supported for IPv6 subscriber interface.

Default none

Parameters *ipv6-address* — Specifies the 128-bit IPv6 address.

Values 128-bit hexadecimal IPv6 address in compressed forM.

prefix-length — Specifies the length of any associated aggregate prefix.

Values 32-63

pd — Specifies that this aggregate is used by IPv6 ESM hosts for DHCPv6 prefix-delegation.

wan-host — Specifies that this aggregate is used by IPv6 ESM hosts for DHCPv6 addresses or by a routing gateway's WAN interface.

allow-unmatching-prefixes

Syntax [**no**] **allow-unmatching-prefixes**

Context configure>service>vprn>sub-if>ipv6
configure>service>ies>sub-if>ipv6

Description This command allows address assignments for IPoEv6 and PPPoEv6 hosts in cases where the subscriber host assigned IPv6 address or prefix falls outside of the subscriber-prefix range explicitly configured for the subscriber-interface (**configure>service>vprn/ies>sub-if>ipv6**> or the subscriber-prefix is not configured at all.

SLAAC hosts will be installed in the FIB as /64 entries, the length of the installed DHCP-PD prefix will be dictated by the prefix-length and the DHCP-NA host will be installed as /128 entries.

IPv4 subscriber hosts are unaffected by this command.

Default no allow-unmatching-prefixes

auto-reply

Syntax [**no**] **auto-reply**

Context configure>service>vprn>sub-if>ipv6
configure>service>ies>sub-if>ipv6

Description This command assists IP-only static hosts to resolve their default gateway and MAC. By default, the BNG anti-spoof filter drops packets from unknown hosts. The auto-reply features first allow hosts to resolve their default gateway and afterwards allow them to forward traffic. Using the data traffic, the BNG can utilize the data-trigger mechanism to learn the host's MAC and populate the full IP+MAC static host entry.

Default no auto-reply

neighbor-solicitation

Syntax [**no**] **neighbor-solicitation**

Context configure>service>vprn>sub-if>ipv6>auto-reply
configure>service>ies>sub-if>ipv6>auto-reply

Description This command enables auto-reply for neighbor solitation.

The **no** form of the command disables auto-reply.

router-solicitation

Syntax [**no**] **router-solicitation**

Context configure>service>vprn>sub-if>ipv6>auto-reply
configure>service>ies>sub-if>ipv6>auto-reply

Description This command enables auto-reply router solitation.

The **no** form of the command disables router-solitcitation.

delegated-prefix-length

Syntax **delegated-prefix-length** *bits*
delegated-prefix-length variable
no delegated-prefix-length

Context configure>service>vprn>sub-if>ipv6
configure>service>ies>sub-if>ipv6

Description This command configures the subscriber-interface level setting for delegated prefix length. The delegated prefix length for a subscriber- interface can be either set to a fixed value that is explicitly configured under the subscriber-interface CLI hierarchy or a variable value that can be obtained from various sources. This command can be changed only when no IPv6 prefixes are configured under the subscriber-interface.

Default no delegated-prefix-length This means that the delegated prefix length is 64.

Parameters *bits* — The delegated prefix length in bits. This value will beapplicable to the entire subscriber-interface. In case that the delegated prefix length is also supplied via other means (LUDB,

RADIUS or DHCP Server), such supplied value must match the value configured under the subscriber-interface. Otherwise the prefix instantiation in 7x50 will fail.

Values 48 — 64

variable — The delegated prefix value can be of any length between 48..64. The value itself can vary between the prefixes and it will be provided at the time of prefix instantiation. The order of priority for the source of the delegated prefix length is:

- LUDB
- RADIUS
- DHCPv6 server

subscriber-prefixes

Syntax [**no**] **subscriber-prefixes**

Context config>services>ies>sub-if>ipv6
config>services>vprn>sub-if>ipv6

Description This command specifies aggregate off-link subscriber prefixes associated with this subscriber interface. Individual prefixes are specified under the prefix context list aggregate routes in which the next-hop is indirect via the subscriber interface.

prefix

Syntax **prefix** *ipv6-address/prefix-length* [**pd**] [**wan-host**]
no prefix *ipv6-address/prefix-length*

Context config>services>ies>sub-if>ipv6>sub-prefixes
config>services>vprn>sub-if>ipv6>sub-prefixes

Description This command allows a list of prefixes(using the prefix command multiple times) to be routed to hosts associated with this subscriber interface. Each prefix will be represented in the associated FIB with a reference to the subscriber interface. Prefixes are defined as being for prefix delegation (pd) or use on a WAN interface or host (wan-host).

Parameters *ipv6-address* — Specifies the 128-bit IPv6 address.

Values 128-bit hexadecimal IPv6 address in compressed form.

prefix-length — Specifies the length of any associated aggregate prefix.

Values 32-63

pd — Specifies that this aggregate is used by IPv6 ESM hosts for DHCPv6 prefix-delegation.

wan-host — Specifies that this aggregate is used by IPv6 ESM hosts for local addressing or by a routing gateway's WAN interface.

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 DHCP6 relay.

lease-populate

Syntax **lease-populate** [*nbr-of-leases*]
lease-populate [*nbr-of-leases*] **route-populate** [**pd**] **na** [**ta**]
lease-populate [*nbr-of-leases*] **route-populate pd** [**na**] [**ta**] [**exclude**]
lease-populate [*nbr-of-leases*] **route-populate** [**pd**] [**na**] **ta**
no lease-populate

Context config>service>ies>if>ipv6>dhcp6-relay
config>service>vprn>if>ipv6>dhcp6-relay

Description This command specifies the maximum number of DHCPv6 lease states allocated by the DHCPv6 relay function, allowed on this interface.

Optionally, by specifying "route-populate" parameter, system could:

- Create routes based on the IA_PD/IA_NA/IA_TA prefix option in relay-reply message.
- Create black hole routes based on OPTION_PD_EXCLUDE in IA_PD in relay-reply message.

These routes could be redistributed into IGP/BGP by using route-policy, following protocol types that could be used in "from protocol":

- dhcpv6-pd
- dhcpv6-na
- dhcpv6-ta
- dhcpv6-pd-excl

The **no** form of the command disables dynamic host lease state management.

Default no lease-populate

Parameters *nbr-of-leases* — 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

route-populate — Create routes based on options in relay-reply messages.

Values pd/na/ta — Create route based on specified option.
exclude — Create blackhole route based on OPTION_PD_EXCLUDE

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.

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.

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

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.

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

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** [*number seconds*]
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 *number* — Specifies the number of “packet-too-big” ICMP messages to send in the time frame specified by the *seconds* parameter.

Values 10 — 1000

Default 100

seconds — Specifies the time frame in seconds that is used to limit the number of “packet-too-big” ICMP messages issued.

Values 1 — 60

Default 10

param-problem

Syntax **param-problem** [*number seconds*]
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.

Default 100 10

number — Specifies the number of "parameter-problem" ICMP messages to send in the time frame specified by the *seconds* parameter.

Values 10 — 1000

Default 100

seconds — Specifies the time frame in seconds that is used to limit the number of "parameter-problem" ICMP messages issued.

Values 1 — 60

Default 10

redirects

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

Context config>service>ies>if>ipv6>icmp6

Description This command configures ICMPv6 redirect messages. When enabled, ICMPv6 redirects are generated when routes are not optimal on this router and another router on the same subnetwork has a better route in order to alert that node that a better route is available.

When disabled, ICMPv6 redirects are not generated.

Default 100 10

number — Specifies the number of version 6 redirects are to be issued in the time frame specified by the *seconds* parameter.

Values 10 — 1000

Default 100

seconds — Specifies the time frame in seconds that is used to limit the number of version 6 redirects issued.

Values 1 — 60

Default 10

time-exceeded

Syntax **time-exceeded** [*number seconds*]
no time-exceeded

Context config>service>ies>if>ipv6>icmp6

Description This command specifies whether "time-exceeded" ICMP messages should be sent. When enabled, ICMPv6 "time-exceeded" messages are generated by this interface.

When disabled, ICMPv6 "time-exceeded" messages are not sent.

Default 100 10

number — Specifies the number of "time-exceeded" ICMP messages are to be issued in the time frame specified by the *seconds* parameter.

Values 10 — 1000

Default 100

seconds — Specifies the time frame in seconds that is used to limit the number of "time-exceeded" ICMP message to be issued.

Values 1 — 60

Default 10

unreachables

Syntax **unreachables** [*number seconds*]
no unreachables

Context config>service>ies>if>ipv6>icmp6

Description This command specifies that ICMPv6 host and network unreachable messages are generated by this interface.

When disabled, ICMPv6 host and network unreachable messages are not sent.

Default 100 10

number — Specifies the number of destination unreachable ICMPv6 messages are issued in the time frame specified by the *seconds* parameter.

Values 10 — 1000

Default 100

seconds — Specifies the time frame in seconds that is used to limit the number of destination unreachable ICMPv6 messages to be issued.

Values 1 — 60

Default 10

local-proxy-nd

Syntax [**no**] **local-proxy-nd**

Context config>service>ies>if>ipv6

Description When this command is enabled, the interface will reply to neighbor solicitation requests when both the hosts are on the same subnet. In this case, ICMP redirects will be disabled. When this command is disabled, the interface will not reply to neighbor solicitation requests if both the hosts are on the same subnet.

Default disabled

neighbor

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: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 This command 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 MACs can be re-learned in this service, before the SAP where the moving MAC was last seen is automatically disabled in order to protect the system against undetected loops or duplicate MACs. The rate is computed as the maximum number of re-learns allowed in a 5 second interval: for example, the default rate of 10 re-learns per second corresponds to 50 re-learns in a 5 second period.

Label	Description (Continued)
Mac Move Timeout	Indicates the time in seconds to wait before a SAP that has been disabled after exceeding the maximum re-learn rate is re-enabled. A value of zero indicates that the SAP will not be automatically re-enabled after being disabled. If after the SAP is re-enabled it is disabled again, the effective retry timeout is doubled in order to avoid thrashing.
Table Size	The maximum number of learned and static entries allowed in the FDB.
Total Count	The current number of entries (both learned and static) in the FDB.
Learned Count	The current number of learned entries in the FDB.
Static Count	The current number of static entries in the FDB.
Remote Age	The number of seconds used to age out FDB entries learned on an SDP. These entries correspond to MAC addresses learned on remote SAPs.
Local Age	The number of seconds used to age out FDB entries learned on local SAPs.
High WaterMark	The utilization of the FDB table of this service at which a 'table full' alarm is raised by the agent.
Low WaterMark	The utilization of the FDB table of this service at which a 'table empty' alarm is raised by the agent.
Mac Learning	Specifies whether the MAC learning process is enabled in this service.
Discard Unknown	Specifies whether frames received with an unknown destination MAC are discarded in this service.
MAC Aging	Specifies whether the MAC aging process is enabled in this service.
MAC Pinning	Specifies whether MAC Pinning is enabled in this service.
Relearn Only	When enabled, indicates that either the FDB table of this service is full or that the maximum system-wide number of MACs supported by the agent has been reached, and thus MAC learning is temporary disabled, and only MAC re-learns can take place.
Total Service FDBs	The current number of service FDBs configured on this node.
Total FDB Size	The sum of configured FDBs.
Total FDB Entries In Use	The total number of entries (both learned and static) in use.

Sample Output

```
A:ALA-12# show service fdb-info
===============================================================================
Forwarding Database (FDB) Information
===============================================================================
Service Id        : 700                 Mac Move          : Disabled
Mac Move Rate     : 10                  Mac Move Timeout  : 10
Table Size        : 250                 Total Count       : 0
Learned Count     : 0                   Static Count      : 0
Remote Age        : 900                 Local Age         : 300
High WaterMark    : 95%                 Low Watermark     : 90%
Mac Aging         : Enabl               Relearn Only      : False

Service Id        : 725                 Mac Move          : Disabled
Mac Move Rate     : 10                  Mac Move Timeout  : 10
Table Size        : 250                 Total Count       : 0
Learned Count     : 0                   Static Count      : 0
Remote Age        : 900                 Local Age         : 300
High WaterMark    : 95%                 Low Watermark     : 90%
Mac Learning      : Enabl               Discard Unknown   : Dsabl
Mac Aging         : Enabl               Relearn Only      : False

Service Id        : 740                 Mac Move          : Disabled
Mac Move Rate     : 10                  Mac Move Timeout  : 10
Table Size        : 250                 Total Count       : 0
Learned Count     : 0                   Static Count      : 0
Remote Age        : 900                 Local Age         : 300
High WaterMark    : 95%                 Low Watermark     : 90%
Mac Learning      : Enabl               Discard Unknown   : Dsabl
Mac Aging         : Enabl               Relearn Only      : False
...
-------------------------------------------------------------------------------
Total Service FDBs : 7
Total FDB Configured Size : 1750
Total FDB Entries In Use  : 0
-------------------------------------------------------------------------------
===============================================================================
*A:ALA-48#
```

fdb-mac

Syntax **fdb-mac** *ieee-address*

Context show>service

Description Displays the FDB entry for a given MAC address.

Parameters *ieee-address* — The 48-bit MAC address for which to display the FDB entry 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.

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** *ingress-label1* [*ingress-label2*]

Context show>service

Description Display services using the range of ingress labels.

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

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

Use the **show router** *vprn-service-id* **ldp bindings** command to display dynamic labels.

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

Values 0, 2048 —131071

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

Default The *ingress-label1* 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 1510 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** [*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.

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** [*sdp-id*[:*vc-id*] | far-end *ip-address*]

Context show>service

Description Display services using SDP or far-end address options.

Parameters *sdp-id* — Displays only services bound to the specified SDP ID.

Values 1 — 17407

vc-id — The virtual circuit identifier.

Values 1 — 4294967295

far-end *ip-address* — Displays only services matching with the specified far-end IP address.

Default Services with any far-end IP address.

Output **Show Service SDP Using X —** The following table describes **sdp-using** output fields.

Label	Description
Svc ID	The value identifying a service.

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 This command displays the services matching certain usage properties.

If no optional parameters are specified, all services defined on the system are displayed.

Parameters **ies** — Displays matching IES instances.

vpls — Displays matching VPLS instances.

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 detail**
active-subscribers mirror
active-subscribers summary

Context show>service

Description This command displays active subscriber information.

Parameters **summary —** Displays active subscriber information in a brief format.

detail — Displays detailed output.

Sample Output

```
*A:Dut-C# show service active-subscribers
===============================================================================
```

```
Active Subscribers
===============================================================================
-------------------------------------------------------------------------------
Subscriber hpolSub1 (hpolSubProf2)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
(1) SLA Profile Instance sap:lag-1:2000.1 - sla:hpolSlaProf1
-------------------------------------------------------------------------------
IP Address
                MAC Address       PPPoE-SID Origin
-------------------------------------------------------
200.1.4.194
                00:01:00:00:03:c1 1         IPCP
-------------------------------------------------------------------------------
(2) SLA Profile Instance sap:lag-1:2000.1 - sla:hpolSlaProf2
-------------------------------------------------------------------------------
IP Address
                MAC Address       PPPoE-SID Origin
-------------------------------------------------------
200.1.4.35
                00:01:00:00:03:22 N/A       ARP-Host
200.1.4.195
                00:01:00:00:03:c2 1         IPCP
-------------------------------------------------------------------------------
Subscriber hpolSub16 (hpolSubProf1)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
(1) SLA Profile Instance sap:[lag-1:2000.2] - sla:hpolSlaProf1
-------------------------------------------------------------------------------
IP Address
                MAC Address       PPPoE-SID Origin
-------------------------------------------------------
200.1.4.224
                00:01:00:00:03:df 1         IPCP
-------------------------------------------------------------------------------
Subscriber hpolSub2 (hpolSubProf1)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
(1) SLA Profile Instance sap:lag-1:2000.1 - sla:hpolSlaProf1
-------------------------------------------------------------------------------
IP Address
                MAC Address       PPPoE-SID Origin
-------------------------------------------------------
200.1.4.196
                00:01:00:00:03:c3 1         IPCP
-------------------------------------------------------------------------------
Number of active subscribers : 3
===============================================================================
*A:Dut-C#
```

credit-control

Syntax **credit-control credit-control** [**subscriber** *sub-ident-string*]
credit-control out-of-credit [**action** *action*] [**summary**]

Context show>service>active-subscribers

Description This command displays active subscriber credit control information.

filter

Syntax **filte**r [**subscriber** *sub-ident-string*] [**origin** *origin*]

Context show>service>active-subscribers

Description This command displays active subscriber filter information.

hierarchy

Syntax **hierarchy** [**subscriber** *sub-ident-string*]

Context show>service>active-subscribers

Description This command displays active subscriber hierarchy information.

host-tracking

Syntax **host-tracking** [**subscriber** *sub-ident-string*]
host-tracking [**subscriber** *sub-ident-string*] **detail**
host-tracking [**subscriber** *sub-ident-string*] **summary**
host-tracking [**subscriber** *sub-ident-string*] **statistics**

Context show>service>active-subscribers

Description This command displays active subscriber host tracking information.

groups

Syntax **groups** [**group** *group-ip-address*]
groups group *group-ip-address*] **detail**
groups group *group-ip-address*] **summary**

Context show>service>active-subscribers

Description This command displays active subscriber host tracking groups information.

igmp

Syntax **igmp** [**subscriber** *sub-ident-string*][**detail**]

Context show>service>active-subscribers

Description This command displays active subscriber IGMP information.

Sample Output

```
*B:Dut-C# show service active-subscribers igmp
===============================================================================
Active Subscribers
===============================================================================
Subscriber                      IGMP-Policy
  HostAddr                      GrpItf                            NumGroups
-------------------------------------------------------------------------------
sub_1                           pol1
  112.112.1.1                   gi_1_1                            1
  112.112.1.2                   gi_1_1                            2
sub_2                           pol1
  112.112.1.3                   gi_1_2                            0
-------------------------------------------------------------------------------
Number of Subscribers : 2
===============================================================================
*B:Dut-C#


*B:Dut-C# show service active-subscribers igmp detail
===============================================================================
Active Subscribers Detail
===============================================================================
Subscriber                      IGMP-Policy
  HostAddr                      GrpItf                            NumGroups
    GrpAddr                     Type            Up-Time           Mode
      SrcAddr                   Type                              Blk/Fwd
-------------------------------------------------------------------------------
sub_1                           pol1
  112.112.1.1                   gi_1_1                            1
    225.0.0.1                   Dynamic         0d 00:00:53       Include
      11.11.0.1                 Dynamic                           Fwd
      11.11.0.2                 Dynamic                           Fwd
  112.112.1.2                   gi_1_1                            2
    225.0.0.1                   Dynamic         0d 00:00:44       Exclude
      11.11.0.1                 Dynamic                           Blk
    225.0.0.2                   Dynamic         0d 00:00:44       Exclude
      11.11.0.1                 Dynamic                           Blk
sub_2                           pol1
  112.112.1.3                   gi_1_2                            0
-------------------------------------------------------------------------------
Number of Subscribers : 2
===============================================================================
*B:Dut-C#

*B:Dut-C# show service active-subscribers igmp subscriber "sub_1" detail
===============================================================================
Active Subscribers Detail
===============================================================================
Subscriber                      IGMP-Policy
```

```
  HostAddr                        GrpItf                          NumGroups
    GrpAddr                       Type            Up-Time         Mode
      SrcAddr                     Type                            Blk/Fwd
-------------------------------------------------------------------------------
sub_1                             pol1
  112.112.1.1                     gi_1_1                          1
    225.0.0.1                     Dynamic         0d 00:01:04     Include
      11.11.0.1                   Dynamic                         Fwd
      11.11.0.2                   Dynamic                         Fwd
  112.112.1.2                     gi_1_1                          2
    225.0.0.1                     Dynamic         0d 00:00:55     Exclude
      11.11.0.1                   Dynamic                         Blk
    225.0.0.2                     Dynamic         0d 00:00:55     Exclude
      11.11.0.1                   Dynamic                         Blk
-------------------------------------------------------------------------------
Number of Subscribers : 1
===============================================================================
B:Dut-C#
```

subscriber

Syntax **subscriber** *sub-ident-string*
subscriber *sub-ident-string* **detail**
subscriber *sub-ident-string* **mirror**
subscriber *sub-ident-string* **sap** *sap-id* **sla-profile** *sla-profile-name*
subscriber *sub-ident-string* **sap** *sap-id* **sla-profile** *sla-profile-name* **detail**
subscriber *sub-ident-string* **sap** *sap-id* **sla-profile** *sla-profile-name* **mirror**

Context show>service>active-subscribers

Description This command displays active subscriber information for a subscriber.

Sample Output

```
*A:Dut-C# show service active-subscribers subscriber "hpolSub1"
===============================================================================
Active Subscribers
===============================================================================
-------------------------------------------------------------------------------
Subscriber hpolSub1 (hpolSubProf2)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
(1) SLA Profile Instance sap:lag-1:2000.1 - sla:hpolSlaProf1
-------------------------------------------------------------------------------
IP Address
                MAC Address       PPPoE-SID Origin
-------------------------------------------------------
200.1.4.194
                00:01:00:00:03:c1 1         IPCP
-------------------------------------------------------------------------------
(2) SLA Profile Instance sap:lag-1:2000.1 - sla:hpolSlaProf2
-------------------------------------------------------------------------------
IP Address
                MAC Address       PPPoE-SID Origin
-------------------------------------------------------
200.1.4.35
```

```
                00:01:00:00:03:22 N/A       ARP-Host
200.1.4.195
                00:01:00:00:03:c2 1         IPCP
===============================================================================
*A:Dut-C#


*A:Dut-C# show service active-subscribers subscriber "hpolSub1"
sap lag-1:2000.1 sla-profile "hpolSlaProf2"
===============================================================================
Active Subscribers
===============================================================================
-------------------------------------------------------------------------------
Subscriber hpolSub1 (hpolSubProf2)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
(1) SLA Profile Instance sap:lag-1:2000.1 - sla:hpolSlaProf2
-------------------------------------------------------------------------------
IP Address
                MAC Address       PPPoE-SID Origin
-------------------------------------------------------
200.1.4.35
                00:01:00:00:03:22 N/A       ARP-Host
200.1.4.195
                00:01:00:00:03:c2 1         IPCP
===============================================================================
*A:Dut-C#


*A:Dut-C# show service active-subscribers subscriber "hpolSub1"
sap lag-1:2000.1 sla-profile "hpolSlaProf1" detail
===============================================================================
Active Subscribers
===============================================================================
-------------------------------------------------------------------------------
Subscriber hpolSub1 (hpolSubProf2)
-------------------------------------------------------------------------------
I. Sched. Policy : N/A
E. Sched. Policy : N/A                              E. Agg Rate Limit: 6071693
I. Policer Ctrl. : N/A
E. Policer Ctrl. : N/A
Q Frame-Based Ac*: Disabled
Acct. Policy     : N/A                              Collect Stats    : Disabled
Rad. Acct. Pol.  : hpolRadAcctPol
Dupl. Acct. Pol. : N/A
ANCP Pol.        : N/A
HostTrk Pol.     : N/A
IGMP Policy      : N/A
Sub. MCAC Policy : N/A
NAT Policy       : N/A
Def. Encap Offset: none                             Encap Offset Mode: auto
Avg Frame Size   : 120
Sub. ANCP-String : "hpolSub1"
Sub. Int Dest Id : "2000"
Host Trk Rate Adj: N/A
RADIUS Rate-Limit: 10220541
Oper-Rate-Limit  : 10220541
-------------------------------------------------------------------------------
(1) SLA Profile Instance
    - sap:lag-1:2000.1 (VPRN 2000 - grp-Vprn-1)
    - sla:hpolSlaProf1
```

```
-------------------------------------------------------------------------------
Description          : SLA Profile Id hpolSlaProf1
Host Limit           : No Limit
Ingress Qos-Policy   : 2                      Egress Qos-Policy : 2 (vport)
Ingress Queuing Type : Service-queuing (non policer)
Ingr IP Fltr-Id      : n/a                    Egr IP Fltr-Id    : n/a
Ingr IPv6 Fltr-Id    : n/a                    Egr IPv6 Fltr-Id  : n/a
Ingress Report-Rate  : N/A
Egress Report-Rate   : N/A
Egress Remarking     : from Sap Qos
Credit Control Pol.  : N/A
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
IP Address
                MAC Address       PPPoE-SID Origin
-------------------------------------------------------
200.1.4.194
                00:01:00:00:03:c1 1         IPCP
-------------------------------------------------------------------------
SLA Profile Instance statistics
-------------------------------------------------------------------------
                         Packets                 Octets
Off. HiPrio           : 0                       0
Off. LowPrio          : 0                       0
Off. Uncolor          : 0                       0

Queueing Stats (Ingress QoS Policy 2)
Dro. HiPrio           : 0                       0
Dro. LowPrio          : 0                       0
For. InProf           : 0                       0
For. OutProf          : 0                       0

Queueing Stats (Egress QoS Policy 2)
Dro. InProf           : 0                       0
Dro. OutProf          : 0                       0
For. InProf           : 0                       0
For. OutProf          : 0                       0
-------------------------------------------------------------------------
SLA Profile Instance per Queue statistics
-------------------------------------------------------------------------
                         Packets                 Octets
Egress Queue 1
Dro. InProf           : 0                       0
Dro. OutProf          : 0                       0
For. InProf           : 0                       0
For. OutProf          : 0                       0

Egress Queue 2
Dro. InProf           : 0                       0
Dro. OutProf          : 0                       0
For. InProf           : 0                       0
For. OutProf          : 0                       0

Egress Queue 3
Dro. InProf           : 0                       0
Dro. OutProf          : 0                       0
For. InProf           : 0                       0
For. OutProf          : 0                       0

Egress Queue 4
Dro. InProf           : 0                       0
```

```
Dro. OutProf          : 0                      0
For. InProf           : 0                      0
For. OutProf          : 0                      0

Egress Queue 5
Dro. InProf           : 0                      0
Dro. OutProf          : 0                      0
For. InProf           : 0                      0
For. OutProf          : 0                      0

Egress Queue 6
Dro. InProf           : 0                      0
Dro. OutProf          : 0                      0
For. InProf           : 0                      0
For. OutProf          : 0                      0

Egress Queue 7
Dro. InProf           : 0                      0
Dro. OutProf          : 0                      0
For. InProf           : 0                      0
For. OutProf          : 0                      0

Egress Queue 8
Dro. InProf           : 0                      0
Dro. OutProf          : 0                      0
For. InProf           : 0                      0
For. OutProf          : 0                      0
-------------------------------------------------------------------------------
SLA Profile Instance per Policer statistics
-------------------------------------------------------------------------------
                         Packets                  Octets
Ingress Policer 1 (Stats mode: minimal)
Off. All              : 0                      0
Dro. All              : 0                      0
For. All              : 0                      0
===============================================================================
*A:Dut-C#
```

id

Syntax **id** *service-id*

Context show>service

Description Enables the context to display information for a particular service-id.

Parameters *service-id* — 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 — Displays basic 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.

Label	Description (Continued)
Admin Path MTU	The desired largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Oper Path MTU	The actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Delivery	Specifies the type of delivery used by the SDP: GRE or MPLS.
Admin State	The administrative state of this SDP.
Oper State	The operational status of the KeepAlive protocol.
Ingress Label	The label used by the far-end device to send packets to this device in this service by this SDP.
Egress Label	The label used by this device to send packets to the far-end device in this service by this SDP.
Ingress Filter	The ID of the ingress filter policy.
Egress Filter	The ID of the egress filter policy.
Far End	Specifies the IP address of the remote end of the GRE or MPLS tunnel defined by this SDP.
Last Changed	The date and time of the most recent change to this customer.
Signaling	Specifies the signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received on this SDP.
Admin State	Specifies the operating status of the SDP.
Oper State	The operational state of the SDP.
Hello Time	Specifies how often the SDP echo request messages are transmitted on this SDP.
Hello Msg Len	Specifies the length of the SDP echo request messages transmitted on this SDP.
Max Drop Count	Specifies the maximum number of consecutive SDP echo request messages that can be unacknowledged before the keepalive protocol reports a fault.
Hold Down Time	Specifies the amount of time to wait before the keepalive operating status is eligible to enter the alive state.
SDP Delivery Mechanism	When the SDP type is MPLS, a list of LSPs used to reach the far-end router displays. All the LSPs in the list must terminate at the IP address specified in the Far End field. If the SDP type is GRE, then the following message displays: `SDP Delivery Mechanism is not MPLS`

Label	Description (Continued)
Admin Path MTU	The desired largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Oper Path MTU	The actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Delivery	Specifies the type of delivery used by the SDP: GRE or MPLS.
Admin State	The administrative state of this SDP.
Oper State	The operational status of the KeepAlive protocol.
Ingress Label	The label used by the far-end device to send packets to this device in this service by this SDP.
Egress Label	The label used by this device to send packets to the far-end device in this service by this SDP.
Ingress Filter	The ID of the ingress filter policy.
Egress Filter	The ID of the egress filter policy.
Far End	Specifies the IP address of the remote end of the GRE or MPLS tunnel defined by this SDP.
Last Changed	The date and time of the most recent change to this customer.
Signaling	Specifies the signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received on this SDP.
Admin State	Specifies the operating status of the SDP.
Oper State	The operational state of the SDP.
Hello Time	Specifies how often the SDP echo request messages are transmitted on this SDP.
Hello Msg Len	Specifies the length of the SDP echo request messages transmitted on this SDP.
Max Drop Count	Specifies the maximum number of consecutive SDP echo request messages that can be unacknowledged before the keepalive protocol reports a fault.
Hold Down Time	Specifies the amount of time to wait before the keepalive operating status is eligible to enter the alive state.
SDP Delivery Mechanism	When the SDP type is MPLS, a list of LSPs used to reach the far-end router displays. All the LSPs in the list must terminate at the IP address specified in the Far End field. If the SDP type is GRE, then the following message displays: `SDP Delivery Mechanism is not MPLS`

Label	Description (Continued)
Number of SDPs	The total number SDPs applied to this service ID.
Service Access Points	
Service Id	The value that identifies a service.
Port Id	The ID of the access port where this SAP is defined.
Description	Generic information about the SAP.
Encap Value	The value of the label used to identify this SAP on the access port.
Admin State	The desired state of the SAP.
Oper State	The operating state of the SAP.
Last Changed	The date and time of the last change.
Admin MTU	The desired largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Oper MTU	The actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Ingress qos-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.

Label	Description (Continued)
Forwarded Out Profile	The number of out-of-profile packets or octets (rate above CIR) forwarded.
Dropped In Profile	The number of in-profile packets or octets discarded.
Dropped Out Profile	The number of out-of-profile packets or octets discarded.
Forwarded In Profile	The number of in-profile packets or octets (rate below CIR) forwarded.
Forwarded Out Profile	The number of out-of-profile packets or octets (rate above CIR) forwarded.
Ingress Queue 1	The index of the ingress QoS queue of this SAP.
High priority offered	The packets or octets count of the high priority traffic for the SAP.
High priority dropped	The number of high priority traffic packets/octets dropped.
Low priority offered	The packets or octets count of the low priority traffic.
Low priority dropped	The number of low priority traffic packets/octets dropped.
In profile forwarded	The number of in-profile packets or octets (rate below CIR) forwarded.
Out profile forwarded	The number of out-of-profile octets (rate above CIR) forwarded.
Egress Queue 1	The index of the egress QoS queue of the SAP.
In profile forwarded	The number of in-profile packets or octets (rate below CIR) forwarded.
In profile dropped	The number of in-profile packets or octets dropped for the SAP.
Out profile forwarded	The number of out-of-profile packets or octets (rate above CIR) forwarded.
Out profile dropped	The number of out-of-profile packets or octets discarded.
State	Specifies whether DHCP relay is enabled on this SAP.
Info Option	Specifies whether Option 82 processing is enabled on this SAP.
Action	Specifies the Option 82 processing on this SAP or interface: keep, replace or drop.
Circuit ID	Specifies whether the If index is inserted in Circuit ID sub-option of Option 82.
Remote ID	Specifies whether the far-end MAC address is inserted in remote ID sub-option of Option 82

Label	Description (Continued)
Managed by Service	Specifies the service-id of the management VPLS managing this SAP.
Managed by SAP	Specifies the sap-id inside the management VPLS managing this SAP.
Prune state	Specifies the STP state inherited from the management VPLS.
Managed by Service	Specifies the service-id of the management VPLS managing this spoke SDP.
Managed by Spoke	Specifies the sap-id inside the management VPLS managing this spoke SDP.
Prune state	Specifies the STP state inherited from the management VPLS.

Sample Output

```
A:ALA-12# show service id 9000 all
===============================================================================
Service Detailed Information
===============================================================================
Service Id        : 9000               Vpn Id            : 0
Service Type      : VPLS
Description       : This is a distributed VPLS.
Customer Id       : 6                  Last Mgmt Change  : 01/18/2007 10:31:58
Adm               : Up                 Oper              : Down
MTU               : 1514               Def. Mesh VC Id   : 750
SAP Count         : 1                  SDP Bind Count    : 3
-------------------------------------------------------------------------------
Split Horizon Group specifics
-------------------------------------------------------------------------------
Split Horizon Group : splitgroup1
-------------------------------------------------------------------------------
Description       : Split horizon group 1
Instance Id       : 1                  Last Changed      : 01/18/2007 10:31:58
-------------------------------------------------------------------------------
Service Destination Points(SDPs)
-------------------------------------------------------------------------------
Sdp Id 2:22  -(10.10.10.104)
-------------------------------------------------------------------------------
Description     : GRE-10.10.10.104
SDP Id             : 2:22                    Type              : Spoke
VC Type            : Ether                   VC Tag            : n/a
Admin Path MTU     : 0                       Oper Path MTU     : 0
Far End            : 10.10.10.104            Delivery          : GRE
Flags              : SdpOperDown
                     NoIngVCLabel NoEgrVCLabel
                     PathMTUTooSmall
Admin State        : Up                      Oper State        : Down
Ingress Label      : 0                       Egress Label      : 0
Ingress Filter     : n/a                     Egress Filter     : n/a
Last Changed       : 01/18/2007 10:31:58     Signaling         : TLDP

KeepAlive Information :
Admin State        : Disabled                Oper State        : Disabled
Hello Time         : 10                      Hello Msg Len     : 0
Max Drop Count     : 3                       Hold Down Time    : 10
```

```
Statistics            :
I. Fwd. Pkts.      : 0                        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
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                   Encap             : q-tag
Dot1Q Ethertype    : 0x8100                    QinQ Ethertype    : 0x8100

Admin State        : Up                        Oper State        : Down
Flags              : PortOperDown
Last Status Change : 04/11/2007 15:56:40
Last Mgmt Change   : 04/11/2007 17:24:54
Max Nbr of MAC Addr: No Limit                  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                        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
-------------------------------------------------------------------------------
VPLS Rapid Spanning Tree Information
-------------------------------------------------------------------------------
VPLS oper state    : Down                      Core Connectivity : Down
Rstp Admin State   : Up                        Rstp Oper State   : Down
Mode               : Rstp                      Vcp Active Prot.  : N/A

Bridge Id          : 80:01.14:30:ff:00:00:01   Bridge Instance Id: 1
Bridge Priority    : 32768                     Tx Hold Count     : 6
Topology Change    : Inactive                  Bridge Hello Time : 2
Last Top. Change   : 0d 00:00:00               Bridge Max Age    : 20
Top. Change Count  : 0                         Bridge Fwd Delay  : 15

Root Bridge        : N/A
```

```
Primary Bridge     : N/A

Root Path Cost     : 0                        Root Forward Delay: 15
Rcvd Hello Time    : 2                        Root Max Age      : 20
Root Priority      : 32769                    Root Port         : N/A
-------------------------------------------------------------------------------
Forwarding Database specifics
-------------------------------------------------------------------------------
Service Id         : 9000               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
===============================================================================
*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 1510 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.

Label	Description (Continued)
	OAM – Entries created by the OAM process.
Age	The time lapsed since the service was enabled.
Interface	The interface applied to the service.
Port	The port where the SAP is applied.

Sample Output

```
A:ALA-12# show service id 2 arp
===============================================================================
ARP Table
===============================================================================
IP Address      MAC Address       Type    Age       Interface         Port
-------------------------------------------------------------------------------
190.11.1.1      00:03:fa:00:08:22 Other   00:00:00  ies-100-190.11.1  1/1/11:0
===============================================================================
A:ALA-12#
```

arp-host

Syntax **arp-host** [**wholesaler** *service-id*] [**sap** *sap-id* | **interface** *interface-name* | **ip-address** *ip-address*[*/mask*] | **mac** *ieee-address* | {[**port** *port-id*] [**no-inter-dest-id** | **inter-dest-id** *inter-dest-id*]}] [**detail**]
arp-host statistics [**sap** *sap-id* | **interface** *interface-name*]
arp-host summary [**interface** *interface-name*]

Context show>service>id

Description This command displays ARP host related information.

Sample Output

```
*A:Dut-C# show service id 2 arp-host
===============================================================================
ARP host table, service 2
===============================================================================
IP Address      Mac Address       Sap Id              Remaining           MC
                                                      Time                Stdby
-------------------------------------------------------------------------------
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.
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.

Label	Description (Continued)
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 1510 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.

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition. See Common Service Commands on page 1510 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
*               225.0.0.7       sap:2/1/5:7            Fwd
-------------------------------------------------------------------------------
Number of entries: 2


*A:rbae_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:rbae_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:rbae_C# show service id 1 mld-snooping all
===============================================================================
MLD Snooping info for service 1
-------------------------------------------------------------------------------
MLD Snooping Base info
-------------------------------------------------------------------------------
Admin State : Up
Querier     : FE80::12 on SAP 2/1/5
-------------------------------------------------------------------------------
Sap/Sdp                  Oper    MRtr  Send      Max Num   MVR        Num
Id                       State   Port  Queries   Groups    From-VPLS  Groups
-------------------------------------------------------------------------------
sap:1/1/4                Up      No    Disabled  No Limit  Local      0
sap:2/1/5                Up      Yes   Disabled  No Limit  Local      0
sdp:31:1                 Up      No    Disabled  No Limit  N/A        0
sdp:36:1                 Up      No    Disabled  No Limit  N/A        0
-------------------------------------------------------------------------------
MLD Snooping Querier info
-------------------------------------------------------------------------------
Sap Id                  : 2/1/5
IP Address              : FE80::12
Expires                 : 11s
Up Time                 : 0d 00:05:05
Version                 : 2

General Query Interval  : 10s
Query Response Interval : 1.0s
Robust Count            : 2
-------------------------------------------------------------------------------
MLD Snooping Multicast Routers
-------------------------------------------------------------------------------
MRouter
                 Sap/Sdp Id             Up Time        Expires   Version
-------------------------------------------------------------------------------
FE80::12
                 2/1/5                  0d 00:05:05    11s       2
-------------------------------------------------------------------------------
Number of mrouters: 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
-------------------------------------------------------------------------------
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:rbae_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:rbae_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:rbae_C#


*A:rbae_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
===============================================================================
```

mvr

Syntax **mvr**

Context show>service>id>mld-snooping

Description This command displays multicast VPLS registration information.

port-db

Syntax **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* [**detail**]
port-db sdp *sdp-id*:*vc-id* **group** *grp-ipv6-address*

Context show>service>id>mld-snooping

Description This command displays MLD snooping information related to a specific SAP.

Sample Output

```
*A:rbae_C# show service id 1 mld-snooping port-db sap 1/1/4
===============================================================================
MLD Snooping SAP 1/1/4 Port-DB for service 1
===============================================================================
Group Address
                Mode    Type    From-VPLS  Up Time         Expires  Num   MC
                                                                    Src   Stdby
-------------------------------------------------------------------------------
FF04::1
                include dynamic local      0d 00:00:19    0s       1
FF04::2
                include dynamic local      0d 00:00:18    0s       1
-------------------------------------------------------------------------------
Number of groups: 2
===============================================================================
*A:rbae_C#

*A:rbae_C# show service id 1 mld-snooping port-db sap 1/1/4 detail
===============================================================================
MLD Snooping SAP 1/1/4 Port-DB for service 1
-------------------------------------------------------------------------------
MLD Group FF04::1
-------------------------------------------------------------------------------
Mode             : include              Type             : dynamic
Up Time          : 0d 00:00:33          Expires          : 0s
Compat Mode      : MLD Version 2
V1 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
V1 Host Expires  : 0s                                    :
MVR From-VPLS    : local                MVR To-SAP       : local
MC Standby       : no
-------------------------------------------------------
Source Address
                Up Time      Expires  Type     Fwd/Blk
-------------------------------------------------------
2011::1
                0d 00:00:33  20s      dynamic  Fwd
-------------------------------------------------------------------------------
Number of groups: 2
===============================================================================
*A:rbae_C#
```

proxy-db

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:rbae_C# show service id 1 mld-snooping proxy-db
===============================================================================
MLD Snooping Proxy-reporting DB for service 1
===============================================================================
Group Address
                  Mode     Up Time         Num Sources
-------------------------------------------------------------------------------
FF04::1
                  include  0d 00:01:01    1
FF04::2
                  include  0d 00:01:00    1
-------------------------------------------------------------------------------
Number of groups: 2
===============================================================================
*A:rbae_C#


*A:rbae_C# show service id 1 mld-snooping proxy-db detail
===============================================================================
MLD Snooping Proxy-reporting DB for service 1
-------------------------------------------------------------------------------
MLD Group FF04::1
-------------------------------------------------------------------------------
Up Time : 0d 00:01:03                  Mode : include
-------------------------------------------------------------------------
Source Address                                                  Up Time
-------------------------------------------------------------------------
2011::1                                                         0d 00:01:03
-------------------------------------------------------------------------------
MLD Group FF04::2
-------------------------------------------------------------------------------
Up Time : 0d 00:01:02                  Mode : include
-------------------------------------------------------------------------
Source Address                                                  Up Time
-------------------------------------------------------------------------
2011::1                                                         0d 00:01:02
-------------------------------------------------------------------------------
Number of groups: 2
===============================================================================
*A:rbae_C#
```

querier

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:rbae_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:rbae_C#
```

statistics

Syntax **statistics**[**sap** *sap-id* | **sdp** *sdp-id:vc-id*]

Context show>service>id>mld-snooping

Description This command displays MLD snooping statistics.

Sample Output

```
*A:rbae_C# show service id 1 mld-snooping statistics
===============================================================================
MLD Snooping Statistics for service 1
===============================================================================
Message Type            Received      Transmitted   Forwarded
-------------------------------------------------------------------------------
General Queries         109           0             327
Group Queries           0             0             0
Group-Source Queries    0             0             0
V1 Reports              0             0             0
V2 Reports              438           87            0
V1 Done                 0             0             0
Unknown Type            0             N/A           0
-------------------------------------------------------------------------------
Drop Statistics
-------------------------------------------------------------------------------
Bad Length              : 0
Bad MLD Checksum        : 0
Bad Encoding            : 0
No Router Alert         : 0
Zero Source IP          : 0
Wrong Version           : 0
Lcl-Scope Packets       : 0
Rsvd-Scope Packets      : 0

Send Query Cfg Drops    : 0
Import Policy Drops     : 0
Exceeded Max Num Groups : 0
MCAC Policy Drops       : 0
MCS Failures            : 0

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

mstp-configuration

Syntax **mstp-configuration**

Context show>service>id

Description This command displays the MSTP specific configuration data. This command is only valid on a managment 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 *sap-id* — The ID that displays SAP information. See Common Service Commands on page 1510 for *sap-id* 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 Ethertype value.
Admin State	The administrative state of the SAP.
Oper State	The operational state of the SAP.
Flags	Specifies the conditions that affect the operating status of this SAP.
Last Status Change	Specifies the time of the most recent operating status change to this SAP
Last Mgmt Change	Specifies the time of the most recent management-initiated change to this SAP.
Admin MTU	The desired largest service frame size (in octets) that can be transmitted through the SAP to the far-end router, without requiring the packet to be fragmented.

Label	Description (Continued)
Oper MTU	The actual largest service frame size (in octets) that can be transmitted through the SAP to the far-end router, without requiring the packet to be fragmented.
Ingress qos-policy	The ingress QoS policy ID assigned to the SAP.
Egress qos-policy	The egress QoS policy ID assigned to the SAP.
Ingress Filter-Id	The ingress filter policy ID assigned to the SAP.
Egress Filter-Id	The egress filter policy ID assigned to the SAP.
Acct. Pol	The accounting policy ID assigned to the SAP.
Collect Stats	Specifies whether collect stats is enabled.
Dropped	The number of packets and octets dropped due to SAP state, ingress MAC or IP filter, same segment discard, bad checksum, etc.
Off. HiPrio	The number of high priority packets and octets, as determined by the SAP ingress QoS policy, offered by the Pchip to the Qchip.
Off. LowPrio	The number of low priority packets and octets, as determined by the SAP ingress QoS policy, offered by the Pchip to the Qchip.
Off. Uncolor	The number of uncolored packets and octets, as determined by the SAP ingress QoS policy, offered by the Pchip to the Qchip.
Dro. HiPrio	The number of high priority packets and octets, as determined by the SAP ingress QoS policy, dropped by the Qchip due to: MBS exceeded, buffer pool limit exceeded, etc.
Dro. LowPrio	The number of low priority packets and octets, as determined by the SAP ingress QoS policy, dropped by the Qchip due to: MBS exceeded, buffer pool limit exceeded, etc.
For. InProf	The number of in-profile packets and octets (rate below CIR) forwarded by the ingress Qchip.
For. OutProf	The number of out-of-profile packets and octets discarded by the egress Qchip due to MBS exceeded, buffer pool limit exceeded, etc.
Dro. InProf	The number of in-profile packets and octets discarded by the egress Qchip due to MBS exceeded, buffer pool limit exceeded, etc.
Dro. OutProf	The number of out-of-profile packets and octets discarded by the egress Qchip due to MBS exceeded, buffer pool limit exceeded, etc.
For. InProf	The number of in-profile packets and octets (rate below CIR) forwarded by the egress Qchip.
For. OutProf	The number of out-of-profile packets and octets (rate above CIR) forwarded by the egress Qchip.

```
Max Nbr of MAC Addr: No Limit                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.)
Oper Path MTU	The actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Far End	Specifies the IP address of the remote end of the GRE or MPLS tunnel defined by this SDP.
Delivery	Specifies the type of delivery used by the SDP: GRE or MPLS.
Admin State	The administrative state of this SDP.
Oper State	The current status of the SDP.

Label	Description (Continued)
Ingress Label	The label used by the far-end device to send packets to this device in this service by this SDP.

Sample Output

```
A:ALA-12#  show service id 9000 sdp 2:22 detail
===============================================================================
Service Destination Point (Sdp Id : 2:22) Details
-------------------------------------------------------------------------------
 Sdp Id 2:22  -(10.10.10.103)
-------------------------------------------------------------------------------
Description     : GRE-10.10.10.103
SDP Id             : 2:22                       Type              : Spoke
Split Horiz Grp : (DSL-group1
VC Type            : Ether                      VC Tag            : n/a
Admin Path MTU     : 4462                       Oper Path MTU     : 4462
Far End            : 10.10.10.103               Delivery          : GRE
Admin State        : Up                         Oper State        : TLDP Down
Ingress Label      : 0                          Egress Label      : 0
Ingress Filter     : n/a                        Egress Filter     : n/a
Last Changed       : 04/11/2007 11:48:20        Signaling         : TLDP

KeepAlive Information :
Admin State        : Disabled                   Oper State        : Disabled
Hello Time         : 10                         Hello Msg Len     : 0
Max Drop Count     : 3                          Hold Down Time    : 10

Statistics            :
I. Fwd. Pkts.      : 0                          I. Dro. Pkts.     : 0
E. Fwd. Pkts.      : 0                          E. Fwd. Octets    : 0

Associated LSP LIST :
SDP Delivery Mechanism is not MPLS
-------------------------------------------------------------------------------
Rstp Service Destination Point specifics
-------------------------------------------------------------------------------
Mac Move           : Disabled
Rstp Admin State   : Up                         Rstp Oper State   : Down
Core Connectivity  : Down
Port Role          : N/A                        Port State        : Discarding
Port Number        : 2049                       Port Priority     : 128
Port Path Cost     : 10                         Auto Edge         : Enabled
Admin Edge         : Disabled                   Oper Edge         : N/A
Link Type          : Pt-pt                      BPDU Encap        : Dot1d
Designated Bridge  : N/A                        Designated Port Id: 0
Active Protocol    : N/A

Fwd Transitions    : 0                          Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                          Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                          TCN BPDUs tx      : 0
RST BPDUs rcvd     : 0                          RST BPDUs tx      : 0
-------------------------------------------------------------------------------
Number of SDPs : 1
===============================================================================
A:ALA-12#
```

gsmp

Syntax **gsmp**

Context show>service>id

Description This enables the command to display GSMP information.

neighbors

Syntax **neighbors group** [*name*] [*ip-address*]

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.

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.

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

Sample Output

These commands show the configured neighbors per service, regardless that there exists an open TCP connection with this neighbor. The admin state is shown because for a neighbor to be admin enabled, the service, gsmp node, group node and the neighbor node in this service must all be in 'no shutdown' state. Session gives the number of session (open TCP connections) for each configured neighbor.

```
A:active>show>service>id>gsmp# neighbors
===============================================================================
GSMP neighbors
===============================================================================
Group                          Neighbor           AdminState  Sessions
-------------------------------------------------------------------------------
dslam1                         192.168.1.2       Enabled     0
dslam1                         192.168.1.3       Enabled     0
```

```
-------------------------------------------------------------------------------
Number of neighbors shown: 2
===============================================================================
A:active>show>service>id>gsmp#

A:active>show>service>id>gsmp# neighbors group dslam1
===============================================================================
GSMP neighbors
===============================================================================
Group                          Neighbor           AdminState  Sessions
-------------------------------------------------------------------------------
dslam1                         192.168.1.2        Enabled     0
dslam1                         192.168.1.3        Enabled     0
-------------------------------------------------------------------------------
Number of neighbors shown: 2
===============================================================================
A:active>show>service>id>gsmp#


A:active>show>service>id>gsmp# neighbors group dslam1 192.168.1.2
===============================================================================
GSMP neighbors
===============================================================================
Group                          Neighbor           AdminState  Sessions
-------------------------------------------------------------------------------
dslam1                         192.168.1.2        Enabled     0
===============================================================================
A:active>show>service>id>gsmp#
```

sessions

Syntax **sessions** [**group** *name*] **neighbor** *ip-address*] [**port** *port-number*] [**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.
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     dslam1
-------------------------------------------------------
Number of GSMP sessions : 1
=======================================================
A:active>show>service>id>gsmp#


A:active>show>service>id>gsmp# sessions neighbor 192.168.1.2 port 40590
===============================================================================
GSMP sessions for service 999 (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), neigbor 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** *sap-id*] [**wholesaler** *service-id*] [**port** *port-id*] [**inter-dest-id** *intermediate-destination-id*] [**detail**]
host [**sap** *sap-id*] [**wholesaler** *service-id*] [**port** *port-id*] **no-inter-dest-id** [**detail**]
host summary

Context show>service>id

Description This command displays information about static host configured on this service.

Parameters **sap** *sap-id* — Specifies the physical port identifier portion of the SAP definition. See Common Service Commands on page 1510 for *sap-id* command syntax.

intermediate-destination-id — Specifies the intermediate destination identifier which is encoded in the identification strings.

Values Up to 32 characters maximum

summary — Displays summary static host information.

detail — Displays detailed static host information.

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

Values 1 — 2147483647

interface

Syntax **interface** [{[*ip-address*|*ip-int-name*] [*interface-type*] [**detail**] [**family**]}|**summary**]

Context show>service>id

Description This command displays service interface information.

Parameters *ip-address* — Displays information for the specified IP address.

ip-int-name — Displays information for the specified interface name.

detail — Displays detailed interface information.

family *family* **—** Specifies to display only information belonging to the address family IPv4 or IPv6. Only these two values will be accepted.

summary — Displays detailed infterface information.

Sample Output

```
A:cses-E11>config>service>vprn# show service id 10 interface "gi-2-01" detail
===============================================================================
```

```
Interface Table
===============================================================================
-------------------------------------------------------------------------------
Interface
-------------------------------------------------------------------------------
If Name      : gi-2-01
Sub If Name  : si-2
Red If Name  :
Admin State  : Up                                Oper (v4/v6)     : Up/Down
Protocols    : None
-------------------------------------------------------------------------------
Details
-------------------------------------------------------------------------------
Description  : (Not Specified)
If Index     : 6                                 Virt. If Index   : 6
Last Oper Chg: 11/27/2012 13:19:28               Global If Index  : 380
Mon Oper Grp : None
Srrp En Rtng : Disabled                          Hold time        : N/A
Group Port   : 1/1/2
TOS Marking  : Trusted                           If Type          : VPRN Grp
SNTP B.Cast  : False
MAC Address  : d2:30:01:01:00:02                 Mac Accounting   : Disabled
Ingress stats: Disabled
Arp Timeout  : 14400                             IPv6 Nbr ReachTi*: 30
IP Oper MTU  : 1500                              ICMP Mask Reply  : True
Arp Populate : Disabled                          Host Conn Verify : Disabled
Cflowd       : None
LdpSyncTimer : None
LSR Load Bal*: system
uRPF Chk     : disabled
uRPF Ipv6 Chk: disabled
Rx Pkts      : 0                                 Rx Bytes         : 0
Rx V4 Pkts   : 0                                 Rx V4 Bytes      : 0
Rx V6 Pkts   : 0                                 Rx V6 Bytes      : 0
Tx Pkts      : 32                                Tx Bytes         : 3392
Tx V4 Pkts   : 32                                Tx V4 Bytes      : 3392
Tx V4 Discar*: 0                                 Tx V4 Discard By*: 0
Tx V6 Pkts   : 0                                 Tx V6 Bytes      : 0
Tx V6 Discar*: 0                                 Tx V6 Discard By*: 0

Proxy ARP Details
Rem Proxy ARP: Disabled                          Local Proxy ARP  : Disabled
Policies     : none

Proxy Neighbor Discovery Details
Local Pxy ND : Disabled
Policies     : none

DHCP no local server

DHCP Details
Description  : (Not Specified)
Admin State  : Down                              Lease Populate   : 0
Gi-Addr      : Not configured                    Gi-Addr as Src Ip: Disabled
Action       : Keep                              Trusted          : Disabled

DHCP Proxy Details
Admin State  : Down
Lease Time   : N/A
Emul. Server : Not configured
```

```
Subscriber Authentication Details
Auth Policy  : None

DHCP6 Relay Details
Description  : (Not Specified)
Admin State  : Down                              Lease Populate   : 0
Oper State   : Down                              Nbr Resolution   : Disabled
If-Id Option : None                              Remote Id        : Disabled
Src Addr     : Not configured

DHCP6 Server Details
Admin State  : Down                              Max. Lease States: 8000

ISA Tunnel redundant next-hop information
Static Next-*:
Dynamic Next*:

ICMP Details
Redirects    : Number - 100                      Time (seconds)   - 10
Unreachables : Number - 100                      Time (seconds)   - 10
TTL Expired  : Number - 100                      Time (seconds)   - 10

IPCP Address Extension Details
Peer IP Addr*: Not configured
Peer Pri DNS*: Not configured
Peer Sec DNS*: Not configured
-------------------------------------------------------------------------------
Qos Details
-------------------------------------------------------------------------------

Ing Qos Poli*: (none)                            Egr Qos Policy   : (none)
Ingress FP Q*: (none)                            Egress Port QGrp : (none)
Ing FP QGrp *: (none)                            Egr Port QGrp In*: (none)
-------------------------------------------------------------------------------
Interfaces : 1
===============================================================================
* indicates that the corresponding row element may have been truncated.
A:cses-E11>config>service>vprn#
```

retailers

Syntax **retailers**

Context show>service>id

Description This command displays service retailer information.

Sample Output

```
*A:ALA-48>config# show service id 101 retailers
===============================================================================
Retailers for service 101
===============================================================================
Retailer Svc ID              Num Static Hosts         Num Dynamic Hosts
-------------------------------------------------------------------------------
102                          3                        1
```

```
105                           0                         1
-------------------------------------------------------------------------------
Number of retailers : 2
===============================================================================
*A:ALA-48>config#
```

wholesalers

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
===============================================================================
DHCP lease state table, service 105
===============================================================================
IP Address      Mac Address       Sap/Sdp Id          Remaining  Lease    MC
                                                      LifeTime   Origin   Stdby
-------------------------------------------------------------------------------
Wholesaler 101 Leases
-------------------------------------------------------------------------------
103.3.2.62      00:00:1f:bd:00:c6 lag-1:105           00h00m39s  RADIUS
-------------------------------------------------------------------------------
Number of lease states : 1
===============================================================================
*A:ALA-48>config#
```

split-horizon-group

Syntax **split-horizon-group** [*group-name*]

Context show>service>id

Description This command displays service split horizon groups.

Parameters *group-name* — 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** *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

Context show>service>id

Description This command displays Display static hosts configured on this service.

Parameters **sap** *sap-id* — Specifies the physical port identifier portion of the SAP definition. See Common Service Commands on page 1510 for *sap-id* command syntax.

intermediate-destination-id — Specifies the intermediate destination identifier which is encoded in the identification strings.

Values Up to 32 characters maximum

summary — Displays summary static host information.

detail — Displays detailed static host information.

wholesaler *service-id* **—** 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 max age    : 20
Bridge priority   : 1                   Topology change   : Inactive
Last Top. change  : 0d 00:00:00         Top. change count : 0

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

Root path cost    : 1                   Root forward delay: 15
Root hello time   : 1                   Root max age      : 20
Root priority     : 0                   Root port         : vcp

-------------------------------------------------------------------------------
Spanning Tree Specifics
-------------------------------------------------------------------------------
SAP Identifier    : 1/1/7:0             RSTP State        : Down
STP Port State    : Forwarding          BPDU encap        : dot1d
Port Number       : 2048                Priority          : 128
Cost              : 10                  Fast Start        : Disabled
Designated Port   : 34816               Designated Bridge : 02:fa:00:04:54:01
===============================================================================
A:ALA-12#
```

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 1510 for *sap-id* 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 1510 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 Subscriber	The ID of the originating subscriber.
Redirection filter id	Displays the Redirection Filter ID number.
Status: active/inactive	Displays the status of one-time HTTP redirection.
Filter-id-source	Displays source of the HTTP filter.

Sample Output

```
*A:Dut-C>#  show service id 3 subscriber-hosts
===============================================================
Subscriber Host table
===============================================================
Sap                    Subscriber
  IP Address
    MAC Address          PPPoE-SID     Origin       Fwding State
---------------------------------------------------------------
2/1/5:2                    TEACAHEH74
  11.11.1.61
    00:80:00:00:00:0a          N/A               ARP-Host        Fwding
[pw-11:11]                  VIACAHEH74
  11.11.1.2
    00:00:11:11:01:02          N/A               ARP-Host        Fwding
[pw-11:12]                  pw-11:12
  11.11.1.3
    00:00:11:11:01:03          N/A               ARP-Host        Fwding
[pw-11:13]                  pw-11:13
  11.11.1.4
    00:00:11:11:01:04          N/A               ARP-Host        Fwding
[pw-22:22]                  XMACAHEH74
  22.22.1.2
    00:00:22:22:01:02          N/A               ARP-Host        Fwding

[pw-33:33]                  IUASAHEH74
  33.33.1.2
    00:00:33:33:01:02          N/A               ARP-Host        Fwding
---------------------------------------------------------------
Number of subscriber hosts: 6
===============================================================
*A:Dut-C>#


A:Dut-A# show service id 100 subscriber-hosts ip 10.100.1.5
===============================================================================
Subscriber Host table
===============================================================================
Sap                    IP Address      MAC Address       Origin(*) Subscriber
-------------------------------------------------------------------------------
1/2/1:102              10.100.1.5      00:10:00:00:00:03 -/D/- alcatel_100
-------------------------------------------------------------------------------
Number of subscriber hosts : 1
===============================================================================
(*) S=Static Host, D=DHCP Lease, N=Non-Sub-Traffic
A:Dut-A#
```

sdp

Syntax **sdp** *sdp-id* **pw-port** [*pw-port-id*]
sdp *sdp-id* **pw-port**
sdp *sdp-id* **pw-port** [*pw-port-id*] [**statistics]**
sdp [**consistent | inconsistent | na] egressifs**
sdp *sdp-id* **keep-alive-history**
sdp far-end *ip-address* | *ipv6-address* **keep-alive-history**
sdp [*sdp-id*] detail

sdp far-end *ip-address* | *ipv6-address* **detail**

Context show>service>sdp

Description This command displays SDP information.

If no optional parameters are specified, a summary SDP output for all SDPs is displayed.

Parameters *sdp-id* — The SDP ID for which to display information.

Default All SDPs.

Values 1 — 17407

pw-port *pw-port-id* **—** Displays the SAP identifier for PW-SAPs.

Values 1 — 10239

far-end *ip-address* **—** Displays only SDPs matching with the specified far-end IP address.

Default SDPs with any far-end IP address.

detail — Displays detailed SDP information.

Default SDP summary output.

keep-alive-history — Displays the last fifty SDP keepalive events for the SDP.

Default SDP summary output.

Sample Output

```
===============================================================================
*A:ALA-12>config>service# show service sdp 1 pw-port
===============================================================================
Service Destination Point (sdp Id 1 Pw-Port)
===============================================================================
Pw-port   VC-Id     Adm     Encap       Opr     VC Type     Egr       Monitor
                                                            Shaper    Oper
                                                            VPort     Group
-------------------------------------------------------------------------------
1         1         up      dot1q       up      ether
2         2         up      qinq        up      ether
3         3         up      dot1q       up      ether
4         4         up      qinq        up      ether
-------------------------------------------------------------------------------
Entries found : 4
===============================================================================

*A:ALA-12>config>service# show service sdp 1 pw-port 3
===============================================================================
Service Destination Point (Sdp Id 1 Pw-Port 3)
===============================================================================
SDP Binding port     : lag-1
VC-Id                : 3                    Admin Status       : up
Encap                : dot1q                Oper Status        : up
VC Type              : ether
Oper Flags           : (Not Specified)
Monitor Oper-Group   : (Not Specified)
===============================================================================
```

```
*A:ALA-12>config>service# show service sdp 1 pw-port 3 statistics
===============================================================================
Service Destination Point (Sdp Id 1 Pw-Port 3)
===============================================================================
SDP Binding port     : lag-1
VC-Id                : 3                    Admin Status       : up
Encap                : dot1q                Oper Status        : up
VC Type              : ether
Oper Flags           : (Not Specified)
Monitor Oper-Group   : (Not Specified)

Statistics            :
I. Fwd. Pkts.        : 0                    I. Dro. Pkts.      : 0
I. Fwd. Octs.        : 0                    I. Dro. Octs.      : 0
E. Fwd. Pkts.        : 0                    E. Fwd. Octets     : 0
===============================================================================
```

subscriber-using

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 1510 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-name*

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

app-profile — Displays the application specified profile.

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

```
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 *ip-address* [**lag** *lag-id*]
multi-chassis mc-lag [**peer** *ip-address* [**lag** *lag-id*]] **statistics**
multi-chassis sync [**peer** *ip-address*] [**detail**]
multi-chassis sync [**peer** *ip-address*] **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
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
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:pc1#

A:pc1# show redundancy multi-chassis sync statistics
===============================================================================
Multi-chassis Peer Sync Stats
===============================================================================
Peer IP Address         : 10.10.10.102
Packets Tx Total        : 511
Packets Tx Hello        : 510
Packets Tx Data         : 0
Packets Tx Other        : 1
Packets Tx Error        : 0
Packets Rx Total        : 511
Packets Rx Hello        : 510
Packets Rx Data         : 0
Packets Rx Other        : 1
Packets Rx Error        : 0
Packets Rx Header Err   : 0
Packets Rx Body Err     : 0
Packets Rx Seq Num Err  : 0
===============================================================================
Peer IP Address         : 10.10.20.1
Packets Tx Total        : 0
Packets Tx Hello        : 0
Packets Tx Data         : 0
Packets Tx Other        : 0
Packets Tx Error        : 0
Packets Rx Total        : 0
Packets Rx Hello        : 0
Packets Rx Data         : 0
Packets Rx Other        : 0
Packets Rx Error        : 0
Packets Rx Header Err   : 0
Packets Rx Body Err     : 0
Packets Rx Seq Num Err  : 0
===============================================================================
A:pc1#

A:pc1# show redundancy multi-chassis sync peer 10.10.10.102 statistics
===============================================================================
Multi-chassis Peer Sync Stats
===============================================================================
Peer IP Address         : 10.10.10.102
Packets Tx Total        : 554
Packets Tx Hello        : 553
Packets Tx Data         : 0
```

```
Packets Tx Other        : 1
Packets Tx Error        : 0
Packets Rx Total        : 554
Packets Rx Hello        : 553
Packets Rx Data         : 0
Packets Rx Other        : 1
Packets Rx Error        : 0
Packets Rx Header Err   : 0
Packets Rx Body Err     : 0
Packets Rx Seq Num Err  : 0
===============================================================================
A:pc1#
```

mc-ipsec

Syntax **mc-ipsec peer** *addr* [**tunnel-group** *group-id*]

Context show>redundancy>multi-chassis

Description This command displays the IPsec multi-chassis states. Optionally, only the states of the specified tunnel-groups will be displayed.

Parameters *addr* — Specifies the address of the peer.

group-id — Specifies the tunnel-group ID.

Sample Output

```
===============================================================================
Multi-Chassis MC-IPSec
===============================================================================
Peer Name       : (Not Specified)
Peer Addr       : 2.2.2.2
Keep Alive Intvl: 10                  Hold on Nbr Fail     : 3
BFD Intf Name   : None
BFD Dest Addr   :
Last update     : 03/20/2012 22:48:55

===============================================================================
Multi-Chassis IPsec Multi Active Tunnel-Group Table
===============================================================================
ID             Peer Group     Priority  Preempt   Admin State    Mastership
-------------------------------------------------------------------------------
1              1              100       Disabled  Up             eligible
-------------------------------------------------------------------------------
Multi Active Tunnel Group Entries found: 1
===============================================================================
```

mc-ring

Syntax **mc-ring peer** *ip-address* **statistics**
mc-ring peer *ip-address* [**ring** *sync-tag* [**detail**|**statistics**]]
mc-ring peer *ip-address* **ring** *sync-tag* **ring-node** [*ring-node-name* [**detail**|**statistics**]]
mc-ring global-statistics

Context show>redundancy>multi-chassis

Description This command displays multi-chassis ring information.

Parameters *ip-address* — Specifies the address of the multi-chassis peer to display.

ring *sync-tag* **—** Specifies a synchronization tag to be displayed that was used while synchronizing this port with the multi-chassis peer.

node *ring-node-name* **—** 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).

Label	Description (Continued)
	localBroken – The inband control connection with the peer is known to be broken due to local failure or local administrative action. shutdown – The ring is shutdown.
Failure Reason	
No. of MC Ring entries	

Sample Output

```
show redundancy multi-chassis mc-ring peer 10.0.0.2 ring ring11 detail
===============================================================================
Multi-Chassis MC-Ring Detailed Information
===============================================================================
Peer           : 10.0.0.2
Sync Tag       : ring11
Port ID        : 1/1/3
Admin State    : inService
Oper State     : connected
Admin Change   : 01/07/2008 21:40:07
Oper Change    : 01/07/2008 21:40:24
Failure Reason : None
-------------------------------------------------------------------------------
In Band Control Path
-------------------------------------------------------------------------------
Service ID     : 10
Interface Name : to_an1
Oper State     : connected
Dest IP        : 10.10.0.2
Src  IP        : 10.10.0.1
-------------------------------------------------------------------------------
VLAN Map B Path Provisioned
-------------------------------------------------------------------------------
range 13-13
range 17-17
-------------------------------------------------------------------------------
VLAN Map Excluded Path Provisioned
-------------------------------------------------------------------------------
range 18-18
-------------------------------------------------------------------------------
VLAN Map B Path Operational
-------------------------------------------------------------------------------
range 13-13
range 17-17
-------------------------------------------------------------------------------
VLAN Map Excluded Path Operational
-------------------------------------------------------------------------------
range 18-18
===============================================================================
```

```
*A:ALA-48>show>redundancy>multi-chassis# mc-ring peer 192.251.10.104
===============================================================================

MC Ring entries
===============================================================================
Sync Tag                          Oper State      Failure Reason
-------------------------------------------------------------------------------
No. of MC Ring entries: 0
===============================================================================

show redundancy multi-chassis mc-ring peer 10.0.0.2

===============================================================================
MC Ring entries
===============================================================================
Sync Tag                          Oper State      Failure Reason
-------------------------------------------------------------------------------
ring11                            connected       None
ring12                            shutdown        None
-------------------------------------------------------------------------------
No. of MC Ring entries: 4
===============================================================================


show redundancy multi-chassis mc-ring peer 10.0.0.2 ring ring11 ring-node an1 detail
===============================================================================
Multi-Chassis MC-Ring Node Detailed Information
===============================================================================
Peer           : 10.0.0.2
Sync Tag       : ring11
Node Name      : an1
Oper State Loc : connected
Oper State Rem : notTested
In Use         : True
Admin Change   : 01/07/2008 21:40:07
Oper Change    : 01/07/2008 21:40:25
Failure Reason : None
-------------------------------------------------------------------------------
Ring Node Connectivity Verification
-------------------------------------------------------------------------------
Admin State    : inService
Service ID     : 11
VLAN Tag       : 11
Dest IP        : 10.11.3.1
Src  IP        : None
Interval       : 1 minutes
Src MAC        : None
===============================================================================

show redundancy multi-chassis mc-ring peer 10.0.0.2 ring ring11 ring-node
===============================================================================
MC Ring Node entries
===============================================================================
Name                              Loc Oper St.      Failure Reason
  In Use                            Rem Oper St.
-------------------------------------------------------------------------------
an1                               connected         None
  Yes                               notTested
an2                               connected         None
  Yes                               notTested
-------------------------------------------------------------------------------
```

```
No. of MC Ring Node entries: 2
===============================================================================
```

show redundancy multi-chassis ring peer statistics Output — The following table describes multi-chassis ring peer output fields

Label	Description
Message	Displays the message type.
Received	Indicates the number of valid MC-Ring signaling messages received from the peer.
Transmitted	Indicates the number of valid MC-Ring signaling messages transmitted from the peer.
MCS ID Request	Displays the number of valid MCS ID requests were received from the peer.
MCS ID Response	Displays the number of valid MCS ID responses were received from the peer.
Ring Exists Request	Displays the number of valid 'ring exists' requests were received from the peer.
Ring Exists Response	Displays the number of valid ring exists' responses were received from the peer.
Keepalive	Displays the number of valid MC-Ring control packets of type 'keepalive' were received from the peer.

```
*A:ALA-48>show>redundancy>multi-chassis# mc-ring peer 192.251.10.104 statistics
===============================================================================
MC Ring statistics for peer 192.251.10.104
===============================================================================
Message                                                 Received     Transmitted
-------------------------------------------------------------------------------
MCS ID Request                                          0            0
MCS ID Response                                         0            0
Ring Exists Request                                     0            0
Ring Exists Response                                    0            0
Keepalive                                               0            0
-------------------------------------------------------------------------------
Total                                                   0            0
===============================================================================
*A:ALA-48>show>redundancy>multi-chassis#
```

show mc-ring ring-node Output

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

Label	Description (Continued)
	notProvisioned – Connection verification is not provisioned. configErr – Connection verification is provisioned but a configuration error prevents it from operating properly. notTested – Connection verification is administratively disabled or is not possible in the current situation. testing – Connection Verification is active, but no results are yet available. connected – The ring node is reachable. disconnected – Connection verification has timed out.
In Use	Displays "True" if the ring node is referenced on an e-pipe or as an inter-dest-id on a static host or dynamic lease.

show mc-ring global-statistics Output

Label	Description
Rx	Displays the number of MC-ring signaling packets were received by this system.
Rx Too Short	Displays the number of MC-ring signaling packets were received by this system that were too short.
Rx Wrong Authentication	Displays the number of MC-ring signaling packets were received by this system with invalid authentication.
Rx Invalid TLV	Displays the number of MC-ring signaling packets were received by this system with invalid TLV.
Rx Incomplete	Displays the number of MC-ring signaling packets were received by this system that were incomplete.
Rx Unknown Type	Displays the number of MC-ring signaling packets were received by this system that were of unknown type.
Rx Unknown Peer	Displays the number of MC-ring signaling packets were received by this system that were related to an unknown peer.
Rx Unknown Ring	Displays the number of MC-ring signaling packets were received by this system that were related to an unknown ring.
Rx Unknown Ring Node	Displays the number of MC-ring signaling packets were received by this system that were related to an unknown ring node.

Label	Description (Continued)
Tx	Displays the number of MC-ring signaling packets were transmitted by this system.
Tx No Buffer	Displays the number of MC-ring signaling packets could not be transmitted by this system due to a lack of packet buffers.
Tx Transmission Failed	Displays the number of MC-ring signaling packets could not be transmitted by this system due to a transmission failure.
Tx Unknown Destination	Displays the number of MC-ring **unknown destination** signaling packets were transmitted by this system.
Missed Configuration Events	Displays the number of missed configuration events on this system.
Missed BFD Events	Displays the number of missed BFD events on this system.

Sample Output

```
*A:ALA-48>show>redundancy>multi-chassis# mc-ring global-statistics
===============================================================================
Global MC Ring statistics
===============================================================================
Rx                          : 0
Rx Too Short                : 0
Rx Wrong Authentication     : 0
Rx Invalid TLV              : 0
Rx Incomplete               : 0
Rx Unknown Type             : 0
Rx Unknown Peer             : 0
Rx Unknown Ring             : 0
Rx Unknown Ring Node        : 0
Tx                          : 36763
Tx No Buffer                : 0
Tx Transmission Failed      : 0
Tx Unknown Destination      : 0
Missed Configuration Events : 0
Missed BFD Events           : 0
===============================================================================
*A:ALA-48>show>redundancy>multi-chassis#
```

lease-state

Syntax **lease-state** [**wholesaler** *service-id*] [**sap** *sap-id*|**sdp** *sdp-id:vc-id*| **interface** *interface-name*|**ip-address** *ip-address*[/*mask*]|**chaddr** *ieee-address*|**mac** *ieee-address*|{[**port** *port-id*] [**no-inter-dest-id** | **inter-dest-id** *inter-dest-id*]}] [**session** {**none**|**ipoe**}] [**detail**]

Context show>service>id>dhcp

Description This command displays DHCP lease state information. Note that the **wholesaler** *service-id* parameter is applicable only in the VPRN context.

Parameters **wholesaler** *service-id* **—** 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 *sap-id* — Specifies the physical port identifier portion of the SAP definition. See Common Service Commands on page 1510 for *sap-id* command syntax.

sdp *sdp-id* — The SDP identifier.

Values 1 — 17407

vc-id — The virtual circuit identifier. This value is used to validate the VC ID portion of each mesh SDP binding defined in the service. The default value of this object is equal to the service ID.

Values 1 — 4294967295

interface *interface-name* **—** Specifies the interface name up to 32 characters in length.

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

port *port-id* **—** The DHCP lease state local specifies that the DHCP lease state is learned by either a SAP or SDP. When the value is SAP, the value indicates the SAP for which this entry contains information.

chaddr — Specifies the MA address of the DHCP lease state.

interface *interface-name* **—** Shows information for the specified IP interface.

detail — Displays detailed lease state information.

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

session — Shows DHCPv4 lease states for hosts that are associated with an IPoE session or for hosts that are not associated with an IPoE session.

Values none, ipoe

detail — Displays detailed information.

Sample Output

```
*A:ALA-48>config# show service id 101 dhcp lease-state
===============================================================================
DHCP lease state table, service 101
===============================================================================
IP Address      Mac Address       Sap/Sdp Id          Remaining  Lease    MC
                                                      LifeTime   Origin   Stdby
-------------------------------------------------------------------------------
102.1.1.52      00:00:1f:bd:00:bb lag-1:101           00h02m56s  DHCP-R
103.3.2.62      00:00:1f:bd:00:c6 lag-1:105           00h02m59s  RADIUS
-------------------------------------------------------------------------------
Number of lease states : 2
===============================================================================
*A:ALA-48>config#


*A:ALA-48>config# show service id 105 dhcp lease-state  wholesaler 101
===============================================================================
```

```
DHCP lease state table, service 105
===============================================================================
IP Address      Mac Address       Sap/Sdp Id          Remaining  Lease    MC
                                                      LifeTime   Origin   Stdby
-------------------------------------------------------------------------------
Wholesaler 101 Leases
-------------------------------------------------------------------------------
103.3.2.62      00:00:1f:bd:00:c6 lag-1:105           00h00m39s  RADIUS
-------------------------------------------------------------------------------
Number of lease states : 1
===============================================================================
*A:ALA-48>config#
```

lease-state

Syntax **lease-state** [**detail**] [**wholesaler** *service-id*] [**session** {**none**|**ipoe**|**ppp**}]|

lease-state [**detail**] **interface** *interface-name* [**wholesaler** *service-id*] [**session** {**none**|**ipoe**|**ppp**}]

lease-state [**detail**] *ipv6-address ipv6-prefix*[/pre*fix-length*] [**wholesaler** *service-id*] [**session** {**none**|**ipoe**|**ppp**}]

lease-state [**detail**] **mac** *ieee-address* [**wholesaler** *service-id*] [**session** {**none**|**ipoe**|**ppp**}]

Context show>service>id>dhcp6

Description This command displays DHCP6 lease state information. Note that the **wholesaler** *service-id* parameter is applicable only in the VPRN context.

Parameters **wholesaler** *service-id* **—** 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 *sap-id* — Specifies the physical port identifier portion of the SAP definition. See Common Service Commands on page 1510 for *sap-id* command syntax.

sdp *sdp-id* — The SDP identifier.

Values 1 — 17407

vc-id — The virtual circuit identifier. This value is used to validate the VC ID portion of each mesh SDP binding defined in the service. The default value of this object is equal to the service ID.

Values 1 — 4294967295

interface *interface-name* **—** Specifies the interface name up to 32 characters in length.

ipv6 *ipv6-address*[/*mask*] **—** v6Shows information for the specified IPv6 address and mask.

port *port-id* **—** The DHCP6 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 DHCP6 lease state.

interface *interface-name* **—** Shows information for the specified IP interface.

detail — Displays detailed lease state information.

inter-dest-id — Indicates the intermediate destination identifier received from either the DHCP6 or the RADIUS server or the local user database.

session — Shows DHCPv6 lease states for clients that are associated with an IPoE session or for clients that are associated with a PPP session or for clients that are not associated with an IPoE session.

Values none, ipoe, ppp

detail — Displays detailed information.

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 1510 for *sap-id* 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#


VPRN:
*A:ALA-49>show>service# id 1 dhcp summary
===============================================================================
DHCP Summary, service 1
===============================================================================
Interface Name                  Arp      Used/                Info    Admin
  SapId/Sdp                     Populate Provided             Option  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-ce1                          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                  Arp      Used/                Info    Admin
  SapId/Sdp                     Populate Provided             Option  State
-------------------------------------------------------------------------------
SpokeTerm                       No       0/0                  Keep    Up
  sdp:spoke-3:3                          0/0
new-if                          No       0/1                  Keep    Up
  sap:1/2/19:0                           0/1
test123                         No       0/0                  Keep    Up
  sap:3/2/4:50/5                         0/0
testabc                         No       0/0                  Keep    Up
  sap:1/2/20:0                           0/0
-------------------------------------------------------------------------------
Interfaces: 4
===============================================================================
*A:ALA-48#


*A:ALA-48# show service id 88 dhcp summary interface SpokeTerm
===============================================================================
DHCP Summary, service 88
===============================================================================
Interface Name                  Arp      Used/                Info    Admin
  SapId/Sdp                     Populate Provided             Option  State
-------------------------------------------------------------------------------
SpokeTerm                       No       0/0                  Keep    Up
  sdp:spoke-3:3                          0/0
-------------------------------------------------------------------------------
Interfaces: 1
===============================================================================
*A:ALA-48#
```

statistics

Syntax **statistics** [**interface** *ip-int-name* | *ip-address*]

Context show>router>dhcp

Description This command displays statistics for DHCP relay and DHCP snooping. If no IP address or interface name is specified, then all configured interfaces are displayed. If an IP address or interface name is specified, then only data regarding the specified interface is displayed.

Parameters **interface** *ip-int-name* | *ip-address* — Displays statistics for the specified IP interface or IP address.

Output **Show DHCP Statistics Output —** The following table describes the output fields for DHCP statistics.

Label	Description
Received Packets	The number of packets received from the DHCP clients.
Transmitted Packets	The number of packets transmitted to the DHCP clients.
Received Malformed Packets	The number of malformed packets received from the DHCP clients.
Received Untrusted Packets	The number of untrusted packets received from the DHCP clients.
Client Packets Discarded	The number of packets received from the DHCP clients that were discarded.
Client Packets Relayed	The number of packets received from the DHCP clients that were forwarded.
Client Packets Snooped	The number of packets received from the DHCP clients that were snooped.
Server Packets Discarded	The number of packets received from the DHCP server that were discarded.
Server Packets Relayed	The number of packets received from the DHCP server that were forwarded.
Server Packets Snooped	The number of packets received from the DHCP server that were snooped.
Client packets proxied (RADIUS)	The number of packets that were generated from RADIUS data and not relayed from a server
Client packets proxied (Lease-Split)	Indicates the total number of client packets proxied by the DHCP relay agent based on data received from a RADIUS server
DHCP RELEASEs spoofed	Indicates the total number of DHCP release messages spoofed by the DHCP relay agent to the DHCP server.
DHCP 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#
```

mrouters

Syntax **mrouters [detail]**

Context show>service>id>igmp-snooping

Description Displays all multicast routers.

Parameters **detail** — Displays detailed information.

Output **Show igmp-snooping mrouters —** The following table describes the show igmp-snooping mrouters output fields:

Label	Description
MRouter	Specifies the multicast router port.
Sap/Sdp Id	Specifies the SAP and SDP ID multicast router ports.
Up Time	Displays the length of time the mrouter has been up.
Expires	Displays the amount of time left before the query interval expires.
Version	Displays the configured version of IGMP running on this interface.

Sample Output

```
*A:ALA-48# show service id 700 igmp-snooping mrouters
===============================================================================
IGMP Snooping Multicast Routers for service 700
===============================================================================
MRouter          Sap/Sdp Id             Up Time        Expires   Version
-------------------------------------------------------------------------------
```

```
Number of mrouters: 0
===============================================================================
*A:ALA-48#
```

mvr

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-*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 1510 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
===============================================================================
```

```
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** *sap-id* | **sdp** *sdp-id:vc-id*]

Context show>service>id>igmp-snooping

Description Displays information on static IGMP snooping source groups for the VPLS service.

Parameters **sap** *sap-id* — Displays static IGMP snooping source groups for a specific SAP. See Common Service Commands on page 1510 for *sap-id* command syntax.

sdp *sdp-id* — Displays the IGMP snooping source groups 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# 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 1510 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** *policy-name* [**detail**]

Context show>mcast-management

Description This command displays multicast path management bandwidth policy information.

Parameters *policy-name* — 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** *router-instance* | **vpls** *service-id*] [**mda** *slot*[/*mda*]] [**group** *ip-address* [**source** *ip-address*]] [**path** *path-type*] [**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#
```

mcast-reporting-dest

Syntax **mcast-reporting-dest** [*mcast-reporting-dest-name*]

Context show>mcast-management

Description This command displays multicast path management reporting destination information.

mda

Syntax **mda** [**slot**[/*mda*]] [**path** *path-type*]

Context show>mcast-management

Description This command displays multicast path management MDA related information.

Parameters *path-type* — Specifies the path type.

Values primary, secondary, ancillary

Output

```
*A:Dut-C# *show mcast-management mda 10/2*
===============================================================================
MDA 10/2
===============================================================================
S/M  Bw-policy                        Type          Limit    In-use-Bw
Admin
-------------------------------------------------------------------------------
10/2 gie                              primary       0 Ms     0 Ms        up
     gie                              secondary     0 Ms     0 Ms        up
     gie                              ancillary     0 Ms      219. 64 Ms up
===============================================================================
*A:Dut-C#
```

group

Syntax **group** [*grp-ip-address*]
group summary

Context show>router>igmp

Description This command displays IGMP group information.

group-interface

Syntax **group-interface** [**fwd-service** *service-id*] [*ip-int-name*] [**detail**]

Context show>router>igmp

Description This command displays IGMP group-interface information.

hosts

Syntax **hosts** [**group** *grp-address*] [**detail**] [**fwd-service** *service-id*] [**grp-interface** *ip-int-name*]
hosts [**host** *ip-address*] [**group** *grp-address*] [**detail**]
hosts summary

Context show>router>igmp

Description This command displaysIGMP hosts information.

interface

Syntax **interface** [*ip-int-name*|*ip-address*] [**group**] [*grp-ip-address*] [**detail**]

Context show>router>igmp

Description This command displays IGMP interface information.

mcast-reporting-statistics

Syntax **mcast-reporting-statistics** [*ip-address*]

Context show>router>igmp

Description This command displays IGMP mcast reporting statistics.

ssm-translate

Syntax **ssm-translate** [*interface-name*]

Context show>router>igmp

Description This command displays SSM translate configuration information.

static

Syntax **static** [*ip-int-name*|*ip-addr*]

Context show>router>igmp

Description This command displaysIGMP static group/source configuration information.

statistics

Syntax **statistics** [*ip-int-name*|*ip-address*]
statistics host [*ip-address*]

Context show>router>igmp

Description This command displaysIGMP statistics information.

status

Syntax **status**

Context show>router>igmp

Description This command displays IGMP status information.

tunnel-interface

Syntax **tunnel-interface**

Context show>router>igmp

Description This command displaysIGMP tunnel-interface information.

Clear Commands

id

Syntax **id** *service-id*

Context clear>service
clear>service>statistics

Description This command clears the identification for a specific service.

Parameters *service-id* — 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** *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*]

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** *msap-id*

Context clear>service>id

Description This command clears Managed SAP information.

Parameters *msap-id* — 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** *msap-policy-name*

Context clear>service>id

Description This command clears Managed SAPs created by the Managed SAP policy.

Parameters *msap-policy-name* — 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** *sub-ident-string*

Context clear>service>statistics

Description This command clears the statistics for a particular subscriber.

Parameters *sub-ident-string* — Clears statistics for the specified subscriber identification string.

fdb

Syntax **fdb** {**all** | **mac** *ieee-address* | **sap** *sap-id*] | **mesh-sdp** *sdp-id*[**:***vc-id*] | **spoke-sdp** *sdp-id*:*vc-id*}

Context clear>service>id

Description This command clears FDB entries for the service.

Parameters **all —** Clears all FDB entries.

mac *ieee-address* **—** Clears only FDB entries in the FDB 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.

sap-id — Specifies the physical port identifier portion of the SAP definition. See Common Service Commands on page 1510 for *sap-id* 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.

sdp-id — The SDP ID for which to clear associated FDB entries.

vc-id — The virtual circuit ID on the SDP ID for which to clear associated FDB entries.

Values			
	sdp-id[:vc-id]	*sdp-id*	1 — 17407
		vc-id	1 — 4294967295
	sdp-id:vc-id	sdp-id	1 — 17407
		vc-id	1 — 4294967295

mesh-sdp

Syntax **mesh-sdp** *sdp-id*[**:***vc-id*] **ingress-vc-label**

Context clear>service>id

Description Clears and resets the mesh SDP bindings for the service.

Parameters *sdp-id* — The mesh SDP ID to be reset.

Values 1 — 17407

vc-id — The virtual circuit ID on the SDP ID to be reset.

Default All VC IDs on the SDP ID.

Values 1 — 4294967295

spoke-sdp

Syntax **spoke-sdp** *sdp-id*:*vc-id* **ingress-vc-label**

Context clear>service>id

Description Clears and resets the spoke SDP bindings for the service.

Parameters *sdp-id* — The spoke SDP ID to be reset.

Values 1 — 17407

vc-id — The virtual circuit ID on the SDP ID to be reset.

Values 1 — 4294967295

sap

Syntax **sap** *sap-id* {**all** | **counters** | **stp**}

Context clear>service>statistics

Description Clears SAP statistics for a SAP.

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition. See Common Service Commands on page 1510 for *sap-id* 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** *sdp-id* [**keep-alive**]

Context clear>service>statistics

Description Clears keepalive statistics associated with the SDP ID.

Parameters *sdp-id* — 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 *sap-id* — Specifies the SAP ID for which to clear statistics. See Common Service Commands on page 1510 for *sap-id* 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 *sdp-id* — The spoke SDP ID for which to clear statistics.

Values 1 — 17407

vc-id — 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** *sap-id* | **spoke-sdp** *sdp-id*[*:vc-id*]}

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.

sap-id — Clears the specified lease state SAP information. See Common Service Commands on page 1510 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

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*[*/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**]

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 *port-id* **—** The DHCP lease state local specifies that the DHCP lease state is learned by either a SAP or SDP. When the value is SAP, the value indicates the SAP for which this entry contains information.

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 1510 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** *sap-id* | **sdp** *[sdp-id*[*:vc-id*]]

Context clear>service>id>dhcp

Description Clears DHCP statistics for this service.

port-db

Syntax **port-db** {**sap** *sap-id* | **sdp** *sdp-id:vc-id*} [**group** *grp-address* [**source** *ip-address*]]

Context clear>service>id>igmp-snooping

Description Clears the information on the IGMP snooping port database for the VPLS service.

Parameters **sap** *sap-id* — Clears IGMP snooping statistics matching the specified SAP ID and optional encapsulation value. See Common Service Commands on page 1510 for *sap-id* command syntax.

sdp *sdp-id* — 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

vc-id — 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 *grp-address* — 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** *sap-id* | **sdp** *sdp-id:vc-id*}]

Context clear>service>id>igmp-snooping

Description Clears IGMP snooping statistics for the VPLS service.

Parameters **sap** *sap-id* — Clears IGMP snooping statistics matching the specified SAP ID and optional encapsulation value. See Common Service Commands on page 1510 for *sap-id* command syntax.

sdp *sdp-id* — 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

vc-id — 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** *grp-address*}

Context clear>service>id>mfib

Description Clears multicast FIB statistics for the VPLS service.

Parameters *grp-address* — 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** *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*

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 *sap-id*
statistics sdp *sdp-id*:*vc-id*

Context clear>service>id>mld-snooping

Description This command clears MLD snooping statistics.

Parameters **all** — Clears all MLD snooping statistics

sap *sap-id* — Clears all MLD snooping SAP statistics. See Common Service Commands on page 1510 for *sap-id* command syntax.

sdp *sdp-id:vc-id* — Clears all MLD snooping SDP statistics. See Common Service Commands on page 1510 for *sap-id* command syntax.

arp

Syntax **arp** {**all** | *ip-address*}
arp interface [*ip-int-name* | *ip-address*]

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.

ip-address — Clears the ARP cache entry for the specified IP address.

interface *ip-int-name* **—** Clears all ARP cache entries for the IP interface with the specified name.

interface *ip-addr* **—** 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** *ip-int-name* | *ip-address*]

Context clear>router>dhcp

Description Clears DHCP statistics.

Debug Commands

mcast-reporting-dest

Syntax [**no**] **mcast-reporting-dest** [*dest-name*]

Context debug>mcast-mgmt

Description This command debugs multicast path management reporting destinations.

igmp

Syntax [**no**] **igmp** [**host** *ip-address*] [**group** *grp-address*]

Context debug>mcast-mgmt>mcast-rprt-dest

Description This command sets mcast reporting dest debug filtering options.

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** *ieee-address*

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** *sap-id*

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** *sdp-id:vc-id*

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.