VPLS Service Configuration Commands

Generic Commands

shutdown

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

Context config>service>vpls
config>service>vpls>snooping
config>service>vpls>igmp-snooping
config>service>vpls>mac-move
config>service>vpls>gsmp
config>service>vpls>gsmp>group
config>service>vpls>gsmp>group>neighbor
config>service>vpls>interface
config>service>vpls>split-horizon-group
config>service>vpls>sap
config>service>vpls>sap>stp
config>service>vpls>sap>arp-host
config>service>vpls>sap>sub-sla-mgmt
config>service>vpls>mesh-sdp
config>service>vpls>spoke-sdp
config>service>vpls>spoke-sdp>stp
config>service>vpls>stp
config>service>vpls>spoke-sdp>stp
config>service>vpls>mrp
config>service>vpls>sap>dhcp>proxy
config>service>vpls>radius-discovery
config>service>vpls>mesh-sdp>eth-cfm>mep
config>service>vpls>sap>eth-cfm>mep
config>service>vpls>>spoke-sdp>eth-cfm>mep
config>service>vpls>bgp-ad
config>service>vpls>eth-cfm>mep
config>service>vpls>mesh-sdp>eth-cfm>mep
config>service>vpls>>spoke-sdp>eth-cfm>mep

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.g

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>gsmp>group
config>service>vpls>gsmp>group>neighbor
config>service>vpls>igmp-snooping>mvr
config>service>vpls>interface
config>service>vpls>split-horizon-group
config>service>vpls>sap
config>service>vpls>mesh-sdp
config>service>vpls>spoke-sdp
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 *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.

VPLS Service Commands

vpls

Syntax **vpls** *service-id* **customer** *customer-id* **vpn** *vpn-id* [**m-vpls**] [**bvpls | i-vpls**] [**create**]
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.

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.

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 — 2147483648
svc-name: 64 characters maximum

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 management VPLS.

b-vpls | i-vpls — Creates a backbone-vpls or ISID-vpls.

backbone-smac

Syntax **backbone-smac** *ieee-address*

Context config>service>vpls

Description This command configures the backbone source MAC address used for PBB. This command allows a per B-VPLS control of the B-SMAC and the B-Mcast MAC. All I-VPLS provisioned under this B-VPLS will share the provisioned value.

Default backbone-smac address is chassis MAC address

Parameters *ieee-address* — Specifies the backbone source MAC address.

backbone-vpls

Syntax **backbone-vpls** *vpls-id*[:*isid*]
no backbone-vpls

Context config>service>vpls

Description This command associated the I-VPLS with the B-VPLS service. The ISID value is used to mux/demux packets for the VPLS flowing through the B-VPLS.

Parameters *vpls-id* — This value represents the VPLS ID value associated with the B-VPLS.

isid — Defines ISID associated with the I-VPLS.

Default The default is the service-id.

Values 0 — 16777215

stp

Syntax [**no**] **stp**

Context config>service>vpls>backbone-vpls

Description This command enables STP on the backbone VPLS service.

The **no** form of the command disables STP on the backbone VPLS service.

block-on-mesh-failure

Syntax [**no**] **block-on-mesh-failure**

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

Description This command enables blocking (brings the entity to an opererationally down state) after all configured SDPs or endpoints 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

bpdu-translation

Syntax **bpdu-translation** {**auto** | **pvst** | **stp**}
no bpdu-translation

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

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.

calling-station-id

Syntax **calling-station-id** {**mac** | **remote-id** | **sap-id** | **sap-string**}
no calling-station-id

Context config>service>vpls>sap

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.

lag-link-map-profile

Syntax **lag-link-map-profile** *link-map-profile-id*
no lag-link-map-profile

Context config>service>vpls>sap

Description This command assigns a pre-configured lag link map profile to a SAP/network interface configured on a LAG or a PW port that exists on a LAG. Once assigned/de-assigned, the SAP/network interface egress traffic will be re-hashed over LAG as required by the new configuration.

The **no** form of this command reverts the SAP/network interface to use per-flow, service or link hash as configured for the service/LAG.

Default **no lag-link-map-profile**

Parameters *link-map-profile-id* — An integer from 1 to 32 that defines a unique lag link map profile on which the LAG the SAP/network interface exist.

l2pt-termination

Syntax **l2pt-termination** [**cdp**] [**dtp**] [**pagp**] [**stp**] [**udld**] [**vtp**]
no l2pt-termination

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

Description This command enables Layer 2 Protocol Tunneling (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

Parameters **cdp —** Specifies the Cisco discovery protocol.

dtp — Specifies the dynamic trunking protocol.

pagp — Specifies the port aggregation protocol.

stp — Specifies all spanning tree protocols: stp, rstp, mstp, pvst (default).

udld — Specifies unidirectional link detection.

vtp — Specifies the virtual trunk protocol.

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 nodes 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

Parameters *vc-id* — Specifies the default mesh vc-id.

Values 1 — 4294967295

default-gtw

Syntax **default-gtw**

Context config>service>vpls

Description This command configures a service default gateway.

ip

Syntax **ip** *ip-address*
no ip

Context config>service>vpls>defgw

Description This command configures the default gateway IP address.

mac

Syntax **mac** *ieee-address*

Context config>service>vpls>defgw

Description This command configures the default gateway MAC address.

disable-aging

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

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

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 [**no**] **disable-learning**

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

Description This command disables 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 **disable-learning** is 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.

Default no disable-learning (Normal MAC learning is enabled)

discard-unknown

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

Context config>service>vpls
config>template>vpls-template

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** — Packets with unknown destination MAC addresses are flooded.

dist-cpu-protection

Syntax **dist-cpu-protection** *policy-name*
no dist-cpu-protection

Context config>service>vpls>sap

Description This command assigns a Distributed CPU Protection (DCP) policy to the SAP. Only a valid created DCP policy can be assigned to a SAP or a network interface. Note that this rule does not apply to templates such as msap-policy.

Default no dist-cpu-protection

endpoint

Syntax **endpoint** *endpoint-name* [**create**]
no endpoint

Context config>service>vpls

Description This command configures a service endpoint.

Parameters *endpoint-name* — Specifies an endpoint name up to 32 characters in length.

create — This keyword is mandatory while creating a service endpoint.

description

Syntax **description** *description-string*
no description

Context config>service>vpls>endpoint

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 *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.

auto-learn-mac-protect

Syntax [**no**] **auto-learn-mac-protect**

Context config>service>vpls>endpoint
config>service>vpls>mesh-sdp
config>service>vpls>sap
config>service>vpls>split-horizon-group
config>service>vpls>spoke-sdp

Description This command specifies whether to enable automatic population of the MAC protect list with source MAC addresses learned on the associated with this SHG. For more information, refer to Auto-Learn MAC Protect on page 616.

The **no** form of the command disables the automatic population of the MAC protect list.

Default auto-learn-mac-protect

ignore-standby-signaling

Syntax [**no**] **ignore-standby-signaling**

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

Description When this command is enabled, the node will ignore standby-bit received from TLDP peers for the given spoke SDP and performs internal tasks withot taking it into account.

This command is present at endpoint level as well as spoke SDP level. If the spoke SDP is part of the explicit-endpoint, it is not possible to change this setting at the spoke SDP level. The existing spoke SDP will become part of the explicit-endpoint only if the setting is not conflicting. The newly created spoke SDP which is a part of the given explicit-endpoint will inherit this setting from the endpoint configuration.

Default enabled

restrict-protected-src

Syntax **restrict-protected-src alarm-only**
restrict-protected-src [**discard-frame**]
no restrict-protected-src

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

config>service>vpls>split-horizon-group
config>service>vpls>spoke-sdp

This command indicates the action to take whenever a relearn request for a protected MAC is received on a restricted SAP belonging to this SHG

When enabled, the agent will protect the MAC from being learned or re-learned on a SAP that has restricted learning enabled.

Default restrict-protected-src

Parameters **alarm-only —** Specifies that the SAP will be left up and only a notification, sapReceivedProtSrcMac, will be generated.

discard-frame — Specifies that the SAP will start discarding the frame in addition to generating sapReceivedProtSrcMac notification.

revert-time

Syntax **revert-time** *revert-time* **| infinite**
no revert-time

Context config>service>vpls>endpoint

Description This command configures the time to wait before reverting to primary spoke SDP.

In a regular endpoint the revert-time setting affects just the pseudowire defined as primary (precedence 0). For a failure of the primary pseudowire followed by restoration the revert-timer is started. After it expires the primary pseudowire takes the active role in the endpoint. This behavior does not apply for the case when both pseudowires are defined as secondary. For example, if the active secondary pseudowire fails and is restored it will stay in standby until a configuration change or a force command occurs.

Parameters *revert-time* — Specifies the time to wait, in seconds, before reverting back to the primary spoke SDP defined on this service endpoint, after having failed over to a backup spoke SDP.

Values 0 — 600

infinite — Specifying this keywork makes endpoint non-revertive.

static-mac

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

Context config>service>vpls>endpoint

Description This command assigns a static MAC address to the endpoint. In the FDB, the static MAC is then associated with the active spoke SDP.

Default none

Parameters *ieee-address* — Specifies the static MAC address to the endpoint.

Values 6-byte mac-address (xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx). Cannot be all zeros.

create — This keyword is mandatory while creating a static MAC.

suppress-standby-signaling

Syntax [**no**] **suppress-standby-signaling**

Context config>service>vpls>endpoint

Description When this command is enabled, the pseudowire standby bit (value 0x00000020) will not be sent to T-LDP peer when the given spoke is selected as a standby. This allows faster switchover as the traffic will be sent over this SDP and discarded at the blocking side of the connection. This is particularly applicable to multicast traffic.

Default enabled

propagate-mac-flush

Syntax [**no**] **propagate-mac-flush**

Context config>service>vpls

Description This command enabled propagation of mac-flush messages received from the given T-LDP on all spoke and mesh-sdps within the context of the VPLS service. The propagation will follow split-horizon principles and any data-path blocking in order to avoid looping of these messages.

Default disabled

fdb-table-high-wmark

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

Context config>service>vpls
config>template>vpls-template

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

Parameters *high-water-mark* — Specify the value to send logs and traps when the threshold is reached.

Values 0— 100

Default 95%

fdb-table-low-wmark

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

Context config>service>vpls
config>template>vpls-template

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

Parameters *low-water-mark* — Specify the value to send logs and traps when the threshold is reached.

Values 0— 100

Default 90%

fdb-table-size

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

Context config>service>vpls
config>template>vpls-template

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* — Specifies the maximum number of MAC entries in the FDB.

Values 1 — 511999
Chassis-mode A or B limit: 131071
Chassis-mode C limit: 196607
Chassis-mode D limit: 511999

interface

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

Context config>service>vpls

Description This command creates an IP interface.

address

Syntax **address** *ip-address*[*/mask*]> [*netmask*]
no address

Context config>service>vpls>interface

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

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

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

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

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

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

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

The operational state is a read-only variable and the only controlling variables are the address and admin states. The address and admin states are independent and can be set independently. If an interface is in an 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 netmask

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.

arp-timeout

Syntax **arp-timeout** *seconds*

no arp-timeout

Context config>service>vpls>interface

Description This command configures the minimum time in seconds an ARP entry learned on the IP interface will be stored in the ARP table. ARP entries are automatically refreshed when an ARP request or gratuitous ARP is seen from an IP host, otherwise, the ARP entry is aged from the ARP table. If **arp-timeout** is set to a value of zero seconds, ARP aging is disabled.

When the **arp-populate** and **lease-populate** commands are enabled on an 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

mac

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

Context config>service>vpls>interface

Description This command assigns a specific MAC address to a VPLS IP interface.

For Routed Central Office (CO), a group interface has no IP address explicitely configured but inherits an address from the parent subscriber interface when needed. For example, a MAC will respond to an ARP request when an ARP is requested for one of the IPs associated with the subscriber interface through the group interface.

The **no** form of the command returns the MAC address of the IP interface to the default value.

Default The system chassis MAC address.

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.

static-arp

Syntax **static-arp** *ieee-mac-addr unnumbered*
no static-arp *unnumbered*

Context config>service>vpls>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. A static ARP can only be configured if it exists on the network attached to the IP interface.

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

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

Default None

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

unnumbered — Specifies the static ARP MAC for an unnumbered interface. Unnumbered interfaces support dynamic ARP. Once this command is configured, it overrides any dynamic ARP.

static-mac

Syntax **static-mac**

Context config>service>vpls

Description A set of conditional Static MAC addresses can be created within a B-VPLS VPLS supporting SPBM. Conditional Static MACs are dependent on the SAP/SDP state.

This command allows assignment of a set of conditional static MAC addresses to a SPBM SAP/spoke-SDP. In the FDB, the static MAC is then associated with the active SAP or spoke SDP.

Static MACs are used for PBB Epipe and I-VPLS services that may terminate external to SPBM. If this is configured under a Control B-VPLS the interface referenced will not use IS-IS for this neighbor. This may also be configured under a User B-VPLS where the corresponding interface is not supported under the Control B-VPLS.

mac

Syntax **mac ieee-address [create] sap** *sap-id* **[monitor fwd-status]**
mac ieee-address [create] spoke-sdp *sdp-id:vc-id***] [monitor fwd-status]**
[no] mac ieee-address

Context config>service>vpls>static-mac

Description This command assigns a conditional static MAC address entry to an SPBM B-VPLS SAP/spoke-SDP allowing external MACs for single and multi-homed operation.

Static MACs are used for PBB Epipe and I-VPLS services that may terminate external to SPBM. If this is configured under a Control B-VPLS the interface referenced will not use IS-IS for this neighbor. This may also be configured under a User B-VPLS where the corresponding interface is not supported under the Control B-VPLS.

Default none

Parameters **ieee-address —** Specifies the static MAC address to an SPBM interface.

Values 6-byte mac-address (xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx). Cannot be all zeros.

create — This keyword is mandatory while creating a static MAC.

monitor fwd-status — Specifies that this static mac is based on the forwarding status of the SAP or spoke SDP for multi-homed operation. Monitoring is optional but is required for multi-homing.

unnumbered

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

Context config>service>ies>if
config>service>vpls>if
config>service>vprn>if

Description This command configures the interface as an unnumbered interface. Unnumbered IP interface is supported on a Sonet/SDH access port with the PPP, ATM, or Frame Relay encapsulation. It is also supported on an Ethernet port. It is not supported on a TDM port or channel.

Parameters *ip-int-name* — Specifies the name of the IP interface. If the string ocntains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

ip-address — Specifies an IP address which must be a valid address of another interface.

isid-policy

Syntax **isid-policy**
no isid-policy

Context config>service>vpls

Description This command configures isid-policies for individual ISIDs or ISID ranges in a B-VPLS using SPBM. The ISIDs may belong to I-VPLS services or may be static-isids defined on this node. Multiple entry statements are allowed under a **isid-policy**. ISIDs that are declared as static do not require and **isid-policy** unless the ISIDs are not to be advertised.

isid-policy allows finer control of ISID multicast but is not typically required for SPBM operation. Use of ISID policies can cause additional flooding of multicast traffic.

Default no default

entry

entry *id* **create**
no entry

Context config>service>vpls>isid-policy

Description This command creates or edits an isid-policy entry. Multiple entries can be created using unique entry-id numbers within the isid-policy.

Default: No entry

entry-id — An entry-id uniquely identifies a ISID range and the corresponding actions. This allows users to insert a new entry in an existing policy without requiring renumbering of all the existing entries.

The following rules govern the usage of multiple entry statements:

- overlapping values are allowed:
  - isid from 301 to 310
  - isid from 305 to 315
  - isid 316
- the minimum and maximum values from overlapping ranges are considered and displayed. The above entries will be equivalent with “isid from 301 to 316” statement.
- there is no consistency check with the content of ISID statements from other entries. The entries will be evaluated in the order of their IDs and the first match will cause the implementation t o execute the associated action for that entry.

no isid - removes all the previous statements under one entry.

no isid value | from value to higher-value - removes a specific ISID value or range. Must match a previously used positive statement: for example, if the command “isid 16 to 100” was used using “no isid 16 to 50”, it will not work but “no isid 16 to 100 will be successful.

Values 1-65535

create — Required when first creating the configuration context. Once the context is created, one can navigate into the context without the create keyword.

advertise-local

Syntax **[no] advertise-local**

Context config>service>vpls>isid-policy>entry

Description The **no advertise-local** option prevents the advertisement of any locally defined I-VPLS ISIDs or static-isids in the range in a B-VPLS. For I-VPLS services or static-isids that are primarily unicast traffic, the use-def-mcast and no advertise-local options allows the forwarding of ISID based multicast frames locally using the default multicast. The **no advertise-local** option also suppresses this range of ISIDs from being advertised in ISIS. When using the **use-def-mcast** and **no advertise-local** policies, the ISIDs configured under this **static-isid** declarations SPBM treats the ISIDs as belonging to the default tree.

Default advertise-local

range

Syntax **range** *isid* **[to** *isid***]**

Context config>service>vpls>isid-policy>entry

Description This command specifies an ISID or a Range of ISIDs in a B-VPLS. One range is allowed per entry.

Default no range

Parameters *isid* — Specifies the ISID value in 24 bits. When singular, ISID identifies a particular ISID to be used for matching.

Values 0..16777215

to *isid* **—** Identifies upper value in a range of ISIDs to be used as matching criteria.

use-def-mcast

Syntax **[no] use-def-mcast**

Context config>service>vpls>isid-policy>entry

Description The **use-def-mcast** option prevents local installation of the ISIDs in the range in the MFIB and uses the default multicast tree instead for a B-VPLS. In a node that does not have I-VPLS or static-isids, this command prevents the building of an MFIB entry for this ISID when received in a SPBM TLV and allows the broadcast of ISID based traffic on the default multicast tree. If an **isid-policy** exists, the core nodes can have this policy to prevent connectivity problems when some nodes are advertizing an ISID and others are not. In a I-VPLS service if the customer MAC (C-MAC) is unknown, a frame will have the Multicast DA for an ISID (PBB-OUI + ISID) flooded on the default multicast tree and not pruned.

Default no use-def-mcast

local-age

Syntax **local-age** *aging-timer*
no local-age

Context config>service>vpls
config>template>vpls-template

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 *aging-timer* — The aging time for local MACs expressed in seconds.

Values 60 — 86400

mac-move

Syntax [**no**] **mac-move**

Context config>service>vpls
config>template>vpls-template

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. 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.

mac-move will disable a VPLS port when the number of relearns detected has reached the number of relearns needed to reach the move-frequency in the 5-second interval. For example, when the move-frequency is configured to 1 (relearn per second) mac-move will disable one of the VPLS ports when 5 relearns were detected during the 5-second interval because then the average move-frequency of 1 relearn per second has been reached. This can already occur in the first second if the real relearn rate is 5 relearns per second or higher.

The **no** form of this command disables MAC move.

mac-protect

Syntax **mac-protect**

Context config>service>vpls

Description This command indicates whether or not this MAC is protected on the MAC protect list. When enabled, the agent will protect the MAC from being learned or re-learned on a SAP, spoke SDP or mesh-SDP that has restricted learning enabled. The MAC protect list is used in conjunction with **restrict-protected-src**, **restrict-unprotected-dst** and **auto-learn-mac-protect**.

Default disabled

mac

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

Context config>service>vpls>mac-protect

Description This command adds a protected MAC address entry.

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

move-frequency

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

Context config>service>vpls>mac-move
config>template>vpls-template>mac-move

Description This command 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 **no** form of the command reverts to the default value.

Default 2 (when mac-move is enabled). For example, 10 relearns in a 5 second period.

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

Values 1 — 100

number-retries

Syntax **number-retries** *number-retries*
no number-retries

Context config>service>vpls>mac-move
config>template>vpls-template>mac-move

Description This command configures the number of times retries are performed for reenabling the SAP/SDP.

Parameters *number-retries*

Specifies number of retries for reenabling the SAP/SDP. A zero (0) value indicates unlimited number of retries.

Values 0 — 255

primary-ports

Syntax **primary-ports**

Context config>service>vpls>mac-move
config>template>vpls-template>mac-move

Description This command enables the context to define primary VPLS ports. VPLS ports that were declared as secondary prior to the execution of this command will be moved from secondary port-level to primary port-level. Changing a port to the tertiary level can only be done by first removing it from the secondary port-level.

cumulative-factor

Syntax **cumulative-factor** *cumulative-factor*
no cumulative-factor

Context configure->service->vpls->mac-move->primary-ports
configure->service->vpls->mac-move->secondary-ports
config>template>vpls-template>mac-move>primary-ports
config>template>vpls-template>mac-move>secondary-ports

Description This command configures a factor for the primary or secondary ports defining how many MAC relearn periods should be used to measure the MAC relearn rate . This rate must be exceeded during consecutive periods before the corresponding ports (SAP and/or spoke-SDP) are blocked by the MAC-move feature.

Parameters *cumulative-factor* — Specifies a MAC relearn period to be used for MAC relearn rate.

Values 3 — 10

sap

Syntax **sap** [**split-horizon-group** *group-name*] [**create**] [**capture-sap**]
no sap *sap-id*

Context config>service>vpls>mac-move>primary-ports
config>service>vpls>mac-move>secondary-ports

Description This command declares a given SAP as a primary (or secondary) VPLS port.

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition. See Common CLI Command Descriptions on page 2569 for command syntax.

spoke-sdp

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

Context config>service>vpls>mac-move>primary-ports
config>service>vpls>mac-move>seocndary-ports

Description This command declares a given spoke SDP as a primary (or secondary) VPLS port.

Parameters *spoke-id* — Specifies the SDP ID to configure as the primary VPLS port.

Values 1 — 17407

vc-id — The virtual circuit identifier.

Values 1 — 4294967295

cumulative-factor

Syntax [**no**] **cumulative-factor** *factor*

Context config>service>vpls>mac-move>primary-ports
config>service>vpls>mac-move>seocndary-ports

Description This command defines a factor defining how many mac-relearn measurement periods can be used to measure mac-relearn rate. The rate must be exceeded during the defined number of consecutive periods before the corresponding port is blocked by the mac-move feature. The cumulative-factor of primary ports must be higher than cumulative-factor of secondary ports.

Default 2 — secondary ports

3 — primary ports

Parameters *factor* — Specifies the factor defining the number of mac-relearn measurement periods can be used to measure mac-relearn rate.

Values 2 — 10

secondary-ports

Syntax **secondary-ports**

Context config>service>vpls>mac-move
config>template>vpls-template>mac-move

Description This command opens configuration context for defining secondary vpls-ports. VPLS ports that were declared as primary prior to the execution of this command will be moved from primary port-level to secondary port-level. Changing a port to the tertiary level can only be done by first removing it from the primary port-level.

retry-timeout

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

Context config>service>vpls>mac-move
config>template>vpls-template>mac-move

Description This indicates the time in seconds to wait before a SAP that has been disabled after exceeding the maximum relearn rate is reenabled.

It is recommended that the retry-timeout value is larger or equal to 5s * cumulative factor of the highest priority port so that the sequential order of port blocking will not be disturbed by re-initializing lower priority ports.

A zero value indicates that the SAP will not be automatically re-enabled after being disabled. If, after the SAP is reenabled it is disabled again, the effective retry timeout is doubled in order to avoid thrashing.

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

Default 10 (when mac-move is enabled)

Parameters *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

mld-snooping

Syntax **mld-snooping**

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

Description This command configures MLD snooping parameters.

remote-age

Syntax **remote-age** *seconds*
no remote-age

Context config>service>vpls
config>template>vpls-template

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

send-bvpls-flush

Syntax **send-bvpls-flush** {[**all-but-mine**] [**all-from-me**]}
no send-bvpls-flush

Context config>service>vpls

Description This command enables generation of LDP MAC withdrawl "flush-all-from-me" in the B-VPLS domain when the following triggers occur in the related IVPLS:

- MC-LAG failure
- Failure of a local SAP
- Failure of a local pseudowire/SDP binding

Note that failure means transition of link SAP/pseudowire to either down or standby status.

This command does not require send-flush-on-failure in B-VPLS to be enabled on an IVPLS trigger to send an MAC flush into the BVPLS.

Default no send-bvpls-flush

Parameters **all-but-mine —** Specifies to send an LDP flush all-but-mine and also sent into the B-VPLS. Note that both parameters can be set together.

all-from-me — Specifies to send an LDP flush-all-from and when STP initiates a flush, it is sent into the B-VPLS using LDP MAC flush all-from-me. Note that both parameters can be set together.

send-flush-on-bvpls-failure

Syntax [**no**] **send-flush-on-bvpls-failure**

Context config>service>vpls ivpls

Description This command enables the generation in the local I-VPLS of a LDP MAC flush-all-from-me following a failure of SAP/the whole endpoint/spoke-SDP in the related B-VPLS. Note that the failure of mesh-SDP in B-VPLS does not generate the I-VPLS MAC flush.

The **no** form of this command disables the generation of LDP MAC flush in I-VPLS on failure of SAP/endpoint/spoke-SDP in the related B-VPLS.

Default no send-flush-on-bvpls-failure

propagate-mac-flush-from-bvpls

Syntax [**no**] **propagate-mac-flush-from-bvpls**

Context config>service>vpls ivpls

Description This command enables the propagation in the local I-VPLS of any regular LDP MAC Flush received in the related B-VPLS. If an LDP MAC flush-all-but-mine is received in the B-VPLS context, the command controls also whether a flush is performed for all the customer MACs in the associated I-VPLS FIB. The command does not have any effect on a PBB MAC Flush (LDP MAC flush with PBB TLV) received in the related B-VPLS context.

The **no** form of this command disables the propagation of LDP MAC Flush in I-VPLS from the related B-VPLS.

Default no propagate-mac-flush-from-bvpls

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. If the endpoint is configured within the VPLS and send-flush-on-failure is enabled, flush-all-from-me messages will be sent out only when all spoke SDPs associated with the endpoint go down.

This feature cannot be enabled on management VPLS.

Default no send-flush-on-failure

service-mtu

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

Context config>service>vpls
config>template>vpls-template

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 (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.

For i-VPLS and Epipes bound to a b-VPLS, the service-mtu must be at least 18 bytes smaller than the b-VPLS service MTU to accomodate the PBB header.

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

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

Values 1 — 9194

service-name

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

Context config>service>vpls

Description This command configures an optional service name, up to 64 characters in length, which adds a name identifier to a given service to then use that service name in configuration references as well as display and use service names in show commands throughout the system. This helps the service provider/administrator to identify and manage services within the 7750 SR, 7450 ESS and 7710 SRplatforms.

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).

allow-ip-int-binding

Syntax [**no**] **allow-ip-int-binding**

Context config>service>vpls

Description The allow-ip-int-binding command that sets a flag on the VPLS or I-VPLS service that enables the ability to attach an IES or VPRN IP interface to the VPLS service in order to make the VPLS service

routable. When the allow-ip-int-binding command is not enabled, the VPLS service cannot be attached to an IP interface.

VPLS Configuration Constraints for Enabling allow-ip-int-binding

When attempting to set the allow-ip-int-binding VPLS flag, the system first checks to see if the correct configuration constraints exist for the VPLS service and the network ports. In Release 8.0 the following VPLS features must be disabled or not configured for the allow-ip-int-binding flag to set:

- SAP ingress QoS policies applied to the VPLS SAPs cannot have MAC match criteria defined
- SDPs used in spoke or mesh SDP bindings cannot be configured as GRE
- The VPLS service type cannot be B-VPLS or M-VPLS, and it cannot be an I-VPLS service bound to a B-VPLS context
- MVR from Routed VPLS and to another SAP is not supported
- Enhanced and Basic Subscriber Management (ESM and BSM) features
- Network domain on SDP bindings

Once the VPLS allow-ip-int-binding flag is set on a VPLS service, the above features cannot be enabled on the VPLS service.

NETWORK PORT HARDWARE CONSTRAINTS

The system also checks to ensure that all ports configured in network mode are associated with FlexPath2 forwarding planes. If a port is currently in network mode and the port is associated with a FlexPath1 forwarding plane, the allow-ip-int-binding command will fail. Once the allow-ip-int-binding flag is set on any VPLS service, attempting to enable network mode on a port associated with a FlexPath1 forwarding plane will fail.

VPLS SAP HARDWARE CONSTRAINTS

Besides VPLS configuration and network port hardware association, the system also checks to that all SAPs within the VPLS are created on Ethernet ports and the ports are associated with FlexPath2 forwarding planes. Certain Ethernet ports and virtual Ethernet ports are not supported which include HSMDA ports and CCAG virtual ports (VSM based). If a SAP in the VPLS exists on an unsupported port type or is associated with a FlexPath1 forwarding plane, the allow-ip-int-binding command will fail. Once the allow-ip-int-binding flag is set on the VPLS service, attempting to create a VPLS SAP on the wrong port type or associated with a FlexPath1 forwarding plane will fail.

VPLS SERVICE NAME BOUND TO IP INTERFACE WITHOUT ALLOW-IP-INT-BINDING FLAG SET

In the event that a service name is applied to a VPLS service and that service name is also bound to an IP interface but the allow-ip-int-binding flag has not been set on the VPLS service context, the system attempt to resolve the service name between the VPLS service and the IP interface will fail. After the allow-ip-int-binding flag is successfully set on the VPLS service, either the service name on the VPLS service must be removed and reapplied or the IP interface must be re-initialized using the shutdown / no shutdown commands. This will cause the system to reattempt the name resolution process between the IP interface and the VPLS service.

The **no** form of the command resets the allow-ip-int-binding flag on the VPLS service. If the VPLS service currently has an IP interface from an IES or VPRN service attached, the no allow-ip-int-binding command will fail. Once the allow-ip-int-binding flag is reset on the VPLS service, the configuration and hardware restrictions associated with setting the flag are removed. The port network mode hardware restrictions are also removed.

site

Syntax **site** *name* [**create**]
no site *name*

Context config>service>vpls

Description This command configures a VPLS site.

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

Parameters *name* — Specifies a site name up to 32 characters in length.

create — This keyword is mandatory while creating a VPLS service.

boot-timer

Syntax **boot-timer** *seconds*
no boot-timer

Context config>service>vpls>site

Description This command configures for how long the service manger waits after a node reboot before running the DF election algorithm. The boot-timer value should be configured to allow for the BGP sessions to come up and for the NLRI information to be refreshed/exchanged.

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

Default 10

Parameters *seconds* — Specifies the site boot-timer in seconds.

Values 0 — 100

failed-threshold

Syntax **failed-threshold** [1..1000]
failed-threshold all

Context config>service>vpls>site

Description This command defines the number of objects should be down for the site to be declared down. Both administrative and operational status must be evaluated and if at least one is down, the related object is declared down.

Default failed-threshold all

Parameters **1** .. **1000 —** Specifies the threshold for the site to be declared down.

mesh-sdp-binding

Syntax [**no**] **mesh-sdp-binding**

Context config>service>vpls>site

Description This command enables applications to all mesh SDPs.

The **no** form of reverts the default.

Default no mesh-sdp-binding

monitor-oper-group

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

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

Description This command specifies the operational group to be monitored by the object under which it is configured. The **oper-group** *name* must be already configured under the **config>service** context before its name is referenced in this command.

The **no** form of the command removes the association.

sap

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

Context config>service>vpls>site

Description This command configures a SAP for the site.

The **no** form of the command removes the SAP ID from the configuration.

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition. See Common CLI Command Descriptions on page 2569 for command syntax.

site-activation-timer

Syntax **site-activation-timer** *seconds*
no site-activation-timer

Context config>service>vpls>site

Description This command configures the time-period the system keeps the local sites in standby status, waiting for BGP updates from remote PEs before running the DF (designated-forwarder) election algorithm

to decide whether the site should be unblocked. This timer if terminated if an update is received for which the remote PE has transitioned from DF to non-DF.

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

Default 2

Parameters *seconds* — Specifies the site activation timer in seconds.

Values 0 — 100

site-id

Syntax **site-id** *value*
no site-id

Context config>service>vpls>site

Description This command configures the identifier for the site in this service.

Parameters *value* — Specifies the site identifier.

Values 1 — 65535

split-horizon-group

Syntax **split-horizon-group** *group-name*
no split-horizon-group

Context config>service>vpls>site

Description This command configures the value of split-horizon group associated with this site.

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

Default no split-horizon-group

Parameters *group-name* — Specifies a split-horizon group name.

spoke-sdp

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

Context config>service>vpls>site

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

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

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. Half are supported in i-VPLS.

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 (cannot be enabled)
- ARP reply agent enabled per default (can be disabled)
- MAC pinning enabled per default (can be disabled)
- Downstream broadcast packets are discarded thus also blocking the unknown, flooded traffic
- Downstream multicast packets are allowed when IGMP snooping is enabled

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.

pppoe-policy

Syntax **pppoe-policy** *pppoe-policy-name*
no pppoe-policy

Context config>service>vpls>sap

Description This command specifies an existing PPPoE policy. These policies are referenced from interfaces configured for PPPoE. Multiple PPPoE policies may be configured.

Default none

Parameters *pppoe-policy-name* — Specifies an existing PPPoE policy name up to 32 characters in length.

auto-learn-mac-protect

Syntax [**no**] **auto-learn-mac-protect**

Context config>service>vpls>sap
config>service>vpls>spoke-sdp
config>service>vpls >mesh-sdp
config>service>vpls>split-horizon-group
config>service>vpls>endpoint
config>service>pw-template
config>service>pw-template>split-horizon-group

Description This command enables the automatic protection of source MAC addresses learned on the associated object. MAC protection is used in conjunction with restrict-protected-src, restrict-unprotected-dst and mac-protect. When this command is applied or removed, the MAC addresses are cleared from the related object.

When the auto-learn-mac-protect is enabled on an SHG the action only applies to the associated SAPs (no action is taken by default for spoke SDPs in the SHG). In order to enable this function for spoke SDPs within a SHG, the auto-learn-mac-protect must be enabled explicitly under the spoke-SDP. If required, auto-learn-mac-protect can also be enabled explicitly under specific SAPs within the SHG.

Default no auto-learn-mac-protect

restrict-protected-src

Syntax **restrict-protected-src** [*alarm-only* | *discard-frame*]
no restrict-protected-src

Context config>service>vpls>sap
config>service>vpls>spoke-sdp
config>service>vpls>mesh-sdp
config>service>vpls>split-horizon-group
config>service>vpls>endpoint
config>service>pw-template>
config>service>pw-template>split-horizon-group

Description This command indicates how the agent will handle relearn requests for protected MAC addresses, either manually added using the mac-protect command or automatically added using the auto-learn-mac-protect command. While enabled all packets entering the configured SAP, spoke-SDP, mesh-SDP , or any SAP that is part of the configured split horizon group (SHG) will be verified not to contain a protected source MAC address. If the packet is found to contain such an address, the action taken depends on the parameter specified on the restrict-protected-src command, namely:

- No parameter

The packet will be discarded, an alarm will be generated and the SAP, spoke-SDP or mesh-SDP will be set operationally down. The SAP, spoke-SDP or mesh-SDP must be shutdown and enabled (no shutdown) for this state to be cleared.

- alarm-only

  The packet will be forwarded, an alarm will be generated but the source MAC is not learned on the SAP/spoke-SDP/mesh-SDP.

- discard-frame

  The packet will be discarded and an alarm generated. The frequency of alarm generation is fixed to be at most one alarm per MAC address per FP2 per 10 minutes in a given VPLS service. This parameter is only applicable to automatically protected MAC addresses.

When the **restrict-protected-src** is enabled on an SHG the action only applies to the associated SAPs (no action is taken by default for spoke SDPs in the SHG). In order to enable this function for spoke SDPs within a SHG, the **restrict-protected-src** must be enabled explicitly under the spoke-SDP. If required, **restrict-protected-src** can also be enabled explicitly under specific SAPs within the SHG.

When this command is applied or removed, with either the alarm-only or discard-frame parameters, the MAC addresses are cleared from the related object.

The use of "**restrict-protected-src discard-frame**" is mutually exclusive with both the "**restrict-protected-src** [**alarm-only**]" command and with the configuration of manually protected MAC addresses within a given VPLS. "restrict-protected-src discard-frame" can only be enabled on SAPs on FP2 or later hardware or on SDPs where all network interfaces are on FP2 or later hardware.

Parameters *alarm-only* — Specifies that the packet will be forwarded, an alarm will be generated but the source MAC is not learned on the SAP/spoke-SDP/mesh-SDP.

Default no alarm-only

discard-frame — Specifies that the packet will be discarded and an alarm generated. The frequency of alarm generation is fixed to be at most one alarm per FP2 per MAC address per 10 minutes within a given VPLS service.

Default no discard-frame

Default no restrict-protected-src

restrict-unprotected-dst

Syntax **restrict-unprotected-dst**
no restrict-unprotected-dst

Context config>service>pw-template>split-horizon-group
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, either manually added using the mac-protect command or automatically added using the auto-learn-mac-protect command. While enabled all packets entering the configured SAP or SAPs within a split-horizon-group (but not spoke or mesh-SDPs) will be verified to contain a protected destination MAC address. If the packet is found to contain a non-protected destination MAC, it will

be discarded. Detecting a non-protected destination MAC on the SAP will not cause the SAP to be placed in the operationally down state. No alarms are generated.

If the destination MAC address is unknown, even if the packet is entering a restricted SAP, with restrict-unprotected-dst enabled, it will be flooded.

Default no restrict-unprotected-dst

vpls-group

Syntax [**no**] **vpls-group** *id*

Context config>service>vpls

Description This command defines a vpls-group index. Multiple vpls-group commands can be specified to allow the use of different VPLS and SAP templates for different ranges of service ids. A vpls-group can be deleted only in shutdown state. Multiple commands under different vpls-group ids can be issued and can be in progress at the same time.

Default no vpls-group

Parameters *id* — Specifies the ID associated with the VPLS group.

Values 1 — 4094

service-range

Syntax **service-range** *startid-endid* [**start-vlan-id** *startvid*]
no service-range *startid-endid*

Context config>service>vpls>vpls-group

Description This command configures the service ID and implicitly the VLAN-ID ranges to be used as input variables for related VPLS and SAP templates to pre-provision "data" VPLS instances and related SAPs using the service ID specified in the command. If the start-vlan-id is not specified then the service-range values are used for vlan-ids. The data SAPs will be instantiated on all the ports used to specify SAP instances under the related control VPLS.

Modifications of the service id and vlan ranges are allowed with the following restrictions.

- service-range increase can be achieved in two ways:
  - Allowed when vpls-group is in shutdown state
  - By creating a new vpls-group
- service-range decrease can be achieved in two ways:
  - Allowed when vpls-group is in shutdown state; when shutdown command is executed the associated service instances are deleted.
  - Allowed when vpls-group is in no shutdown state and has completed successfully instantiating services.

- Note that in both cases only the services that do not have user configured SAPs will be deleted. Otherwise the above commands are rejected. Existing declarations or registrations do not prevent service deletion.

- start-vlan-id change can be achieved in two ways:
  - Allowed when vpls-group is in shutdown state
  - At the time of range decrease by increasing the start-vlan-id which can be done when vpls-group is in no shutdown state and has completed successfully instantiating services

The **no** form of this command removes the specified ranges and deletes the pre-provisioned VPLS instances and related SAPs. The command will fail if any of the VPLS instances in the affected ranges have a provisioned SAP.

Default no service-range

Parameters *startid-endid* — Specifies the range of service IDs.

Values 1—2147483647

startvid — Specifies the starting VLAN ID; it provides a way to set aside a service ID range that is not the same as the VLAN range and allows for multiple MVRP control-VPLSes to control same VLAN range on different ports.

Values 1—4094

vpls-template-binding

Syntax **vpls-template-binding** *name/id*
no vpls-template-binding

Context config>service>vpls>vpls-group

Description This command configures the binding to a VPLS template to be used to instantiate pre-provisioned data VPLS using as input variables the service IDs generated by the vid-range command.

The **no** form of this command removes the binding and deletes the related VPLS instances. The command will fail if any of the affected VPLS instances have either a provisioned SAP or an active MVRP declaration/registration or if the related vpls-group id is in no shutdown state. Any changes to the vpls-template-binding require the vpls-group to be in shutdown state.

Default no vpls-template-binding

Parameters *name/id* — Specifies the name or the ID of the VPLS template.

Values 1—1024

vpls-sap-template-binding

Syntax **vpls-sap-template-binding** *name/id*
no vpls-sap-template-binding

Context config>service>vpls>vpls-group

Description This command configures the binding to a SAP template to be used to instantiate SAPs in the data VPLS using as input variables the VLAN IDs generated by the vid-range command.

The **no** form of this command removes the binding and deletes the related SAP instances. The command will fail if any of the affected VPLS instances have either a provisioned SAP or an active MVRP declaration/registration registration or if the related vpls-group is in no shutdown state. Any changes to the vpls-sap-template-binding require the vpls-group to be in shutdown state. New control SAP additions to the management VPLS are allowed as long as data VPLS instantiations/removals for vpls-groups are not in progress. Control SAPs can be removed at any time generating the removal of related data SAPs from the data VPLS. The shutdown or no shutdown state for the control SAPs does not have any effect on data SAPs instantiated with this command.

Default no vpls-sap-template-binding

Parameters *name* — Specifies the name of the VPLS template.

Values ASCII character string

id — Specifies the ID of the VPLS template

Values 1—8196

mvrp-control

Syntax [**no**] **mvrp-control**

Context config>service>vpls>vpls-group

Description This command enables MVRP control in the VPLS instances instantiated using the templates for the specified vpls-group. That means the flooding FIB will be created empty and will be populated with endpoints whenever MVRP receives a declaration and a registration on a specific endpoint. Also the VLAN ID associated by the control VPLS with the instantiated VPLS will be declared on service activation by MVRP on all virtual MVRP ports in the control VPLS. Service activation takes place when at least one other SAP is provisioned and brought up under the data VPLS. This is usually a customer facing SAP or a SAP leading outside of the MVRP controlled domain.

The **no** form of this command disallows MVRP control over this VPLS. The VPLS will be created with a regular FIB and will become as a result active upon creation time. Command change is allowed only when the related vpls-group is in shutdown state.

Default no mvrp-control

mvrp

Syntax **mvrp**

Context config>service>vpls>mrp
config>service>vpls>sap>mrp

Description This object consolidates the MVRP attributes. MVRP is only supported initially in the management VPLS so the object is not supported under BVPLS, IVPLS or regular VPLS not marked with the m-vpls tag.

hold-time

Syntax **hold-time** *value*
no hold-time

Context config>service>vpls>mrp>mvrp

Description This command enables the dampening timer and applies to both types of provisioned SAPs – end-station and UNI. When a value is configured for the timer, it controls the delay between detecting that the last provisioned SAP in VPLS goes down and reporting it to the MVRP module. The CPM will wait for the time specified in the value parameter before reporting it to the MVRP module. If the SAP comes up before the hold-timer expires, the event will not be reported to MVRP module.

The non-zero hold-time does not apply for SAP transition from down to up, This kind of transition is reported immediately to MVRP module without waiting for hold-time expiration. Also this parameter applies only to the provisioned SAPs. It does NOT apply to the SAPs configured with the vpls-sap-template command. Also when endstation QinQ SAPs are present only the "no hold-time" configuration is allowed.

The **no** form of this command disables tracking of the operational status for the last active SAP in the VPLS. MVRP will stop declaring the VLAN only when the last provisioned customer (UNI) SAP associated locally with the service is deleted. Also MVRP will declare the associated VLAN attribute as soon as the first provisioned SAP is created in the associated VPLS instance, regardless of the operational state of the SAP.

Default no hold-time

Parameters *value* — Specifies the hold time in minutes

Values 1—30 minutes

endstation-vid-group

Syntax **endstation-vid-group** *id* **vlan-id** *startvid-endvid*
no endstation-vid-group *id*

Context config>service>vpls>mrp>mvrp

Description This command specifies the range of VLAN IDs that are controlled by MVRP on the port associated with the parent SAP. When the command is present under a certain SAP, the MVRP will treat the associated virtual port as an endstation.

MVRP endstation behavior means that configuration of a new data SAP with the outer tag in the configured endstation-vid-group will generate down that virtual port a MVRP declaration for the new [outer] VLAN attribute. Also registration received for the VLAN attribute in the range will be accepted but not propagated in the rest of MVRP context.

Note that VPLS-groups are not allowed under the associated Management VPLS (MVPLS) once the endstation is configured under one SAP. VPLS-groups can be supported in the chassis using a different MVPLS.

The **no** form of the command removes the specified group id.

Default no endstation-vid-group

Parameters *id* — Specifies the range index.

Values 1—4094

starvid-endvid — Specifies the range of VLANs to be controlled by MVRP.

Values 1—4094

root-guard

Syntax [**no**] **root-guard**

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

static-host

Syntax **static-host ip** *ip-address* [**mac** *ieee-address*] [**create**]
static-host mac *ieee-address* [**create**]
no static-host [**ip** *ip-address*>] **mac** *ieee-address*>
no static-host all [**force**]
no static-host ip *ip-address*

Context config>service>vpls>sap

Description This command configures a static host on this SAP.

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

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

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

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

create — This keyword is mandatory while configuring a static host.

ancp-string

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

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

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>vpls>sap>static-host

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>vpls>sap>static-host

Description Specifies to which intermediate destination (for example a DSLAM) this host belongs.

Parameters *intermediate-destination-id* — Specifies the intermediate destination ID.

sla-profile

Syntax **sla-profile** *sla-profile-name*
no sla-profile

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

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

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

tod-suite

Syntax **tod-suite** *tod-suite-name*
no tod-suite

Context config>service>vpls>sap

Description This command applies a time-based policy (filter or QoS policy) to the service SAP. The suite name must already exist in the **config>cron** context.

Default no tod-suite

Parameters *tod-suite-name* — Specifies collection of policies (ACLs, QoS) including time-ranges that define the full or partial behavior of a SAP. The suite can be applied to more than one SAP.

trigger-packet

Syntax **trigger-packet** [**dhcp**] [**pppoe**] [**arp**] [**dhcp6**] [**ppp**]
no trigger-packet

Context config>service>vpls>sap

Description This command enables triggering packet to initiate RADIUS authentication that provides a service context. The authentication, together with the service context for this request, creates a managed SAP. The VLAN is the same as the triggering packet. This SAP behaves as a regular SAP but the configuration is not user-editable and not maintained in the configuration file. The managed SAP remains active as long as the session is active.

Default none

Parameters **dhcp —** Specifies whether the receipt of DHCP trigger packets on this VPLS SAP when the keyword **capture-sap** is specified in the **sap** command creation string, will result in a RADIUS authentication that will provide a service context and the creation of a SAP with a value of 'managed'.

pppoe — Specifies whether the receipt of PPPoE trigger packets on this VPLS SAP when the keyword **capture-sap** is specified in the **sap** command creation string, will result in a RADIUS authentication that will provide a service context and the creation of a SAP with a value of 'managed'.

arp — Indicates that ARP is the type of trigger packets for this entry.

dhcp6 — Indicates that DHCP6 is the type of trigger packets for this entry.

ppp — Indicates that PPP is the type of trigger packets for this entry.

VPLS Interface Commands

interface

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

Context config>service>vpls

Description This command creates a logical IP routing interface for a VPLS service. Once created, attributes such as IP address and service access points (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 the VPLS service IDs. The IP interface created is associated with the VPLS management routing instance.This instance does not support routing.

Interface names are case-sensitive and must be unique within the group of defined IP interfaces defined for the network core router instance. Interface names in the dotted decimal notation of an IP address are not allowed. 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. Duplicate interface names can exist in different router instances.

Enter a new name to create a logical router interface. When an existing interface name is entered, the user enters the router interface context for editing and configuration.

By default, no default IP interface names are defined within the system. All VPLS IP interfaces must be explicitly defined in an enabled state.

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

For VPLS services, the IP interface must be shutdown before the SAP on that interface is removed.

For VPLS service, ping and traceroute are the only applications supported.

Parameters *ip-int-name* — Specifies the name of the IP interface. Interface names must be unique within the group of defined IP.

An interface name:

- Should not be in the form of an IP address.
- Can be from 1 to 32 alphanumeric characters.
- If the string contains special characters (such as #,$,spaces), the entire string must be enclosed within double quotes.

If ip-int-name already exists within the service ID, the context changes to maintain that IP interface. If ip-int-name already exists within another service ID, an error occurs and the context does not change to that IP interface. If ip-int-name does not exist, the interface is created and the context is changed to that interface for further command processing.

address

Syntax **address** {*ip-address/mask* | *ip-address netmask*}
address *ip-address mask*

Context config>service>vpls>interface

Description This command assigns an IP address and an IP subnet, to a VPLS IP router interface. Only one IP address can be associated with an IP interface. An IP address must be assigned to each VPLS 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.

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.

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 ipaddress, 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. The values allowed are integers in the range 0 – 30. Note that a mask length of 32 is reserved for system IP 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.

General Switch Management Protocol Commands

gsmp

Syntax **gsmp**

Context config>service>vpls

Description This command enables the context to configure General Switch Management Protocol (GSMP) connections maintained in this service.

Default not enabled

group

Syntax **[no] group *name***

Context config>service>vpls>gsmp

Description This command specifies a GSMP name. A GSMP group name is unique only within the scope of the service in which it is defined.

ancp

Syntax **ancp**

Context config>service>vpls>gsmp>group

Description This command configures Access Node Control Protocol (ANCP) parameters for this GSMP group.

dynamic-topology-discover

Syntax [**no**] **dynamic-topology-discover**

Context config>service>vpls>gsmp>group>ancp

Description This command enables the ANCP dynamic topology discovery capability.

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

idle-filter

Syntax **idle-filter**
no idle-filter

Context config>service>vpls>gsmp
config>service>vprn>gsmp

Description This command when applied will filter out new subscriber's ANCP messages from subscriber with "DSL-line-state" IDLE

Default no idle-filter

line-configuration

Syntax [**no**] **line-configuration**

Context config>service>vpls>gsmp>group>ancp

Description This command enables the ANCP line-configuration capability.

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

oam

Syntax [**no**] **oam**

Context config>service>vpls>gsmp>group>ancp

Description This command specifies whether or not the GSMP ANCP OAM capability should be negotiated at startup of the GSMP connection.

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

hold-multiplier

Syntax **hold-multiplier** *multiplier*
no hold-multiplier

Context config>service>vpls>gsmp>group

Description This command configures the hold-multiplier for the GSMP connections in this group.

Parameters *multiplier* — Specifies the GSMP hold multiplier value.

Values 1 — 100

keepalive

Syntax **keepalive** *seconds*
no keepalive

Context config>service>vpls>gsmp>group

Description This command configures keepalive values for the GSMP connections in this group.

Parameters *seconds* — Specifies the GSMP keepalive timer value in seconds.

Values 1 — 25

neighbor

Syntax [**no**] **neighbo**r *ip-address*

Context config>service>vpls>gsmp>group

Description This command configures a GSMP ANCP neighbor.

Parameters *ip-address* — Specifies the IP address of the GSMP ANCP neighbor.

local-address

Syntax **local-address** *ip-address*
no local-address

Context config>service>vpls>gsmp>group>neighbor

Description This command configures the source ip-address used in the connection towards the neighbor. The local address is optional. If specified the node will accept connections only for that address in the service running ANCP. The address may be created after the reference but connections will not be accepted until it is created. If the local address is not used, the system accepts connections on any interface within the routing context.

Parameters *ip-address* — Specifies the source IP address to be used in the connection toward the neighbor.

priority-marking

Syntax **priority-marking dscp** *dscp-name*
priority-marking prec *ip-prec-value*
no priority-marking

Context config>service>vpls>gsmp>group>neighbor

Description This command configures the type of priority marking to be used.

Parameters **dscp** *dscp-name* — Specifies the DSCP code-point to be used.

Values be, cp1, cp2, cp3, cp4, cp5, cp6, cp7, cs1, cp9, af11, cp11, af12, cp13, af13, cp15, cs2, cp17, af21, cp19, af22, cp21, af23, cp23, cs3, cp25, af31, cp27, af32, cp29, af33, cp31, cs4, cp33, af41, cp35, af42, cp37, af43, cp39, cs5, cp41, cp42, cp43, cp44, cp45, ef, cp47, nc1, cp49, cp50, cp51, cp52, cp53, cp54, cp55, nc2, cp57, cp58, cp59, cp60, cp61, cp62, cp63

prec *ip-prec-value* — Specifies the precedence value to be used.

Values 0 — 7

persistency-database

Syntax **persistency-database**
no persistency-database

Context config>service>vpls <service id>gsmp
config>service>vprn<service id>gsmp

Description This command enables the system to store DSL line information in memory. If the GSMP connection terminates, the DSL line information will remain in memory and accessible for Radius authentication and accounting.

Default no persistency-database

VPLS DHCP Commands

dhcp

Syntax **dhcp**

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

Description This command enables the context to configure DHCP parameters.

lease-populate

Syntax **lease-populate** [*nmb-of-entries*]
no lease-populate

Context config>service>vpls>sap>dhcp

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

The optional number-of-entries parameter is used to define the number of lease state table entries allowed for this SAP or IP interface. If number-of-entries is omitted, only a single entry is allowed. Once the maximum number of entries has been reached, subsequent lease state entries are not allowed and subsequent DHCP ACK messages are discarded.

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

- populate a SAP based anti-spoof filter table to provide dynamic anti-spoof filtering. If the system is unable to populate the dynamic host information in the anti-spoof filter table on the SAP, the DHCP ACK message must be discarded without adding a new lease state entry or updating an existing lease state entry.
- generate dynamic ARP replies if **arp-reply-agent** is enabled.

Default no lease-populate

Parameters *nbr-of-entries* — Specifies the number of DHCP leases allowed.

Values 1 — 8000

option

Syntax [**no**] **option**

Context config>service>vpls>sap>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** [*dhcp-action*]
no action

Context config>service>vpls>sap>dhcp>option

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 *dhcp-action* — Specifies the DHCP option action.

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 sent on towards the client.

The behavior is slightly different in case of Vendor Specific Options (VSOs). When the keep parameter is specified, the router will insert his own VSO into the Option 82 field. This will only be done when the incoming message has already an Option 82 field.

If no Option 82 field is present, the router will not create the Option 82 field. In this in that case, no VSO will be added to the message.

circuit-id

Syntax **circuit-id** [**ascii-tuple** | **vlan-ascii-tuple**]

Context config>service>vpls>sap>dhcp>option

Description When enabled, the router sends an ASCII-encoded tuple in the **circuit-id** suboption of the DHCP packet. This ASCII-tuple consists of the access-node-identifier, service-id, and SAP-ID, separated by

"|". If no keyword is configured, then the circuit-id suboption will not be part of the information option (Option 82).

If disabled, the **circuit-id** suboption of the DHCP packet will be left empty.

Default no 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 and qinq 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 [**no**] **remote-id** [**mac** | **string** *string*]

Context config>service>vpls>sap>dhcp>option

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

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

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

Default no remote-id

Parameters **mac —** This keyword specifies the MAC address of the remote end is encoded in the suboption.

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

vendor-specific-option

Syntax [**no**] **vendor-specific-option**

Context config>service>vpls>sap>dhcp>option
config>service>ies>if>dhcp>option

Description This command configures the vendor specific suboption of the DHCP relay packet.

client-mac-address

Syntax [**no**] **client-mac-address**

Context config>service>vpls>sap>dhcp>option>vendor

Description This command enables the sending of the MAC address in the vendor specific suboption of the DHCP relay packet.

The **no** form of the command disables the sending of the MAC address in the vendor specific suboption of the DHCP relay packet.

sap-id

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

Context config>service>vpls>sap>dhcp>option>vendor

Description This command enables the sending of the SAP ID in the vendor specific suboption of the DHCP relay packet.

The **no** form of the command disables the sending of the SAP ID in the vendor specific suboption of the DHCP relay packet.

service-id

Syntax [**no**] **service-id**

Context config>service>vpls>sap>dhcp>option>vendor

Description This command enables the sending of the service ID in the vendor specific suboption of the DHCP relay packet.

The **no** form of the command disables the sending of the service ID in the vendor specific suboption of the DHCP relay packet.

string

Syntax [**no**] **string** *text*

Context config>service>vpls>sap>dhcp>option>vendor

Description This command specifies the string in the vendor specific suboption of the DHCP relay packet.

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

Parameters *text* — The string can be any combination of ASCII characters up to 32 characters in length. If spaces are used in the string, enclose the entire string in quotation marks (“ ”).

system-id

Syntax [**no**] **system-id**

Context config>service>vpls>sap>dhcp>option>vendor

Description This command specifies whether the system-id is encoded in the vendor specific sub-option of Option 82.

proxy-server

Syntax **proxy-server**

Context config>service>vpls>sap>dhcp

Description This command configures the DHCP proxy server.

emulated-server

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

Context config>service>vpls>sap>dhcp>proxy

Description This command configures the IP address which will be used as the DHCP server address in the context of this VPLS SAP. Typically, the configured address should be in the context of the subnet represented by the VPLS.

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

Parameters *ip-address* — Specifies the emulated server address.

lease-time

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

Context config>service>vpls>sap>dhcp>proxy

Description This command defines the length of lease time that will be provided to DHCP clients. By default, the local-proxy-server will always make use of the lease-time information provide by either a RADIUS or DHCP server.

The **no** form of this command disables the use of the lease-time command. The local proxy server will use the lease time offered by either a RADIUS or DHCP server.

Default 7 days 0 hours 0 seconds

Parameters *days* — Specifies the number of days that the given IP address is valid.

Values 0 — 3650

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

Values 0 — 23

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

Values 0 — 59

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

Values 0 — 59

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 VPLS 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

VPLS STP Commands

stp

Syntax **stp**

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

Description This command enables the context to configure the Spanning Tree Protocol (STP) parameters. Alcatel-Lucent's STP is simply the Spanning Tree Protocol (STP) with a few modifications to better suit the operational characteristics of VPLS services. The most evident change is to the root bridge election. Since the core network operating between Alcatel-Lucent's service routers should not be blocked, the root path is calculated from the core perspective.

auto-edge

Syntax **auto-edge**
no auto-edge

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

Description This command configures automatic detection of the edge port characteristics of the SAP or spoke SDP.

If auto-edge is enabled, and STP concludes there is no bridge behind the spoke SDP, the OPER_EDGE variable will dynamically be set to true. If auto-edge is enabled, and a BPDU is received, the OPER_EDGE variable will dynamically be set to true (see edge-port on page 910).

The **no** form of this command returns the auto-detection setting to the default value.

Default auto-edge

edge-port

Syntax [**no**] **edge-port**

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

Description This command configures the SAP or SDP as an edge or non-edge port. If **auto-edge** is enabled for the SAP, this value will be used only as the initial value.

NOTE: The function of the **edge-port** command is similar to the **rapid-start** command. It tells RSTP that it is on the edge of the network (for example, there are no other bridges connected to that port)

and, as a consequence, it can immediately transition to a forwarding state if the port becomes available.

RSTP, however, can detect that the actual situation is different from what **edge-port** may indicate.

Initially, the value of the SAP or spoke SDP parameter is set to edge-port. This value will change if:

- A BPDU is received on that port. This means that after all there is another bridge connected to this port. Then the edge-port becomes disabled.
- If auto-edge is configured and no BPDU is received within a certain period of time, RSTP concludes that it is on an edge and enables the edge-port.

The **no** form of this command returns the edge port setting to the default value.

Default no edge-port

forward-delay

Syntax **forward-delay** *seconds*
no forward-delay

Context config>service>vpls>stp
config>template>vpls-template>stp

Description RSTP, as defined in the IEEE 802.1D-2004 standards, will normally transition to the forwarding state via a handshaking mechanism (rapid transition), without any waiting times. If handshaking fails (e.g. on shared links, see below), the system falls back to the timer-based mechanism defined in the original STP (802.1D-1998) standard.

A shared link is a link with more than two nodes (for example, a shared 10/100BaseT segment). The `port-type` command is used to configure a link as point-to-point or shared.

For timer-based transitions, the 802.1D-2004 standard defines an internal variable forward-delay, which is used in calculating the default number of seconds that a SAP or spoke SDP spends in the discarding and learning states when transitioning to the forwarding state.

The value of the forward-delay variable depends on the STP operating mode of the VPLS instance:

- in `rstp` or `mstp` mode, but only when the SAP or spoke SDP has not fallen back to legacy STP operation, the value configured by the `hello-time` command is used;
- in all other situations, the value configured by the `forward-delay` command is used.

Default 15 seconds

Parameters *seconds* — The forward delay timer for the STP instance in seconds.

Values 4 — 30

hello-time

Syntax **hello-time** *hello-time*
no hello-time

Context config>service>vpls>stp

config>template>vpls-template>stp

Description This command configures the Spanning Tree Protocol (STP) hello time for the Virtual Private LAN Service (VPLS) STP instance.

The hello time parameter defines the default timer value that controls the sending interval between BPDU configuration messages by this bridge, on ports where this bridge assumes the designated role.

The active hello time for the spanning tree is determined by the root bridge (except when the STP is running in RSTP mode, then the hello time is always taken from the locally configured parameter).

The configured hello-time can also be used to calculate the forward delay. See auto-edge on page 910.

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

Default 2 seconds

Parameters *hello-time* — The hello time for the STP instance in seconds.

Values 1 — 10

hold-count

Syntax **hold-count** *BDPU tx hold count*
no hold-count

Context config>service>vpls>stp
config>template>vpls-template>stp

Description This command configures the peak number of BPDUs that can be transmitted in a period of one second.

The **no** form of this command returns the hold count to the default value

Default 6

Parameters *BDPU tx hold count* — The hold count for the STP instance in seconds.

Values 1 — 10

link-type

Syntax **link-type** {**pt-pt** | **shared**}
no link-type

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

Description This command instructs STP on the maximum number of bridges behind this SAP or spoke SDP. If there is only a single bridge, transitioning to forwarding state will be based on handshaking (fast transitions). If more than two bridges are connected via a shared media, their SAP or spoke SDPs should all be configured as shared, and timer-based transitions are used.

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

Default pt-pt

mst-instance

Syntax **mst-instance** *mst-inst-number*

Context config>service>vpls>sap>stp

Description This command enables the context to configure MSTI related parameters at SAP level. This context can be open only for existing mst-instances defined at the service level (see **mst-instance**).

Default none

Parameters *mst-inst-number* — Specifies an existing Multiple Spanning Tree Instance number.

Values 1 — 4094

mst-path-cost

Syntax **mst-path-cost** *inst-path-cost*
no mst-path-cost

Context config>service>vpls>sap>stp>mst-instance

Description This commands specifies path-cost within a given instance, expressing probability that a given port will be put into the forwarding state in case a loop occurs (the highest value expresses lowest priority).

The **no** form of this command sets port-priority to its default value.

Default The path-cost is proportional to link speed.

Parameters *inst-path-cost* — Specifies the contribution of this port to the MSTI path cost of paths towards the spanning tree regional root which include this port.

Values 1 — 200000000

mst-priority

Syntax **mst-priority** *stp-priority*
no mst-priority

Context config>service>vpls>sap>stp>mst-instance

Description This commands specifies the port priority within a given instance, expressing probability that a given port will be put into the forwarding state if a loop occurs.

The **no** form of this command sets port-priority to its default value.

Default 128

Parameters *stp-priority* — Specifies the value of the port priority field.

max-age

Syntax **max-age** *seconds*
no max-age

Context config>service>vpls>stp
config>template>vpls-template>stp

Description This command indicates how many hops a BPDU can traverse the network starting from the root bridge. The message age field in a BPDU transmitted by the root bridge is initialized to 0. Each other bridge will take the message_age value from BPDUs received on their root port and increment this value by 1. The message_age thus reflects the distance from the root bridge. BPDUs with a message age exceeding max-age are ignored.

STP uses the max-age value configured in the root bridge. This value is propagated to the other bridges via the BPDUs.

The **no** form of this command returns the max ageto the default value.

Default 20 seconds

Parameters *seconds* — The max info age for the STP instance in seconds. Allowed values are integers in the range 6 to 40.

mode

Syntax **mode** {**rstp** | **comp-dot1w** | **dot1w** | **mstp** | **pmstp**}
no mode

Context config>service>vpls>stp
config>template>vpls-template>stp

Description This command specifies the version of Spanning Tree Protocol the bridge is currently running.

See section Spanning Tree Operating Modes on page 622 for details on these modes.

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

Default rstp

Parameters **rstp —** Corresponds to the Rapid Spanning Tree Protocol specified in IEEE 802.1D/D4-2003.

dot1w — Corresponds to the mode where the Rapid Spanning Tree is backward compatible with IEEE 802.1w.

compdot1w — Corresponds to the Rapid Spanning Tree Protocol fully conformant to IEEE 802.1w.

mstp — Sets MSTP as the STP mode of operation. Corresponds to the Multiple Spanning Tree Protocol specified in 802.1Q REV/D5.0-09/2005

pmstp — The PMSTP mode is only supported in VPLS services where the mVPLS flag is configured.

mst-instance

Syntax [**no**] **mst-instance** *mst-inst-number*

Context config>service>vpls>stp

Description This command creates the context to configure MST instance (MSTI) related parameters. Up to 16 instances will be supported by MSTP. The instance 0 is mandatory by protocol and therefore, it cannot be created by the CLI. The software will maintain this instance automatically.

Default none

Parameters *mst-inst-number* — Specifies the Multiple Spanning Tree instance.

Values 1 — 4094

mst-priority

Syntax **mst-priority** *bridge-priority*
no mst-priority

Context config>service>vpls>stp>mst-instance

Description This command specifies the bridge priority for this specific Multiple Spanning Tree Instance for this service. The *bridge-priority* value reflects likelihood that the switch will be chosen as the regional root switch (65535 represents the least likely). It is used as the highest 4 bits of the Bridge ID included in the MSTP BPDU's generated by this bridge.

The priority can only take on values that are multiples of 4096 (4k). If a value is specified that is not a multiple of 4K, then the value will be replaced by the closest multiple of 4K, which is lower than the value entered.

The **no** form of this command sets the bridge-priority to its default value.

Default 32768 — All instances created by **vlan-range** command and not having explicit definition of bridge-priority will inherit default value.

Parameters *bridge-priority* — Specifies the priority of this specific Multiple Spanning Tree Instance for this service.

Values 0 — 65535

vlan-range

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

Context config>service>vpls>stp>mst-instance

Description This command specifies a range of VLANs associated with a certain MST-instance. This range applies to all SAPs of the mVPLS.

Every VLAN range that is not assigned within any of the created **mst-instance** is automatically assigned to mst-instance 0. This instance is automatically maintained by the software and cannot be

modified. Changing the VLAN range value can be performed only when the given mst-instance is shutdown.

The **no** form of this command removes the **vlan-range** from given **mst-instance**.

Parameters *vlan-range* — The first VLAN range specifies the left-bound (i.e., minimum value) of a range of VLANs that are associated with the mVPLS SAP. This value must be smaller than (or equal to) the second VLAN range value. The second VLAN range specifies the right-bound (i.e., maximum value) of a range of VLANs that are associated with the mVPLS SAP.

Values 1 to 4094 — 1 to 4094

mst-max-hops

Syntax **mst-max-hops** *hops-count*
no mst-max-hops

Context config>service>vpls>stp

Description This command specifies the number of hops in the region before BPDU is discarded and the information held for the port is aged out. The root bridge of the instance sends a BPDU (or M-record) with remaining-hop-count set to configured <*max-hops*>. When a bridge receives the BPDU (or M-record), it decrements the received remaining-hop-count by 1 and propagates it in BPDU (or M-record) it generates.

The **no** form of this command sets the *hops-count* to its default value.

Default 20

Parameters *hops-count* — Specifies the maximum number of hops.

Values 1 — 40

mst-name

Syntax **mst-name** *region-name*
no mst-name

Context config>service>vpls>stp

Description This command defines an MST region name. Two bridges are considered as a part of the same MST region as soon as their configuration of the MST region name, the MST-revision and VLAN-to-instance assignment is identical.

The **no** form of this command removes *region-name* from the configuration.

Default no mst-name

Parameters *region-name* — Specifies an MST-region name up to 32 characters in length.

mst-revision

Syntax **mst-revision** *revision-number*

Context config>service>vpls>stp

Description This command defines the MST configuration revision number. Two bridges are considered as a part of the same MST region as soon as their configuration of MST-region name, MST-revision and VLAN-to-instance assignment is identical.

The **no** form of this command returns MST configuration revision to its default value.

Default 0

Parameters *revision-number* — Specifies the MSTP region revision number to define the MSTP region.

Values 0 — 65535

path-cost

Syntax **path-cost** *sap-path-cost*
no path-cost

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

Description This command configures the Spanning Tree Protocol (STP) path cost for the SAP or spoke SDP.

The path cost is used by STP to calculate the path cost to the root bridge. The path cost in BPDUs received on the root port is incremented with the configured path cost for that SAP or spoke SDP. When BPDUs are sent out other egress SAPs or spoke SDPs, the newly calculated root path cost is used. These are the values used for CIST when running MSTP.

STP suggests that the path cost is defined as a function of the link bandwidth. Since SAPs and spoke SDPs are controlled by complex queuing dynamics, in the 7750 SR the STP path cost is a purely static configuration.

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

path-cost — The path cost for the SAP or spoke SDP.

Values 1 — 200000000 (1 is the lowest cost)

Default 10

port-num

Syntax [**no**] **port-num** *virtual-port-number*

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

Description This command configures the virtual port number which uniquely identifies a SAP within configuration bridge protocol data units (BPDUs). The internal representation of a SAP is unique to a

system and has a reference space much bigger than the 12 bits definable in a configuration BPDU. STP takes the internal representation value of a SAP and identifies it with it's own virtual port number that is unique to every other SAP defined on the TLS. The virtual port number is assigned at the time that the SAP is added to the TLS. Since the order that the SAP was added to the TLS is not preserved between reboots of the system, the virtual port number may change between restarts of the STP instance.

The virtual port number cannot be administratively modifed.

priority

Syntax **priority** *bridge-priority*
no priority

Context config>service>vpls>stp
config>template>vpls-template>stp

Description The bridge-priority command is used to populate the priority portion of the bridge ID field within outbound BPDUs (the most significant 4 bits of the bridge ID). It is also used as part of the decision process when determining the best BPDU between messages received and sent. All values will be truncated to multiples of 4096, conforming with IEEE 802.1t and 802.1D-2004.

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

Default By default, the bridge priority is configured to 4096 which is the highest priority.

Parameters *bridge-priority* — The bridge priority for the STP instance.

Values Allowed values are integers in the range of 4096 — 65535 with 4096 being the highest priority. The actual bridge priority value stored/used is the number entered with the lowest 12 bits masked off which means the actual range of values is 4096 to 61440 in increments of 4096.

priority

Syntax **priority** *stp-priority*
no priority

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

Description This command configures the Alcatel-Lucent Spanning Tree Protocol (STP) priority for the SAP or spoke SDP.

STP priority is a configurable parameter associated with a SAP or spoke SDP. When configuration BPDUs are received, the priority is used in some circumstances as a tie breaking mechanism to determine whether the SAP or spoke SDP will be designated or blocked.

In traditional STP implementations (802.1D-1998), this field is called the port priority and has a value of 0 to 255. This field is coupled with the port number (0 to 255 also) to create a 16 bit value. In the latest STP standard (802.1D-2004) only the upper 4 bits of the port priority field are used to encode the SAP or spoke SDP priority. The remaining 4 bits are used to extend the port ID field into a 12 bit

virtual port number field. The virtual port number uniquely references a SAP or spoke SDP within the STP instance.

STP computes the actual priority by taking the input value and masking out the lower four bits.The result is the value that is stored in the SDP priority parameter. For instance, if a value of 0 is entered, masking out the lower 4 bits results in a parameter value of 0. If a value of 255 is entered, the result is 240.

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

Default 128

Parameters *stp-priority* — The STP priority value for the SAP or spoke SDP. Allowed values are integer in the range of 0 to 255, 0 being the highest priority. The actual value used for STP priority (and stored in the configuration) will be the result of masking out the lower 4 bits, thus the actual value range is 0 to 240 in increments of 16.

Default 128

VPLS SAP Commands

sap

Syntax **sap** *sap-id* [**split-horizon-group** *group-name*][**capture-sap**] [**create**] [**eth-ring** *ring-index*]
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 Enhanced 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. The limits of the number of SAPs and SDPs supported in a VPLS service depends on the hardware used. Each SDP must have a unique destination or an error will be generated. Split horizon groups can only be created in the scope of a VPLS service.

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 CLI Command Descriptions on page 2569 for command syntax.

create — Keyword used to create a SAP instance. The **create** keyword requirement can be enabled/disabled in the **environment>create** context.

eth-ring — When used with Ethernet Rings control the split-horizon-group accepts the major ring instance "value". The split horizon group prevents loops in the cases where a Ethernet Virtual Ring is miss configured on the main ring. Each path a and path b major ring are configured in the group and associated with the sub-ring control instance in the VPLS service.

ring-index — Specifies the ring index of the Ethernet ring.

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

capture-sap — Specifies a capturing SAP in which triggering packets will be sent to the CPM. Non-triggering packets captured by the capture SAP will be dropped.

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 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 931) 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**

ETH-CFM Service Commands

eth-cfm

Syntax **eth-cfm**

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

Description This command enables the context to configure ETH-CFM parameters.

eth-tunnel

Syntax **eth-tunnel**

Context config>service>vpls>sap

Description The command enables the context to configure Ethernet Tunnel SAP parameters.

eth-ring

Syntax **eth-ring** *ring-id*
no eth-ring

Context config>service>vpls

Description This command configures a VPLS Sap to be associated with an Ethernet ring. The Sap port-id is associated with the corresponding Ethernet ring path configured on the same port-id. The encapsulation type must be compatible with the Eth-ring path encapsulation.

The **no** form of this command removes eth-ring from this SAP

Default no eth-ring

Parameters *ring-id* — Specifies the ring ID.

Values 1-128

path

Syntax **path** *path-index* **tag** *qtag*[.*qtag*]
no path *path-index*

Context config>service>vpls>sap>eth-tunnel

Description This command configures Ethernet tunnel SAP path parameters.

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

Default none

Parameters *path-index* — Specifies the path index value.

Values 1 — 16

tag *qtag*[.*qtag*] **—** Specifies the qtag value.

Values 0 — 4094, *

mep

Syntax **mep** *mep-id* **domain** *md-index* **association** *ma-index* [**direction** {**up** | **down**}] **primary-valn-enable** [**vlan** vlan-id]
no mep *mep-id* **domain** *md-index* **association** *ma-index*

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

Description This command configures the ETH-CFM maintenance endpoint (MEP). A MEP created at the VPLS service level **vpls>eth-cfm** creates a virtual MEP.

The no version of the command will remove the MEP.

Parameters *mep-id* — Specifies the maintenance association end point identifier.

Values 1 — 8191

md-index — Specifies the maintenance domain (MD) index value.

Values 1 — 4294967295

ma-index — Specifies the MA index value.

Values 1 — 4294967295

direction up| down — Indicates the direction in which the maintenance association (MEP) faces on the bridge port. Direction is not supported when a MEP is created directly under the vpls>eth-cfm construct (vMEP).

down — Sends ETH-CFM messages away from the MAC relay entity.

up — Sends ETH-CFM messages towards the MAC relay entity.

primary-vlan-enable — Provides a method for linking the MEP with the primary VLAN configured under the bridge-identifier for the MA. MEPs can not be changed from or to primary vlan functions. This must be configured as part of the creation step and can only be changed by deleting the MEP and recreating it. Primary VLANs are only supported under Ethernet SAPs.

vlan — A required parameter when including primary-vlan-enable. Provides a method for associating the VLAN under the bride-identifier under the MA with the MEP.

vlan-id — Must match the vlan-id under the bridge-identifier for the MA that is appropriate for this service

Values 0 — 4094

mip

Syntax **mip** [**mac** *mac-address*] **primary-vlan-enable** [*vlan vlan-id*]
mip default-mac
no mip

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

Description This command allows Maintenance Intermediate Points (MIPs). The creation rules of the MIP are dependant on the mhf-creation configuration for the MA. This MIP option is only available for default and static mhf-creation methods.

Parameters *mac-address* — Specifies the MAC address of the MEP.

Values 6-byte mac-address in the form of xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx of the MIP. The MAC must be unicast. Using the all zeros address is equivalent to the **no** form of this command.

default-mac — Using the no command deletes the MIP. If the operator wants to change the mac back to the default mac without having to delete the MIP and reconfiguring this command is useful.

primary-vlan-enable — Provides a method for linking the MIP with the primary VLAN configured under the bridge-identifier for the MA. This is only allowed if the mhf-creation method is static. MIPs can not be changed from or to primary vlan functions without first being deleted. This must be configured as part of the creation step and can only be changed by deleting the MEP and recreating it. Primary VLANs are only supported under Ethernet SAPs.

vlan — A required parameter when including primary-vlan-enable. Provides a method for associating the VLAN under the bride-identifier under the MA with the MIP.

vlan-id — Must match the vlan-id under the bridge-identifier for the MA that is appropriate for this service.

Values 0 — 4094

Default **no** mip

mip

Syntax **mip primary-vlan-enable** [**vlan** *vlan-id*]
no mip

Context config>service>template>vpls-sap-template>eth-cfm

Description This command allows Maintenance Intermediate Points (MIPs). The creation rules of the MIP are dependant on the mhf-creation configuration for the MA. This MIP option is only available for default and static mhf-creation methods.

Parameters **primary-vlan-enable —** Provides a method for linking the MIP with the primary VLAN configured under the bridge-identifier for the MA. This is only allowed if the mhf-creation method is static. MIPs can not be changed from or to primary vlan functions without first being deleted. This must be configured as part of the creation step and can only be changed by deleting the MEP and recreating it. Primary VLANs are only supported under Ethernet SAPs.

vlan — A required parameter when including primary-vlan-enable. Provides a method for associating the VLAN under the bride-identifier under the MA with the MIP.

vlan-id — Must match the vlan-id under the bridge-identifier for the MA that is appropriate for this service

Values 0 — 4094

ais-enable

Syntax **[no] ais-enable**

Context config>service>vpls>mesh-sdp>eth-cfm>mep
config>service>vpls>spoke-sdp>eth-cfm>mep

Description This command enables the generation and the reception of AIS messages.

client-meg-level

Syntax **client-meg-level** [[*level* [*level* ...]]
no client-meg-level

Context config>service>vpls>mesh-sdp>eth-cfm>mep>ais-enable
config>service>vpls>spoke-sdp>eth-cfm>mep>ais-enable

Description This command configures the client maintenance entity group (MEG) level(s) to use for AIS message generation. Up to 7 levels can be provisioned with the restriction that the client MEG level must be higher than the local MEG level.

Parameters *level* — Specifies the client MEG level.

Values 1 — 7

Default 1

interval

Syntax **interval** {**1** | **60**}
no interval

Context config>service>vpls>mesh-sdp>eth-cfm>mep>ais-enable
config>service>vpls>spoke-sdp>eth-cfm>mep>ais-enable

Description This command specifies the transmission interval of AIS messages in seconds.

Parameters **1 | 60 —** The transmission interval of AIS messages in seconds.

Default 1

priority

Syntax **priority** *priority-value*
no priority

Context config>service>vpls>mesh-sdp>eth-cfm>mep>ais-enable
config>service>vpls>spoke-sdp>eth-cfm>mep>ais-enable

Description This command specifies the priority of AIS messages originated by the node.

Parameters *priority-value* — Specify the priority value of the AIS messages originated by the node.

Values 0 — 7

Default 1

ccm-enable

Syntax [**no**] **ccm-enable**

Context config>service>vpls>eth-cfm>mep
config>service>vpls>sap>eth-cfm>mep
config>service>vpls>mesh-sdp>mep
config>service>vpls>spoke-sdp>eth-cfm>mep

Description This command enables the generation of CCM messages.

The **no** form of the command disables the generation of CCM messages.

ccm-ltm-priority

Syntax **ccm-ltm-priority** *priority*
no ccm-ltm-priority

Context config>service>vpls>eth-cfm>mep
config>service>vpls>sap>eth-cfm>mep
config>service>vpls>mesh-sdp>mep
config>service>vpls>spoke-sdp>eth-cfm>mep

Description This command specifies the priority value for CCMs and LTMs transmitted by the MEP.

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

Default The highest priority on the bridge-port.

Parameters *priority* — Specifies the priority of CCM and LTM messages.

Values 0 — 7

eth-test-enable

Syntax [**no**] **eth-test-enable**

Context config>service>vpls>sap>eth-cfm>mep
config>service>vpls>spoke-sdp>eth-cfm>mep
config>service>vpls>mesh-sdp>eth-cfm>mep

Description For ETH-test to work, operators need to configure ETH-test parameters on both sender and receiver nodes. The ETH-test then can be done using the following OAM commands:

oam eth-cfm eth-test *mac-address* mep *mep-id* domain *md-index* association *ma-index* [priority *priority*] [data-length *data-length*]

A check is done for both the provisioning and test to ensure the MEP is an Y.1731 MEP (MEP provisioned with domain format none, association format icc-based). If not, the operation fails. An error message in the CLI and SNMP will indicate the problem.

test-pattern

Syntax **test-pattern** {**all-zeros** | **all-ones**} [**crc-enable**]
no test-pattern

Context config>service>vpls>sap>eth-cfm>mep>eth-test-enable
config>service>vpls>spoke-sdp>eth-cfm>mep>eth-test-enable
config>service>vpls>mesh-sdp>eth-cfm>mep>eth-test-enable

Description This command configures the test pattern for eth-test frames.

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

Parameters **all-zeros** — Specifies to use all zeros in the test pattern.

all-ones — Specifies to use all ones in the test pattern.

crc-enable — Generates a CRC checksum.

Default all-zeros

fault-propagation-enable

Syntax **fault-propagation-enable** {**use-if-tlv** | **suspend-ccm**}
no fault-propagation-enable

Context config>service>vpls>mesh-sdp>eth-cfm>mep
config>service>vpls>sap>eth-cfm>mep
config>service>vpls>spoke-sdp>eth-cfm>mep

Description This command configures the fault propagation for the MEP.

Parameters **use-if-tlv** — Specifies to use the interface TLV.

suspend-ccm — Specifies to suspend the continuity check messages.

low-priority-defect

Syntax **low-priority-defect** {**allDef**|**macRemErrXcon**|**remErrXcon**|**errXcon**|**xcon**|**noXcon**}

Context config>service>vpls>mesh-sdp>eth-cfm>mep
config>service>vpls>>spoke-sdp>eth-cfm>mep

Description This command specifies the lowest priority defect that is allowed to generate a fault alarm.

Default macRemErrXcon

Values

allDef	DefRDICCM, DefMACstatus, DefRemoteCCM, DefErrorCCM, and DefXconCCM
macRemErrXcon	Only DefMACstatus, DefRemoteCCM, DefErrorCCM, and DefXconCCM
remErrXcon	Only DefRemoteCCM, DefErrorCCM, and DefXconCCM
errXcon	Only DefErrorCCM and DefXconCCM
xcon	Only DefXconCCM; or
noXcon	No defects DefXcon or lower are to be reported

mac-address

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

Context config>service>vpls>eth-cfm>mep
config>service>vpls>sap>eth-cfm>mep
config>service>vpls>>spoke-sdp>eth-cfm>mep
config>service>vpls>mesh-sdp>eth-cfm>mep

Description This command specifies the MAC address of the MEP.

The **no** form of this command reverts the MAC address of the MEP back to that of the port (if the MEP is on a SAP) or the bridge (if the MEP is on a spoke).

Parameters *mac-address* — Specifies the MAC address of the MEP.

Values 6-byte mac-address in the form of xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx of the MEP. Must be unicast. Using the all zeros address is equivalent to the no form of this command.

one-way-delay-threshold

Syntax **one-way-delay-threshold** *seconds*

Context config>service>vpls>sap>eth-cfm>mep

Description This command enables/disables eth-test functionality on MEP.

Parameters *seconds* — Specifies the one way delay threshold, in seconds.

Values 0..600

Default 3

tunnel-fault

Syntax **tunnel-fault {accept | ignore}**

Context config>service>vpls>eth-cfm
config>service>vpls>sap>eth-cfm

Description Allows the individual service SAPs to react to changes in the tunnel MEP state. When tunnel-fault accept is configured at the service level, the SAP will react according to the service type, Epipe will set the operational flag and VPLS, IES and VPRN SAP operational state will become down on failure or up on clear. This command triggers the OAM mapping functions to mate SAPs and bindings in an Epipe service as well as setting the operational flag. If AIS generation is the requirement for the Epipe services this command is not required. See the command ais-enable under epipe>sap>eth-cfm>ais-enable for more details. This works in conjunction with the tunnel-fault accept on the individual SAPs. Both must be set to accept to react to the tunnel MEP state. By default the service level command is "ignore" and the sap level command is "accept". This means simply changing the service level command to "accept" will enable the feature for all SAPs. This is not required for Epipe services that only wish to generate AIS on failure.

Parameters **accept —** Share fate with the facility tunnel MEP

ignore — Do not share fate with the facility tunnel MEP

Default **ignore** (Service Level)

accept (SAP Level for Epipe and VPLS)

vmep-extensions

Syntax **vmep-extensions**
[no] vmep-extensions

Context config>service>vpls>eth-cfm

Description This command enables and disables enhanced Virtual Maintenance Endpoints functionality. This must manually be configured for a B-VPLS to change the legacy behavior and cannot be disable for VPLS contexts that are not BVPLS based.

The **no** form of the command reverts to the default values. This is not applicable to a VPLS contexts that is not B-VPLS based.

Default no vmep-extensions (for B-VPLS)

vmep-extensions (for VPLS contexts not B-VPLS based)

vmep-filter

Syntax **vmep-filter**
[no] vmep-filter

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

Description Suppress eth-cfm PDUs based on level lower than or equal to configured Virtual MEP. This command is not supported under a B-VPLS context. This will also delete any MIP configured on the SAP or Spoke-SDP.

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

Default no vmep-filter

limit-mac-move

Syntax **limit-mac-move** [**blockable** | **non-blockable**]
no limit-mac-move

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

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

Description Enabling this command will disable re-learning of MAC addresses on other SAPs within the VPLS. The MAC address will remain attached to a given SAP for duration of its age-timer.
The age of the MAC address entry in the FIB is set by the age timer. If mac-aging is disabled on a given VPLS service, any MAC address learned on a SAP/SDP with mac-pinning enabled will remain in the FIB on this SAP/SDP forever. Every event that would otherwise result in re-learning will be

logged (MAC address; original-SAP; new-SAP).

Note that MAC addresses learned during DHCP address assignment (DHCP snooping enabled) are not impacted by this command. MAC-pinning for such addresses is implicit.

Default When a SAP or spoke SDP is part of a Residential Split Horizon Group (RSHG), MAC pinning is activated at creation of the SAP. Otherwise, MAC pinning is not enabled by default.

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
config>service>vpls>endpoint

Description This command specifies the maximum number of FDB entries for both learned and static MAC addresses for this SAP, spoke SDP or endpoint.

When the configured limit has been reached, and discard-unknown-source has been enabled for this SAP or spoke SDP (see discard-unknown-source on page 921), 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* — Specifies the maximum number of learned and static entries allowed in the FDB of this service.

Values 1 — 511999

Chassis-mode C limit: 196607

Chassis-mode D limit: 511999

mc-endpoint

Syntax **mc-endpoint** *mc-ep-id*
mc-endpoint

Context config>service>vpls>endpoint

Description This command specifies the identifier associated with the multi-chassis endpoint. This value should be the same on both MC-EP peers for the pseudowires that must be part of the same group.

The **no** form of this command removes the endpoint from the MC-EP. Single chassis behavior applies.

Default no mc-endpoint

Parameters *mc-ep-id* — Specifies a multi-chassis endpoint ID.

Values 1 — 4294967295

mc-ep-peer

Syntax **mc-ep-peer** *name*
mc-ep-peer *ip-address*
no mc-ep-peer

Context config>service>vpls>endpoint>mc-ep

Description This command adds multi-chassis endpoint object.

The **no** form of this command removes the MC-Endpoint object.

Default mc-endpoint is not provisioned.

Parameters *name* — Specifies the name of the multi-chassis end-point peer.

ip-address — Specifies the IP address of multi-chassis end-point peer.

msap-defaults

Syntax **msap-defaults**

Context config>service>vpls>sap

Description This command configures the msap-defaults.

service

Syntax [**no**] **service** *service-id*

Context config>service>vpls>sap>msap-defaults

Description This command sets default service for all subscribers created based on trigger packets received on the given capture SAP in case the corresponding VSA is not included in RADIUS authentication response. This command is applicable to capture SAP only.

Default no service.

policy

Syntax **policy** *msap-policy-name*
no policy

Context config>service>vpls>sap>msap-defaults

Description This command sets default msap-policy for all subscribers created based on trigger packets received on the given capture-sap in case the corresponding VSA is not included in the RADIUS authentication response. This command is applicable to capture SAP only.

Default no policy

multi-service-site

Syntax **multi-service-site** *customer-site-name*
no multi-service-site

Context config>service>vpls>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.

precedence

Syntax **precedence** [*precedence-value* | primary]
no precedence

Context config>service>vpls>spoke-sdp

Description This command configures the precedence of this SDP bind when there are multiple SDP binds attached to one service endpoint. When an SDP bind goes down, the next highest precedence SDP bind begins forwarding traffic.

Parameters *precendence-value* — Specifies the precedence of this SDP bind.

Values 1 — 4

primary — Assigns this as the primary spoke-sdp.

static-isid

Syntax [**no**] **static-isid range** *entry-id isid* [**to** *isid*] [**create**]

Context config>service>vpls><instance> b-vpls>sap
config>service>vpls><instance> b-vpls>spokeSdp

Description This command identifies a set of ISIDs for I-VPLS services that are external to SPBM. These ISIDs are advertised as supported locally on this node unless an altered by an isid-policy. This allows communication from I-VPLS services external to SPBM through this node. The SAP may be a regular SAP or MC-LAG SAP. The spoke SDP may be a active/standby spoke. When used with MC-Lag or active/stand-by PWs the conditional static-mac must be configured. ISIDs declared this way become part of the ISID multicast and consume MFIBs. Multiple SPBM static-isid ranges are allowed under a SAP/spoke SDP.

The static-isids are associated with a remote BMAC that must be declared as a static-mac for unicast traffic. ISIDs are advertised as if they were attached to the local BMAC. Only remote I-VPLS ISIDs need to be defined. In the MFIB, the group MACs are then associated with the active SAP or spoke SDP. An ISID policy may be defined to suppress the advertisement of an ISID if the ISID is primary used for unicast services. The following rules govern the usage of multiple ISID statements:

- overlapping values are allowed:
  - isid from 301 to 310
  - isid from 305 to 315
  - isid 316
- the minimum and maximum values from overlapping ranges are considered and displayed. The above entries will be equivalent with “ISID from 301 to 316” statement.
- there is no consistency check with the content of ISID statements from other entries. The entries will be evaluated in the order of their IDs and the first match will cause the implementation t o execute the associated action for that entry.

no isid - removes all the previous statements under one interface

no isid value | from value to higher-value - removes a specific ISID value or range. Must match a previously used positive statement: for example if the command “isid 316 to 400” was used using “no isid 316 to 350” will not work but “no isid 316 to 400 will be successful.

Parameters *entry-id* — Sets context for specified entry ID for the static-isids.

Values 1— 65535

isid — Configures the ISID or the start of an ISID range. Specifies the ISID value in 24 bits. When just one present identifies a particular ISID to be used for matching.

Values 0..16777215

to isid — Identifies upper value in a range of ISIDs to be used as matching criteria.

Values 0..16777215

static-mac

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

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

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

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.

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

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

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

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

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

Parameters *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.

create — This keyword is mandatory when specifying a static MAC address.

managed-vlan-list

Syntax **managed-vlan-list**

Context config>service>vpls>sap

Description This command enables the context to configure 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 managed-vlan-list is not used when STP mode is MSTP in which case the vlan-range is taken from the **config>service>vpls>stp>msti** configuration.

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

default-sap

Syntax [**no**] **default-sap**

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

Description This command adds a default SAP to the managed VLAN list.

The **no** form of the command removes the default SAP to the managed VLAN list.

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. See Modifying VPLS Service Parameters on page 812.

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: 0 — 4094
end-vlan: 0 — 4094

VPLS SAP ATM Commands

atm

Syntax **atm**

Context config>service>vpls>sap

Description This command enables access to the context to configure ATM-related attributes. This command can only be used when a given context (for example, a channel or SAP) supports ATM functionality such as:

- Configuring ATM port or ATM port-related functionality on MDAs supporting ATM functionality
- Configuring ATM-related configuration for ATM-based SAPs that exist on MDAs supporting ATM functionality.

If ATM functionality is not supported for a given context, the command returns an error.

egress

Syntax **egress**

Context config>service>vpls>sap>atm

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

encapsulation

Syntax **encapsulation** *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, a PVCC's operational status is affected when a PVCC goes into an AIS or RDI state because of an AIS/RDI processing (assuming nothing else affects PVCC's operational status, for example, if the PVCC goes DOWN, or enters a fault state and comes back UP, or exits that fault state). 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, the PVCC's operational status is no longer affected by the 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 from DOWN due to alarm-cell processing). RDI cells are not generated as result of PVCC going into an AIS or RDI state, however, the PVCC's OAM status will record OAM faults as described above.

Default Enabled for PVCCs delimiting VPLS SAPs.

VPLS Filter and QoS Policy Commands

egress

Syntax **egress**

Context config>service>vpls>sap

Description This command enables the context to configure egress filter policies.

If no sap-egress QoS policy is defined, the system default sap-egress QoS policy is used for egress processing. If **no** egress filter is defined, no filtering is performed.

ingress

Syntax **ingress**

Context config>service>vpls>sap

Description This command enables the context to configure ingress SAP Quality of Service (QoS) policies and filter policies.

If no sap-ingress QoS policy is defined, the system default sap-ingress QoS policy is used for ingress processing. If no ingress filter is defined, no filtering is performed.

agg-rate-limit

Syntax **agg-rate-limit** *agg-rate* [**queue-frame-based-accounting**]
no agg-rate-limit

Context config>service>vpls>sap>egress
config>service>vpls>sap>egress>encap-defined-qos>encap-group

Description This command defines a maximum total rate for all egress queues on a service SAP or multi-service site. The **agg-rate-limit** command is mutually exclusive with the egress scheduler policy. When an egress scheduler policy is defined, the **agg-rate-limit** command will fail. If the **agg-rate-limit** command is specified, an attempt to bind a **scheduler-policy** to the SAP or multi-service site will fail.

A multi-service site must have a port scope defined that ensures all queues associated with the site are on the same port or channel. If the scope is not set to a port, the **agg-rate-limit** command will fail. Once an agg-rate-limit has been assigned to a multi-service site, the scope cannot be changed to card level.

A port scheduler policy must be applied on the egress port or channel the SAP or multi-service site are bound to in order for the defined agg-rate-limit to take effect. The egress port scheduler enforces the aggregate queue rate as it distributes its bandwidth at the various port priority levels. The port scheduler stops offering bandwidth to member queues once it has detected that the aggregate rate limit has been reached.

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

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

Parameters *agg-rate* — Defines the rate, in kilobits-per-second, that the maximum aggregate rate that the queues on the SAP or MSS can operate.

Values 1 — 40000000, max

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

Default Frame based accounting is disabled by default

agg-rate-limit

Syntax **agg-rate-limit** *agg-rate*
no agg-rate-limit

Context config>service>vpls>sap>ingress

Description This command defines a maximum total rate for all ingress queues on a service SAP or multi-service site. The **agg-rate-limit** command is mutually exclusive with the egress scheduler policy. When an egress scheduler policy is defined, the **agg-rate-limit** command will fail. If the **agg-rate-limit** command is specified, an attempt to bind a **scheduler-policy** to the SAP or multi-service site will fail.

A multi-service site must have a port scope defined that ensures all queues associated with the site are on the same port or channel. If the scope is not set to a port, the **agg-rate-limit** command will fail. Once an agg-rate-limit has been assigned to a multi-service site, the scope cannot be changed to card level.

A port scheduler policy must be applied on the egress port or channel the SAP or multi-service site are bound to in order for the defined agg-rate-limit to take effect. The egress port scheduler enforces the aggregate queue rate as it distributes its bandwidth at the various port priority levels. The port scheduler stops offering bandwidth to member queues once it has detected that the aggregate rate limit has been reached.

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

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

Parameters *agg-rate* — Defines the rate, in Kbps, that the maximum aggregate rate that the queues on the SAP or MSS can operate.

Values 1 — 40000000, max

encap-defined-qos

Syntax **encap-defined-qos**

Context config>service>vpls>sap>egress

Description This command creates a new QoS sub-context in B-VPLS SAP egress context. The user can define encapsulation groups, referred to as encap-group, based on the ISID value in the packet's encapsulation and assign a QoS policy and a scheduler policy or aggregate rate limit to the group.

encap-group

Syntax **encap-group** *group-name* [**type** *group-type*] [**qos-per-member**] [**create**]
no encap-group *group-name*

Context config>service>vpls>sap>egress>encap-defined-qos

Description This command defines an encapsulation group which consists of a group of ISID values. All packets forwarded on the egress of a B-VPLS SAP which payload header matches one of the ISID value in the encap-group will use the same QoS policy instance and scheduler policy or aggregate rate limit instance.

The user adds or removes members to the encap-group one at a time or as a range of contiguous values using the member command. However, when the qos-per-member option is enabled, members must be added or removed one at a time. These members are also referred to as ISID contexts.

The user can configure one or more encap-groups in the egress context of the same B-SAP, thus defining different ISID values and applying each a different SAP egress QoS policy, and optionally a different scheduler policy/agg-rate-limit. Note that ISID values are unique within the context of a B-SAP. The same ISID value cannot be re-used in another encap-group under the same B-SAP but can be re-used in an encap-group under a different B-SAP. Finally, if the user adds to an encap-group an ISID value which is already a member of this encap-group, the command causes no effect. The same if the user attempts to remove an ISID value which is not a member of this encap-group.

Once a group is created, the user will assign a SAP egress QoS policy, and optionally a scheduler policy or aggregate rate limit, using the following commands:

config>service> vpls>sap>egress>encap-defined-qos>encap-group>qos sap-egress-policy-id

config>service> vpls>sap>egress>encap-defined-qos>encap-group>scheduler-policy scheduler-policy-name

config>service> vpls>sap>egress>encap-defined-qos>encap-group>agg-rate-limit kilobits-per-second

Note that a SAP egress QoS policy must first be assigned to the created encap-group before the user can add members to this group. Conversely, the user cannot perform no qos command until all members are deleted from the encap-group.

An explicit or the default SAP egress QoS policy will continue to be applied to the entire B-SAP but this will serve to create the set of egress queues which will be used to store and forward a packet which does not match any of the defined ISID values in any of the encap-groups for this SAP.

Only the queue definition and fc-to-queue mapping from the encap-group SAP egress QoS policy is applied to the ISID members. All other parameters configurable in a SAP egress QoS policy must be inherited from egress QoS policy applied to the B-SAP.

Furthermore, any other CLI option configured in the egress context of the B-SAP will continue to apply to packets matching a member of any encap-group defined in this B-SAP.

The keyword qos-per-member allows the user to specify that a separate queue set instance and scheduler/agg-rate-limit instance will be created for each ISID value in the encap-group. By default, shared instances will be created for the entire encap-group.

Note that when the B-SAP is configured on a LAG port, the ISID queue instances defined by all the encap-groups applied to the egress context of the SAP will be replicated on each member link of the LAG. The set of scheduler/agg-rate-limit instances will be replicated per link or per IOM depending if the adapt-qos option is set to link mode or distribute mode. This is the same behavior as that applied to the entire B-SAP in the current implementation.

The **no** form of this command deletes the encap-group.

Parameters *group-name* — Specifies the name of the encap-group and can be up to 32 ASCII characters in length.

type — This specifies the type of the encapsulation ID used by this encap-group.

Values isid

Default None

qos-per-member — Specifies that a separate queue set instance and scheduler/agg-rate-limit instance will be created for each ISID value in the encap-group.

agg-rate-limit

Syntax **agg-rate-limit** *kilobits-per-second* [**queue-frame-based-accounting**]
no agg-rate-limit

Context config>service>vpls>sap>egress>encap-defined-qos>encap-group

Description

member

Syntax [**no**] **member** *encap-id* [**to** *encap-id*]

Context config>service>vpls>sap>egress>encap-defined-qos>encap-group

Description This command adds or removes a member ISID or a range of contiguous ISID members to an encap-group. The user can add or remove members to the encap-group one at a time or as a range of contiguous values using the member command. However, when the **qos-per-member** option is enabled, members must be added or removed one at a time.

The **no** form of this command removes the single or range of ISID values from the encap-group.

Parameters *encap-id* — The value of the single encap-id or the start encap-id of the range. ISID is the only encap-id supported.

to *encap-id* — The value of the end encap-id of the range. ISID is the only encap-id supported

qos

Syntax **qos** *policy-id*
no qos

Context config>service>vpls>sap>egress>encap-defined-qos>encap-group

Description This command configures the QoS ID.

scheduler-policy

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

Context config>service>vpls>sap>egress>encap-defined-qos>encap-group

Description This command configures the scheduler policy.

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>egress
config>service>vpls>mesh-sdp>ingress
config>service>vpls>spoke-sdp>egress
config>service>vpls>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

hsmda-queue-override

Syntax [**no**] **hsmda-queue-override**

Context config>service>vpls>sap>egress

Description This command enables the context to configure HSMDA queue overrides.

queue

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

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

Description This command configures overrides for a HSMDA queue. The actual valid values are those defined in the given SAP QoS policy.

Parameters *queue-id* — Specifies the queue ID to override.

Values 1 — 8

create — This keyword is mandatory while creating a new queue override.

packet-byte-offset

Syntax **packet-byte-offset** {**add** *add-bytes* | **subtract** *sub-bytes*}
no packet-byte-offset

Context config>service>vpls>sap>egress>hsmda-queue-over

Description This command adds or subtracts the specified number of bytes to the accounting function for each packet handled by the HSMDA queue. Normally, the accounting and leaky bucket functions are based on the Ethernet DLC header, payload and the 4 byte CRC (everything except the preamble and inter-frame gap). As an example, the packet-byte-offset command can be used to add the frame encapsulation overhead (20 bytes) to the queues accounting functions.

The accounting functions affected include:

- Offered High Priority / In-Profile Octet Counter
- Offered Low Priority / Out-of-Profile Octet Counter
- Discarded High Priority / In-Profile Octet Counter
- Discarded Low Priority / Out-of-Profile Octet Counter
- Forwarded In-Profile Octet Counter
- Forwarded Out-of-Profile Octet Counter
- Peak Information Rate (PIR) Leaky Bucket Updates
- Committed Information Rate (CIR) Leaky Bucket Updates
- Queue Group Aggregate Rate Limit Leaky Bucket Updates

The secondary shaper leaky bucket, scheduler priority level leaky bucket and the port maximum rate updates are not affected by the configured packet-byte-offset. Each of these accounting functions are frame based and always include the preamble, DLC header, payload and the CRC regardless of the configured byte offset.

The packet-byte-offset command accepts either add or subtract as valid keywords which define whether bytes are being added or removed from each packet traversing the queue. Up to 31 bytes may be added to the packet and up to 32 bytes may be removed from the packet. An example use case for subtracting bytes from each packet is an IP based accounting function. Given a Dot1Q encapsulation, the command packet-byte-offset subtract 14 would remove the DLC header and the Dot1Q header from the size of each packet for accounting functions only. The 14 bytes are not actually removed from the packet, only the accounting size of the packet is affected.

As inferred above, the variable accounting size offered by the packet-byte-offset command is targeted at the queue and queue group level. The packet-byte-offset, when set, applies to all queues in the queue group. The accounting size of the packet is ignored by the secondary shapers, the scheduling priority level shapers and the scheduler maximum rate. The actual on-the-wire frame size is used for these functions to allow an accurate representation of the behavior of the subscribers packets on an Ethernet aggregation network.

The packet-byte-offset value may be overridden at the queue-group level.

Parameters **add** *add-bytes* **—** Indicates that the byte value should be added to the packet for queue and queue group level accounting functions. Either the **add** or **subtract** keyword must be specified. The corresponding byte value must be specified when executing the packet-byte-offset command. The **add** keyword is mutually exclusive with the **subtract** keyword.

Values 0 — 31

subtract *sub-bytes* **—** Indicates that the byte value should be subtracted from the packet for queue and queue group level accounting functions. The **subtract** keyword is mutually exclusive with the **add** keyword. Either the **add** or **subtract** keyword must be specified. The corresponding byte value must be specified when executing the packet-byte-offset command.

Values 1 — 32

slope-policy

Syntax **slope-policy** *hsmda-slope-policy-name*
no slope-policy

Context config>service>vpls>sap>egress>hsmda-queue-over>queue

Description This command specifies an existing slope policy name.

rate

Syntax **rate** *pir-rate*
no rate

Context config>service>vpls>sap>egress>hsmda-queue-over

Description This command specifies the administrative PIR by the user.

Parameters *pir-rate* — Configures the administrative PIR specified by the user.

Values 1 — 40000000, max

wrr-weight

Syntax **wrr-weight** *value*
no wrr-weight

Context config>service>vpls>sap>egress>hsmda-queue-over>queue

Description This command assigns the weight value to the HSMDA queue.

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

Parameters *percentage* — Specifies the weight for the HSMDA queue.

Values 1— 32

wrr-policy

Syntax **wrr-policy** *hsmda-wrr-policy-name*
no wrr-policy

Context config>service>vpls>sap>egress>hsmda-queue-over

Description This command associates an existing HSMDA weighted-round-robin (WRR) scheduling loop policy to the HSMDA queue.

Parameters *hsmda-wrr-policy-name* — Specifies the existing HSMDA WRR policy name to associate to the queue.

secondary-shaper

secondary-shaper *secondary-shaper-name*
no secondary-shaper

Context config>service>vpls>sap>egress>hsmda-queue-over

Description This command configures an HSMDA secondary shaper. Note that an shaper override can only be configured on an HSMDA SAP.

Parameters *secondary-shaper-name* — Specifies a secondary shaper name up to 32 characters in length.

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.

qinq-mark-top-only

Syntax [**no**] **qinq-mark-top-only**

Context config>service>vpls>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/DEI bit to mark during packet egress. When disabled, both set of P-bits/DEI bit are marked. When enabled, only the P-bits/DEI bit in the top Q-tag are marked.

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

Default no qinq-mark-top-only

policer-control-override

Syntax **policer-control-override** [**create**]
no policer-control-override

Context config>service>vpls>sap>egress

Description This command, within the SAP ingress or egress contexts, creates a CLI node for specific overrides to the applied policer-control-policy. A policy must be applied for a policer-control-overrides node to be created. If the policer-control-policy is removed or changed, the policer-control-overrides node is automatically deleted from the SAP.

The **no** form of the command removes any existing policer-control-policy overrides and the policer-control-overrides node from the SAP.

Default no policer-control-override

Parameters **create —** The create keyword is required when the policer-control-overrides node is being created and the system is configured to expect explicit confirmation that a new object is being created. When the system is not configured to expect explicit confirmation, the create keyword is not required.

max-rate

Syntax **max-rate** {*rate* | **max**}

Context config>service>vpls>sap>egress

Description This command, within the SAP ingress and egress contexts, overrides the root arbiter parent policer max-rate that is defined within the policer-control-policy applied to the SAP.

When the override is defined, modifications to the policer-control-policy max-rate parameter have no effect on the SAP's parent policer until the override is removed using the **no max-rate** command within the SAP.

Parameters *rate* | **max —** Specifies the max rate override in kilobits-per-second or use the maximum.

Values 1 — 20000000 Kbps, max

priority-mbs-thresholds

Syntax **priority-mbs-thresholds**

Context config>service>vpls>sap>egress

Description This command overrides the CLI node contains the configured min-thresh-separation and the various priority level mbs-contribution override commands.

min-thresh-separation

Syntax **min-thresh-separation** *size* [**bytes** | **kilobytes**]

Context config>service>vpls>sap>egress

Description This command within the SAP ingress and egress contexts is used to override the root arbiter's parent policer min-thresh-separation parameter that is defined within the policer-control-policy applied to the SAP.

When the override is defined, modifications to the policer-control-policy min-thresh-separation parameter have no effect on the SAP's parent policer until the override is removed using the no min-thresh-separation command within the SAP.

The no form of the command removes the override and allows the min-thresh-separation setting from the policer-control-policy to control the root arbiter's parent policer's minimum discard threshold separation size.

Default no min-thresh-separation

Parameters **bytes —** Signifies that size is expressed in bytes. The bytes and kilobytes keywords are mutually exclusive and are optionally used to qualify whether size is expressed in bytes or kilobytes. The default is kilobytes.

kilobytes — The size parameter is required when specifying the min-thresh-separation override. It is specified as an integer representing either a number of bytes or kilobytes that are the minimum separation between the parent policer's priority level discard thresholds.

Values 0 — 4194304

Default kilobytes

priority

Syntax [**no**] **priority** *level*

Context config>service>vpls>sap>egress

Description The priority-level level override CLI node contains the specified priority level's mbs-contribution override value.

This node does not need to be created and will not be output in show or save configurations unless an mbs-contribution override exist for level.

Parameters *level* — The level parameter is required when specifying priority-level and identifies which of the parent policer instances priority level's the mbs-contribution is overriding.

Values 1 — 8

mbs-contribution

Syntax **mbs-contribution** *size* [**bytes** | **kilobytes**]

Context config>service>vpls>sap>egress

Description The mbs-contribution override command within the SAP ingress and egress contexts is used to override a parent policer's priority level's mbs-contribution parameter that is defined within the policer-control-policy applied to the SAP. This override allow the priority level's burst tolerance to be tuned based on the needs of the SAP's child policers attached to the priority level.

When the override is defined, modifications to the policer-control-policy priority level's mbs-contribution parameter have no effect on the SAP's parent policer priority level until the override is removed using the no mbs-contribution command within the SAP.

The **no** form of the command removes the override and allows the mbs-contribution setting from the policer-control-policy to control the parent policer's priority level's burst tolerance.

Default no mbs-contribution

Parameters **bytes —** This keyword signifies that size is expressed in bytes.

kilobytes — The optional kilobytes keyword signifies that size is expressed in kilobytes.

Values 1 — 32,000,000,000

policer-control-policy

Syntax **policer-control-policy** *policy-name* [**create**]
no policer-control-policy

Context config>service>vpls>sap>egress

Description This command, within the qos CLI node, is used to create, delete or modify policer control policies. A policer control policy is very similar to the scheduler-policy which is used to manage a set of queues by defining a hierarchy of virtual schedulers and specifying how the virtual schedulers interact to provide an aggregate SLA. In a similar fashion, the policer-control-policy controls the aggregate bandwidth available to a set of child policers. Once created, the policy can be applied to ingress or egress SAPs. The policy may also be applied to the ingress or egress context of a sub-profile.

Policer Control Policy Instances

On the SAP side, an instance of a policy is created each time a policy is applied. When applied to a sub-profile, an instance of the policy is created each time a subscriber successfully maps one or more hosts to the profile per ingress SAP.

Each instance of the policer-control-policy manages the policers associated with the object that owns the policy instance (SAP or subscriber). If a policer on the object is parented to an appropriate arbiter name that exists within the policy, the policer will be managed by the instance. If a policer is not parented or is parented to a non-existent arbiter, the policer will be orphaned and will not be subject to bandwidth control by the policy instance.

Maximum Rate and Root Arbiter

The policer-control-policy supports an overall maximum rate (max-rate) that defines the total amount of bandwidth that may be distributed to all associated child policers. By default, that rate is set to max which provides an unlimited amount of bandwidth to the policers. Once the policy is created, an actual rate should be configured in order for the policy instances to be effective. At the SAP level, the maximum rate may be overridden on a per instance basis. For subscribers, the maximum rate may only be overridden on the subscriber profile which will then be applied to all instances associated with the profile.

The maximum rate is defined within the context of the root arbiter which is always present in a policer-control-policy. The system creates a parent policer which polices the output of all child policers attached to the policy instance to the configured rate. Child policers may be parented directly to the root arbiter (parent root) or parented to one of the tiered arbiters (parent arbiter-name). Since each tiered arbiter must be parented to either another tiered arbiter or the root arbiter (default), every parented child policer is associated with the root arbiter and thus the root arbiter's parent policer.

Parent Policer PIR Leaky Bucket Operation

The parent policer is a single leaky bucket that monitors the aggregate throughput rate of the associated child policers. Forwarded packets increment the bucket by the size of each packet. The rate of the parent policer is implemented as a bucket decrement function which attempts to drain the bucket. If the rate of the packets flowing through the bucket is less than the decrement rate, the bucket does not accumulate depth. Each packet that flows through the bucket is accompanied by a derived discard threshold. If the current depth of the bucket is less than the discard threshold, the packet is allowed to pass through, retaining the colors derived from the packet's child policer. If the current depth is equal to or greater than the threshold value, the packet is colored red and the bucket depth is not incremented by the packet size. Also, any increased bucket depths in the child policer are canceled making any discard event an atomic function between the child and the parent.

Due to the fact that multiple thresholds are supported by the parent policer, the policer control policy is able to protect the throughput of higher priority child policers from the throughput of the lower priority child policers within the aggregate rate.

Tier 1 and Tier 2 Arbiters

As stated above, each child is attached either to the always available root arbiter or to an explicitly created tier 1 or tier 2 arbiter. Unlike the hardware parent policer based root arbiter, the arbiters at tier 1 and tier 2 are only represented in software and are meant to provide an arbitrary hierarchical bandwidth distribution capability. An arbiter created on tier 2 must parent to either to an arbiter on tier 1 or to the root arbiter. Arbiters created on tier 1 always parent to the root arbiter. In this manner, every arbiter ultimately is parented or grand-parented by the root arbiter.

Each tiered arbiter supports an optional rate parameter that defines a rate limit for all child arbiters or child policers associated with the arbiter. Child arbiters and policers attached to the arbiter have a level attribute that defines the strict level at which the child is given bandwidth by the arbiter. Level 8 is the highest and 1 is the lowest. Also a weight attribute defines each child's weight at that strict level in order to determine how bandwidth is distributed to multiple children at that level when insufficient bandwidth is available to meet each child's required bandwidth.

Fair and Unfair Bandwidth Control

Each child policer supports three leaky buckets. The PIR bucket manages the policer's peak rate and maximum burst size, the CIR leaky bucket manages the policer's committed rate (in-profile / out-of-profile) and committed burst size. The third leaky bucket is used by the policer control policy instance to manage the child policer's fair rate (FIR). When multiple child policers are attached to the root arbiter at the same priority level, the policy instance uses each child's FIR bucket rate to control how much of the traffic forwarded by the policer is fair and how much is unfair.

In the simplest case where all the child policers in the same priority level are directly attached to the root arbiter, each child's FIR rate is set according to the child's weight divided by the sum of the active children's weights multiplied by the available bandwidth at the priority level. The result is that the FIR bucket will mark the appropriate amount of traffic for each child as □air□based on the weighted fair output of the policy instance.

The fair/unfair forwarding control in the root parent policer is accomplished by implementing two different discard thresholds for the priority. The first threshold is discard-unfair and the second is discard-all for packet associated with the priority level. As the parent policer PIR bucket fills (due the aggregate forwarded rate being greater than the parent policers PIR decrement rate) and the bucket depth reaches the first threshold, all unfair packets within the priority are discarded. This leaves room in the bucket for the fair packets to be forwarded.

In the more complex case where one or more tiered arbiters are attached at the priority level, the policer control policy instance must consider more than just the child policer weights associated with the attached arbiter. If the arbiter is configured with an aggregate rate limit that its children cannot exceed, the policer control policy instance will switch to calculating the rate each child serviced by the arbiter should receive and enforces that rate using each child policers PIR leaky bucket.

When the child policer PIR leaky bucket is used to limit the bandwidth for the child policer and the child's PIR bucket discard threshold is reached, packets associated with the child policer are discarded. The child policer's discarded packets do not consume depth in the child policer's CIR or FIR buckets. The child policers discarded packets are also prevented from impacting the parent policer and will not consume the aggregate bandwidth managed by the parent policer.

Parent Policer Priority Level Thresholds

As stated above, each child policer is attached either to the root arbiter or explicitly to one of the tier 1 or tier 2 arbiters. When attached directly to the root arbiter, its priority relative to all other child policers is indicated by the parenting level parameter. When attached through one of the tiered arbiters, the parenting hierarchy of the arbiters must be traced through to the ultimate attachment to the root arbiter. The parenting level parameter of the arbiter parented to the root arbiter defines the child policer's priority level within the parent policer.

The priority level is important since it defines the parent policer discard thresholds that will be applied at the parent policer. The parent policer has 8 levels of strict priority and each priority level has its own discard-unfairand discard-all thresholds. Each priority's thresholds are larger than the thresholds of the lower priority levels. This ensures that when the parent policer is discarding, it will be priority sensitive.

To visualize the behavior of the parent policer, picture that when the aggregate forwarding rate of all child policers is currently above the decrement rate of the parent PIR leaky bucket, the bucket depth will increase over time. As the bucket depth increases, it will eventually cross the lowest priority's discard-unfair threshold. If this amount of discard sufficiently lowers the remaining aggregate child policer rate, the parent PIR bucket will hover around this bucket depth. If however, the remaining aggregate child rate is still greater than the decrement rate, the bucket will continue to rise and eventually reach the lowest priority's discard-all threshold which will cause all packets associated with the priority level to be discarded (fair and unfair). Again, if the remaining aggregate child rate is less than or equal to the bucket decrement rate, the parent PIR bucket will hover around this higher bucket depth. If the remaining aggregate child rate is still higher than the decrement rate, the bucket will continue to rise through the remaining priority level discards until equilibrium is achieved.

As noted above, each child's rate feeding into the parent policer is governed by the child policer's PIR bucket decrement rate. The amount of bandwidth the child policer offers to the parent policer will not exceed the child policer's configured maximum rate.

Root Arbiter's Parent Policer's Priority Aggregate Thresholds

Each policer-control-policy root arbiter supports configurable aggregate priority thresholds which are used to control burst tolerance within each priority level. Two values are maintained per priority level; the shared-portion and the fair-portion. The shared-portion represents the amount of parent PIR bucket depth that is allowed to be consumed by both fair and unfair child packets at the priority level. The fair-portion represents the amount of parent PIR bucket depth that only the fair child policer packets may consume within the priority level. It should be noted that the fair and unfair child packets associated with a higher parent policer priority level may also consume the bucket depth set aside for this priority.

While the policy maintains a parent policer default or explicit configurable values for shared-portion and fair-portion within each priority level, it is possible that some priority levels will not be used within the parent policer. Most parent policer use cases require fewer than eight strict priority levels.

In order to derive the actual priority level discard-unfair and discard-all thresholds while only accounting for the actual in-use priority levels, the system maintains a child policer to parent policer association counter per priority level for each policer control policy instance. As a child policer is parented to either the root or a tiered arbiter, the system determines the parent policer priority level for the child policer and increments the association counter for that priority level on the parent policer instance.

The shared-portion for each priority level is affected by the parent policer global min-thresh-separation parameter that defines the minimum separation between any in-use discard thresholds. When more than one child policer is associated with a parent policer priority level, the shared-portion for that priority level will be the current value of min-thresh-separation. When only a single child

policer is associated, the priority level's shared-portion is zero since all packets from the child will be marked fair and the discard-unfair threshold is meaningless. When the association counter is zero, both the shared-portion and the fair-portion for that priority level are zero since neither discard thresholds will be used. Whenever the association counter is greater than 0, the fair-portion for that priority level will be derived from the current value of the priority's mbs-contribution parameter and the global min-thresh-separation parameter.

Each priority level's discard-unfair and discard-alld thresholds are calculated based on an accumulation of lower priorities shared-portions and fair-portions and the priority level's own shared-portion and fair-portion. The base threshold value for each priority level is equal to the sum of all lower priority level's shared-portions and fair-portions. The discard-unfair threshold is the priority level's base threshold plus the priority level's shared-portion. The discard-all threshold for the priority level is the priority level's base threshold plus both the shared-portion and fair-portion values of the priority. As can be seen, an in-use priority level's thresholds are always greater than the thresholds of lower priority levels.

Policer Control Policy Application

A policer-control-policy may be applied on any Ethernet ingress or egress SAP that is associated with a port (or ports in the case of LAG) on an Ethernet MDA attached to an IOM3-XP or IMM module.

The **no** form of the command removes a non-associated policer control policy from the system. The command will not execute when policer-name is currently associated with any SAP or subscriber management sub-profile context.

Default none

Parameters *policy-name* — Each policer-control-policy must be created with a unique policy name. The name must given as policy-name must adhere to the system policy ASCII naming requirements. If the defined policy-name already exists, the system will enter that policy's context for editing purposes. If policy-name does not exist, the system will attempt to create a policy with the specified name. Creating a policy may require use of the create parameter when the system is configured for explicit object creation mode.

create — The keyword is required when a new policy is being created and the system is configured for explicit object creation mode.

policer-override

Syntax [**no**] **policer-override**

Context config>service>vpls>sap>egress

Description This command, within the SAP ingress or egress contexts, is used to create a CLI node for specific overrides to one or more policers created on the SAP through the sap-ingress or sap-egress QoS policies.

The **no** form of the command is used to remove any existing policer overrides.

Default no policer-overrides

policer

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

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

Description This command, within the SAP ingress or egress contexts, is used to create a CLI node for specific overrides to a specific policer created on the SAP through a sap-ingress or sap-egress QoS policy.

The **no** form of the command is used to remove any existing overrides for the specified policer-id.

Parameters *policer-id* — The policer-id parameter is required when executing the policer command within the policer-overrides context. The specified policer-id must exist within the sap-ingress or sap-egress QoS policy applied to the SAP. If the policer is not currently used by any forwarding class or forwarding type mappings, the policer will not actually exist on the SAP. This does not preclude creating an override context for the policer-id.

create — The create keyword is required when a policer policer-id override node is being created and the system is configured to expect explicit confirmation that a new object is being created. When the system is not configured to expect explicit confirmation, the create keyword is not required.

cbs

Syntax **cbs** *size* [**bytes** | *kilobytes*]
no cbs

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

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 setting 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 — 131072 or default

mbs

Syntax **mbs** *size* [**bytes** | **kilobytes**]
no mbs

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

Description This command, within the SAP ingress and egress policer-overrides contexts, is used to override the sap-ingress and sap-egress QoS policy configured mbs parameter for the specified policer-id.

The **no** form of the command is used to restore the policer mbs setting to the policy defined value.

Default no mbs

Parameters **size** — The size parameter is required when specifying mbs override and is expressed as an integer representing the required size in either bytes or kilobytes. The default is kilobytes. The optional byte and kilobyte keywords are mutually exclusive and are used to explicitly define whether size represents bytes or kilobytes.

byte — When byte is defined, the value given for size is interpreted as the queue MBS value given in bytes. When kilobytes is defined, the value is interpreted as the queue MBS value given in kilobytes.

packet-byte-offset

Syntax **packet-byte-offset** {**add** *add-bytes* | **subtract** *sub-bytes*}

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

Description This command, within the SAP ingress and egress policer-overrides contexts, is used to override the sap-ingress and sap-egress QoS policy configured packet-byte-offset parameter for the specified policer-id.

The **no** packet-byte-offset command is used to restore the policer packet-byte-offset setting to the policy defined value.

Default no packet-byte-offset

Parameters **add** *add-bytes* — The add keyword is mutually exclusive to the subtract keyword. Either add or subtract must be specified. When add is defined the corresponding bytes parameter specifies the number of bytes that is added to the size each packet associated with the policer for rate metering, profiling and accounting purposes. From the policer's perspective, the maximum packet size is increased by the amount being added to the size of each packet.

Values 1 — 32

subtract *sub-bytes* — The subtract keyword is mutually exclusive to the add keyword. Either add or subtract must be specified. When subtract is defined the corresponding bytes parameter specifies the number of bytes that is subtracted from the size of each packet associated with the policer for rate metering, profiling and accounting purposes. From the policer's perspective, the maximum packet size is reduced by the amount being subtracted from the size of each packet.

Values 1 — 32

rate

Syntax **rate** {*rate* | **max**} [**cir** {**max** | *rate*}]

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

Description This command within the SAP ingress and egress policer-overrides contexts is used to override the sap-ingress and sap-egress QoS policy configured rate parameters for the specified policer-id.

The **no** rate command is used to restore the policy defined metering and profiling rate to a policer.

Parameters {rate | max} — Specifying the keyword max or an explicit kilobits-per-second parameter directly following the rate override command is required and identifies the policer instance metering rate for the PIR leaky bucket. The kilobits-per-second value must be expressed as an integer and defines the rate in Kilobits-per-second. The integer value is multiplied by 1,000 to derive the actual rate in bits-per-second.

Values 1 — 100,000,000, max

[**cir** {**max** | *rate*} — The optional cir keyword is used to override the policy derived profiling rate of the policer. Specifying the keyword max or an explicit kilobits-per-second parameter directly following the cir keyword is required. The kilobits-per-second value must be expressed as an integer and defines the rate in Kilobits-per-second. The integer value is multiplied by 1,000 to derive the actual rate in bits-per-second.

Values 1 — 100,000,000, max

stat-mode

Syntax **stat-mode** s*tat-mode*
no stat-mode

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

Description The sap-egress QoS policy's policer stat-mode command is used to configure the forwarding plane counters that allow offered, output and discard accounting to occur for the policer. An egress policer has multiple types of offered packets (soft in-profile and out-of-profile from ingress and hard in-profile and out-of-profile due to egress profile overides) and each of these offered types is interacting with the policers metering and profiling functions resulting in colored output packets (green, yellow and red). Due to the potential large number of egress policers, it is not economical to allocate counters in the forwarding plane for all possible offered packet types and output conditions. Many policers will not be configured with a CIR profiling rate and not all policers will receive explicitly re-profiled offered packets. The stat-mode command allows provisioning of the number of counters each policer requires and how the offered packet types and output conditions should be mapped to the counters.

While a no-stats mode is supported which prevents any packet accounting, the use of the policer's parent command requires at the policer's stat-mode to be set at least to the minimal setting so that offered stats are available for the policer's Fair Information Rate (FIR) to be calculated. Once a policer has been made a child to a parent policer, the stat-mode cannot be changed to no-stats unless the policer parenting is first removed.

Each time the policer's stat-mode is changed, any previous counter values are lost and any new counters are set to zero.

Each mode uses a certain number of counters per policer instance that are allocated from the forwarding plane's policer counter resources. If insufficient counters exist to implement a mode on any policer instance, the stat-mode change will fail and the previous mode will continue unaffected for all instances of the policer.

The default stat-mode when a policer is created within the policy is no-stats.

The stat-mode setting defined for the policer in the QoS policy may be overridden on an sla-profile or SAP where the policy is applied. If insufficient policer counter resources exist to implement the override, the stat-mode override command will fail. The previous stat-mode setting active for the policer will continue to be used by the policer.

The no stat-mode command attempts to return the policer's stat-mode setting to no-stats. The command will fail if the policer is currently configured as a child policer using the policer's parent command. The no parent command must first be executed for the no stat-mode command to succeed.

Parameters *stat-mode* — Specifies the mode of statistics collected by this policer.

Values no-stats, minimal, offered-profile-no-cir, offered-profile-cir, offered-total-cir

no-stats — Counter resource allocation: 0
The no-stats mode is the default stat-mode for the policer. The policer does not have any forwarding plane counters allocated and cannot provide offered, discard and forward statistics. A policer using no-stats cannot be a child to a parent policer and the policers parent command will fail.
When collect-stats is enabled, the lack of counters causes the system to generate the following statistics:

a. offered-in = 0
b. offered-out = 0
c. discard-in = 0
d. discard-out = 0
e. forward-in = 0
f. forward-out = 0

Counter 0 indicates that the accounting statistic returns a value of zero.

minimal — Counter resource allocation: 1 The minimal mode allocates 1 forwarding plane offered counter and one traffic manager discard counter. The forwarding counter is derived by subtracting the discard counter from the offered counter. The counters do not differentiate possible offered types (soft or hard profile) and do not count green or yellow output. This does not prevent the policer from supporting different offered packet types and does not prevent the policer from supporting a CIR rate.
This counter mode is useful when only the most basic accounting information is required.
The counters are used in the following manner:

1. offered <= soft-in-profile-out-of-profile, profile in/out
2. discarded <= Same as 1
3. forwarded <= Derived from 1 – 2

When collect-stats is enabled, the counters are used by the system to generate the following statistics:

a. offered-in = 1
b. offered-out = 0
c. discard-in = 2
d. discard-out = 0
e. forward-in = 3
f. forward-out = 0

Counter 0 indicates that the accounting statistic returns a value of zero.

offered-profile-no-cir — Counter resource allocation: 2
The offered-profile-no-cir mode allocates two forwarding plane offered counters and two traffic manager discard counters.
The offered-profile-no-cir mode is most useful when profile based offered, discard and forwarding stats are required from the ingress policer, but a CIR is not being used to recolor the soft in-profile and out-of-profile packets. This mode does not prevent the policer from being configured with a CIR rate.
The counters are used in the following manner:

1. offered-in <= soft-in-profile, profile in
2. offered-out <= soft-out-of-profile, profile out
3. dropped-in <= Same as 1
4. dropped-out <= Same as 2
5. forwarded-in <= Derived from 1 – 3
6. forwarded-out <= Derived from 2 – 4

When collect-stats is enabled, the counters are used by the system to generate the following statistics:

a. offered-in = 1
b. offered-out = 2
c. discard-in = 3
d. discard-out = 4
e. forward-in = 5
f. forward-out = 6

offered-profile-cir — Counter resource allocation: 3

The offered-profile-cir mode allocates three forwarding plane offered counters and three traffic manager discard counters.

The offered-profile-cir mode is most useful when profile based offered, discard and forwarding stats are required from the ingress policer and a CIR rate is being used to recolor the soft in-profile and out-of-profile packets.

The counters are used in the following manner:

1. offered-in-that-stayed-green-or-turned-red <= profile in
2. offered-soft-that-turned-green <= soft-in-profile-out-of-profile
3. offered-soft-or-out-that-turned-yellow-or-red <= soft-in-profile-out-of-profile, profile out
4. dropped-in-that-stayed-green-or-turned-red <= Same as 1
5. dropped-soft-that-turned-green <= Same as 2
6. dropped-soft-or-out-that-turned-yellow-or-red <= Same as 3

7. forwarded-in-that-stayed-green <= Derived from 1 – 4
8. forwarded-soft-that-turned-green <= Derived from 2 – 5
9. forwarded-soft-or-out-that-turned-yellow <= Derived from 3 – 6

When collect-stats is enabled, the counters are used by the system to generate the following statistics:

a. offered-in = 1
b. offered-out = 2 + 3
c. discard-in = 4
d. discard-out = 5 + 6
e. forward-in = 7 + 8
f. forward-out = 9

offered-total-cir — Counter resource allocation: 2

The offered-total-cir mode allocates two forwarding plane offered counters and two traffic manager discard counters.

The offered-total-cir mode is most useful when profile based offered stats are not required from the ingress policer and a CIR rate is being used to recolor the soft in-profile and out-of-profile packets.

The counters are used in the following manner:

1. offered-that-turned-green <= soft-in-profile-out-of-profile, profile in/out
2. offered- that-turned-yellow-or-red<= soft-in-profile-out-of-profile, profile in/out
3. dropped-offered-that-turned-green<= Same as 1
4. dropped-offered-that-turned-yellow-or-red<= Same as 2
5. forwarded-offered-that-turned-green<= Derived from 1 – 3
6. forwarded-offered-that-turned-yellow<= Derived from 2 – 4

When collect-stats is enabled, the counters are used by the system to generate the following statistics:

a. offered-in = 1 + 2 (Or 1 and 2 could be summed on b)
b. offered-out = 0
c. discard-in = 3
d. discard-out = 4
e. forward-in = 5
f. forward-out = 6

Counter 0 indicates that the accounting statistic returns a value of zero.

qos

Syntax **qos** *policy-id* [**shared-queuing** | **multipoint-shared**] [**fp-redirect-group** *queue-group-name* **instance** *instance-id*]
no qos

Context config>service>vpls>sap>ingress

Description This command associates a Quality of Service (QoS) policy with an ingress Service Access Point (SAP).

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 ingress and egress policies on SAP 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 at one time. Attempts to associate a second QoS policy of a given type will return an error.

When an ingress QoS policy is defined on IES ingress IP interface that is bound to a VPLS, the policy becomes associated with every SAP on the VPLS and augments the QoS policy that is defined on each SAP. Packets that are bridged will be processed using the policy defined on the VPLS SAP; packets that are routed will be processed using the policy defined in the IES IP interface-binding context.

By default, if no specific QoS policy is associated with the SAP for ingress or egress, the default QoS policy is used.

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

Default none

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

Values 1 — 65535

shared-queuing — This keyword can only be specified on SAP ingress. Specify the ingress shared queue policy used by this SAP. When the value of this object is null, the SAP will use individual ingress QoS queues, instead of the shared ones.

multipoint-shared — This keyword can only be specified on SAP ingress. Multipoint shared queuing is a superset of shared queuing. When multipoint shared queuing keyword is set, in addition to the unicast packets, multipoint packets also used shared queues.

Ingress unicast service queues are mapped one-for-one with hardware queues and unicast packets traverse the ingress forwarding plane twice, similar to the shared-queuing option. In addition, the multipoint queues defined in the ingress SAP QoS policy are not created. Instead, multipoint packets (broadcast, multicast and unknown unicast destined) are treated to the same dual pass ingress forwarding plane processing as unicast packets.

When the value of this object is null, the SAP will use individual ingress QoS queues, instead of the shared ones.

When the value of this object is null, the SAP will use individual ingress QoS queues, instead of the shared ones.

Values Multipoint or not present.

Default Present (the queue is created as non-multipoint).

fp-redirect-group — This keyword creates an instance of a named queue group template on the ingress forwarding plane of a given IOM/IMM/XMA. The queue-group-name and instance instance-id are mandatory parameters when executing the command. The named queue group template can contain only policers. If it contains queues, then the command will fail.

queue-group-name — Specifies the name of the queue group template to be instantiated on the forwarding plane of the IOM/IMM/XMA, up to 32 characters in length. The queue-group-name must correspond to a valid ingress queue group template name, configured under config>qos>queue- group-templates.

instance-id — Specifies the instance of the named queue group to be created on the IOM/IMMXMA ingress forwarding plane.

qos

Syntax **qos policy-id** [**port-redirect-group** *queue-group-name* **instance** *instance-id*]
no qos

Context config>service>vpls>sap>egress

Description This command associates a Quality of Service (QoS) policy with an egress Service Access Point (SAP).

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. 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 ingress and egress policies on SAP 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 at one time. Attempts to associate a second QoS policy of a given type will return an error.

When an egress QoS policy is associated with an IES IP interface that has been bound to a VPLS, the policy becomes associated with every SAP on the VPLS and augments the egress QoS policy that is defined on each SAP. Packets that are bridged will be processed using the policy defined on the VPLS SAP; packets that are routed will be processed using the policy defined in the IES IP interface-binding context.

By default, if no specific QoS policy is associated with the SAP 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, and the QoS policy reverts to the default.

Default none

Parameters **port-redirect-group —** This keyword associates a SAP egress with an instance of a named queue group template on the egress port of a given IOM/IMM/XMA. The queue-group-name and instance instance-id are mandatory parameters when executing the command.

queue-group-name — Specifies the name of the egress port queue group of the IOM/IMM/XMA, up to 32 characters in length. The queue-group-name must correspond to a valid egress queue group, created under *config>port>ethernet>access>egress*.

instance *instance-id* **—** Specifies the instance of the named egress port queue group on the IOM/IMM/XMA.

Values 1 — 40960

Default 1

queue-override

Syntax **[no] queue-override**

Context config>service>vpls>sap>egress
config>service>vpls>sap>ingress
config>service>vpls>sap>egress>hsmda-queue-over>queue
config>service>vpls>sap>ingress>hsmda-queue-over>queue

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.

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.

adaptation-rule — Specifies the criteria to use to compute the operational CIR and PIR values for this queue, while maintaining a minimum offset.

Values **max** — The **max** (maximum) keyword 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) keyword 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.

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 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 calculate 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>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 setting 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 — 131072 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
config>service>vpls>sap>egress>hsmda-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 packet's RED slope will not force the discard of the packet. Setting proper CBS parameters and controlling CBS oversubscription is one major safeguard to queue starvation (when a queue does not receive its fair share of buffers). Another is properly setting the RED slope parameters for the needs of services on this port or channel.

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

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

Default default

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

Values 0 — 131072 or default

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

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

Default default

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

Values 0 — 131072 or default

rate

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

Context config>service>vpls>sap>egress>queue-override>queue
config>service>vpls>sap>ingress>queue-override>queue
config>service>vpls>sap>egress>hsmda-queue-over>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 — 100000000

Default max

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

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

Default 0

queue-override

Syntax [**no**] **queue-override**

Context config>service>vpls>sap>egress
config>service>vpls>sap>ingress
config>service>vpls>sap>egress>hsmda-queue-over>queue
config>service>vpls>sap>ingress>hsmda-queue-over>queue

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.

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.

adaptation-rule — Specifies the criteria to use to compute the operational CIR and PIR values for this queue, while maintaining a minimum offset.

Values **max** — The **max** (maximum) keyword 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) keyword 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.

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 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 calculate 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>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 setting 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 — 131072 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 packet's RED slope will not force the discard of the packet. Setting proper CBS parameters and controlling CBS oversubscription is one major safeguard to queue starvation (when a queue does not receive its fair share of buffers). Another is properly setting the RED slope parameters for the needs of services on this port or channel.

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

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

Default default

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

Values 0 — 131072 or default

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

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

Default default

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

Values 0 — 131072 or default

rate

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

Context config>service>vpls>sap>egress>queue-override>queue
config>service>vpls>sap>ingress>queue-override>queue
config>service>vpls>sap>egress>hsmda-queue-over>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 — 100000000

Default max

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

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

Default 0

wred-queue-policy

Syntax **wred-queue-policy** *slope-policy-name*
no wred-queue-policy

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

Description The wred-queue-policy command is used on an egress SAP to override the slope policy associated with a WRED queue. When specified, the SAP egress QoS policy derived slope policy is ignored and the configured override slope policy is applied to the WRED queue. The specified *queue-id* must be a WRE- enabled queue to be successful.

The **no** form of the command removes the slope policy override for the WRED queue on the egress SAP.

Parameters *slope-policy-name* — Overrides the SAP Egress QoS policy derived WRED slope policy for the specified queue-id. The defined slope policy must exist or the command will fail.

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 controls that limit each child (other schedulers and queues) associated with the scheduler. Scheduler objects are created within the hierarchical tiers of the policy. It is assumed that each scheduler created will have queues or other schedulers defined as child associations. The scheduler can be a child (take bandwidth from a scheduler in a higher tier, except for schedulers created in tier 1). 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. 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 ate 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 queues 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 — 10000000 or the keyword **max** or **sum** 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

scheduler-policy

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

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

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.

vlan-translation

Syntax **vlan-translation** {**vlan-id | copy-outer**}
no vlan-translation

Context config>service>vpls>sap>ingress

Description This command configures ingress VLAN translation. If enabled with an explicit VLAN value, the preserved VLAN ID will be overwritten with this value. This setting is applicable to Dot1q-encapsulated ports. If enabled with the **copy-outer** keyword, the outer VLAN ID will be copied to the

inner position on QinQ-encapsulated ports. The feature is not supported on default-dot1q SAPs (1/1/1:* and 1/1/1:0), as well as on TopQ SAPs.

The **no** form of this command sets the default value, and no action will be taken.

Default per default the preserved VLAN values will not be overwritten

Parameters *vlan-id* — Specifies the to use the VLAN ID of the SAP.

Values 0 — 4094

copy-outer — Specifies that the outer VLAN ID will be copied to the inner position on QinQ-encapsulated ports

match-qinq-dot1p

Syntax **match-qinq-dot1p** {**top** | **bottom**}
no match-qinq-dot1p de

Context config>service>vpls>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 setting also applies to classification based on the DE indicator bit.

The **no** form of this command reverts the dot1p and de bits matching to the default tag.

By default, the bottom most service delineating Dot1Q tags Dot1P bits are used. Table 10defines the default behavior for Dot1P evaluation.

Table 10: 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

Table 10: Default QinQ and TopQ SAP Dot1P Evaluation (Continued)

Port / SAP Type	Existing Packet Tags	PBits Used for Match
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. The folowing table defines the dot1p evaluation behavior when the top parameter is specified.

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	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. The following table defines the dot1p evaluation behavior when the bottom parameter is specified.

Table 11: 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	TopQ PBits
Null	TopQ (No BottomQ)	TopQ PBits

Table 11: Bottom Position QinQ and TopQ SAP Dot1P Evaluation (Continued)

Port / SAP Type	Existing Packet Tags	PBits Used for Match
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

Egress SAP Type	Ingress Packet Preserved Dot1P State	Marked (or Remarked) PBits
Null	No preserved Dot1P bits	None
Null	Preserved Dot1P bits	Preserved tag PBits remarked using dot1p-value
Dot1Q	No preserved Dot1P bits	New PBits marked using dot1p-value
Dot1Q	Preserved Dot1P bits	Preserved tag PBits remarked using dot1p-value
TopQ	No preserved Dot1P bits	TopQ PBits marked using dot1p-value
TopQ	Preserved Dot1P bits (used as TopQ and BottomQ PBits)	TopQ PBits marked using dot1p-value, BottomQ PBits preserved
QinQ	No preserved Dot1P bits	TopQ PBits and BottomQ PBits marked using dot1p-value
QinQ	Preserved Dot1P bits (used as TopQ and BottomQ PBits)	TopQ PBits and BottomQ PBits marked using dot1p-value

The QinQ and TopQ SAP PBit/DEI bit marking follows the default behavior defined in the table above when **qinq-mark-top-only** is not specified.

The dot1p *dot1p-value* command must be configured without the qinq-mark-top-only parameter to remove the TopQ PBits only marking restriction.

Note that a QinQ-encapsulated Ethernet port can have two different sap types:

- For a TopQ SAP type, only the outer (top) tag is explicitly specified. For example, **sap 1/1/1:10.***
- For QinQ SAP type, both inner (bottom) and outer (top) tags are explicitly specified. For example, **sap 1/1/1:10.100**.

policer-control-policy

Syntax **policer-control-policy** *policy-name* [**create**]
no policer-control-policy

Context config>service>vpls>sap>egress

Description This command, within the qos CLI node, is used to create, delete or modify policer control policies. A policer control policy is very similar to the scheduler-policy which is used to manage a set of queues by defining a hierarchy of virtual schedulers and specifying how the virtual schedulers interact to provide an aggregate SLA. In a similar fashion, the policer-control-policy controls the aggregate bandwidth available to a set of child policers. Once created, the policy can be applied to ingress or egress SAPs. The policy may also be applied to the ingress or egress context of a sub-profile.

Policer Control Policy Instances

On the SAP side, an instance of a policy is created each time a policy is applied. When applied to a sub-profile, an instance of the policy is created each time a subscriber successfully maps one or more hosts to the profile per ingress SAP.

Each instance of the policer-control-policy manages the policers associated with the object that owns the policy instance (SAP or subscriber). If a policer on the object is parented to an appropriate arbiter name that exists within the policy, the policer will be managed by the instance. If a policer is not parented or is parented to a non-existent arbiter, the policer will be orphaned and will not be subject to bandwidth control by the policy instance.

Maximum Rate and Root Arbiter

The policer-control-policy supports an overall maximum rate (max-rate) that defines the total amount of bandwidth that may be distributed to all associated child policers. By default, that rate is set to max which provides an unlimited amount of bandwidth to the policers. Once the policy is created, an actual rate should be configured in order for the policy instances to be effective. At the SAP level, the maximum rate may be overridden on a per instance basis. For subscribers, the maximum rate may only be overridden on the subscriber profile which will then be applied to all instances associated with the profile.

The maximum rate is defined within the context of the root arbiter which is always present in a policer-control-policy. The system creates a parent policer which polices the output of all child policers attached to the policy instance to the configured rate. Child policers may be parented directly to the root arbiter (parent root) or parented to one of the tiered arbiters (parent arbiter-name). Since each tiered arbiter must be parented to either another tiered arbiter or the root arbiter (default), every parented child policer is associated with the root arbiter and thus the root arbiter's parent policer.

Parent Policer PIR Leaky Bucket Operation

The parent policer is a single leaky bucket that monitors the aggregate throughput rate of the associated child policers. Forwarded packets increment the bucket by the size of each packet. The rate of the parent policer is implemented as a bucket decrement function which attempts to drain the bucket. If the rate of the packets flowing through the bucket is less than the decrement rate, the bucket does not accumulate depth. Each packet that flows through the bucket is accompanied by a derived discard threshold. If the current depth of the bucket is less than the discard threshold, the packet is

allowed to pass through, retaining the colors derived from the packet's child policer. If the current depth is equal to or greater than the threshold value, the packet is colored red and the bucket depth is not incremented by the packet size. Also, any increased bucket depths in the child policer are canceled making any discard event an atomic function between the child and the parent.

Due to the fact that multiple thresholds are supported by the parent policer, the policer control policy is able to protect the throughput of higher priority child policers from the throughput of the lower priority child policers within the aggregate rate.

Tier 1 and Tier 2 Arbiters

As stated above, each child is attached either to the always available root arbiter or to an explicitly created tier 1 or tier 2 arbiter. Unlike the hardware parent policer based root arbiter, the arbiters at tier 1 and tier 2 are only represented in software and are meant to provide an arbitrary hierarchical bandwidth distribution capability. An arbiter created on tier 2 must parent to either to an arbiter on tier 1 or to the root arbiter. Arbiters created on tier 1 always parent to the root arbiter. In this manner, every arbiter ultimately is parented or grand-parented by the root arbiter.

Each tiered arbiter supports an optional rate parameter that defines a rate limit for all child arbiters or child policers associated with the arbiter. Child arbiters and policers attached to the arbiter have a level attribute that defines the strict level at which the child is given bandwidth by the arbiter. Level 8 is the highest and 1 is the lowest. Also a weight attribute defines each child's weight at that strict level in order to determine how bandwidth is distributed to multiple children at that level when insufficient bandwidth is available to meet each child's required bandwidth.

Fair and Unfair Bandwidth Control

Each child policer supports three leaky buckets. The PIR bucket manages the policer's peak rate and maximum burst size, the CIR leaky bucket manages the policer's committed rate (in-profile / out-of-profile) and committed burst size. The third leaky bucket is used by the policer control policy instance to manage the child policer's fair rate (FIR). When multiple child policers are attached to the root arbiter at the same priority level, the policy instance uses each child's FIR bucket rate to control how much of the traffic forwarded by the policer is fair and how much is unfair.

In the simplest case where all the child policers in the same priority level are directly attached to the root arbiter, each child's FIR rate is set according to the child's weight divided by the sum of the active children's weights multiplied by the available bandwidth at the priority level. The result is that the FIR bucket will mark the appropriate amount of traffic for each child as □air□based on the weighted fair output of the policy instance.

The fair/unfair forwarding control in the root parent policer is accomplished by implementing two different discard thresholds for the priority. The first threshold is discard-unfair and the second is discard-all for packet associated with the priority level. As the parent policer PIR bucket fills (due the aggregate forwarded rate being greater than the parent policers PIR decrement rate) and the bucket depth reaches the first threshold, all unfair packets within the priority are discarded. This leaves room in the bucket for the fair packets to be forwarded.

In the more complex case where one or more tiered arbiters are attached at the priority level, the policer control policy instance must consider more than just the child policer weights associated with the attached arbiter. If the arbiter is configured with an aggregate rate limit that its children cannot exceed, the policer control policy instance will switch to calculating the rate each child serviced by the arbiter should receive and enforces that rate using each child policers PIR leaky bucket.

When the child policer PIR leaky bucket is used to limit the bandwidth for the child policer and the child's PIR bucket discard threshold is reached, packets associated with the child policer are discarded. The child policer's discarded packets do not consume depth in the child policer's CIR or

FIR buckets. The child policers discarded packets are also prevented from impacting the parent policer and will not consume the aggregate bandwidth managed by the parent policer.

Parent Policer Priority Level Thresholds

As stated above, each child policer is attached either to the root arbiter or explicitly to one of the tier 1 or tier 2 arbiters. When attached directly to the root arbiter, its priority relative to all other child policers is indicated by the parenting level parameter. When attached through one of the tiered arbiters, the parenting hierarchy of the arbiters must be traced through to the ultimate attachment to the root arbiter. The parenting level parameter of the arbiter parented to the root arbiter defines the child policer's priority level within the parent policer.

The priority level is important since it defines the parent policer discard thresholds that will be applied at the parent policer. The parent policer has 8 levels of strict priority and each priority level has its own discard-unfairand discard-all thresholds. Each priority's thresholds are larger than the thresholds of the lower priority levels. This ensures that when the parent policer is discarding, it will be priority sensitive.

To visualize the behavior of the parent policer, picture that when the aggregate forwarding rate of all child policers is currently above the decrement rate of the parent PIR leaky bucket, the bucket depth will increase over time. As the bucket depth increases, it will eventually cross the lowest priority's discard-unfair threshold. If this amount of discard sufficiently lowers the remaining aggregate child policer rate, the parent PIR bucket will hover around this bucket depth. If however, the remaining aggregate child rate is still greater than the decrement rate, the bucket will continue to rise and eventually reach the lowest priority's discard-all threshold which will cause all packets associated with the priority level to be discarded (fair and unfair). Again, if the remaining aggregate child rate is less than or equal to the bucket decrement rate, the parent PIR bucket will hover around this higher bucket depth. If the remaining aggregate child rate is still higher than the decrement rate, the bucket will continue to rise through the remaining priority level discards until equilibrium is achieved.

As noted above, each child's rate feeding into the parent policer is governed by the child policer's PIR bucket decrement rate. The amount of bandwidth the child policer offers to the parent policer will not exceed the child policer's configured maximum rate.

Root Arbiter's Parent Policer's Priority Aggregate Thresholds

Each policer-control-policy root arbiter supports configurable aggregate priority thresholds which are used to control burst tolerance within each priority level. Two values are maintained per priority level; the shared-portion and the fair-portion. The shared-portion represents the amount of parent PIR bucket depth that is allowed to be consumed by both fair and unfair child packets at the priority level. The fair-portion represents the amount of parent PIR bucket depth that only the fair child policer packets may consume within the priority level. It should be noted that the fair and unfair child packets associated with a higher parent policer priority level may also consume the bucket depth set aside for this priority.

While the policy maintains a parent policer default or explicit configurable values for shared-portion and fair-portion within each priority level, it is possible that some priority levels will not be used within the parent policer. Most parent policer use cases require fewer than eight strict priority levels.

In order to derive the actual priority level discard-unfair and discard-all thresholds while only accounting for the actual in-use priority levels, the system maintains a child policer to parent policer association counter per priority level for each policer control policy instance. As a child policer is parented to either the root or a tiered arbiter, the system determines the parent policer priority level for the child policer and increments the association counter for that priority level on the parent policer instance.

The shared-portion for each priority level is affected by the parent policer global min-thresh-separation parameter that defines the minimum separation between any in-use discard thresholds. When more than one child policer is associated with a parent policer priority level, the shared-portion for that priority level will be the current value of min-thresh-separation. When only a single child policer is associated, the priority level's shared-portion is zero since all packets from the child will be marked fair and the discard-unfair threshold is meaningless. When the association counter is zero, both the shared-portion and the fair-portion for that priority level are zero since neither discard thresholds will be used. Whenever the association counter is greater than 0, the fair-portion for that priority level will be derived from the current value of the priority's mbs-contribution parameter and the global min-thresh-separation parameter.

Each priority level's discard-unfair and discard-alld thresholds are calculated based on an accumulation of lower priorities shared-portions and fair-portions and the priority level's own shared-portion and fair-portion. The base threshold value for each priority level is equal to the sum of all lower priority level's shared-portions and fair-portions. The discard-unfair threshold is the priority level's base threshold plus the priority level's shared-portion. The discard-all threshold for the priority level is the priority level's base threshold plus both the shared-portion and fair-portion values of the priority. As can be seen, an in-use priority level's thresholds are always greater than the thresholds of lower priority levels.

Policer Control Policy Application

A policer-control-policy may be applied on any Ethernet ingress or egress SAP that is associated with a port (or ports in the case of LAG).

The **no** form of the command removes a non-associated policer control policy from the system. The command will not execute when policer-name is currently associated with any SAP or subscriber management sub-profile context.

Default none

Parameters *policy-name* — Each policer-control-policy must be created with a unique policy name. The name must given as policy-name must adhere to the system policy ASCII naming requirements. If the defined policy-name already exists, the system will enter that policy's context for editing purposes. If policy-name does not exist, the system will attempt to create a policy with the specified name. Creating a policy may require use of the create parameter when the system is configured for explicit object creation mode.

create — The keyword is required when a new policy is being created and the system is configured for explicit object creation mode.

authentication-policy

Syntax **authentication-policy** *name*
no authentication-policy

Context config>service>vpls>sap

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.

accounting-policy

Syntax **accounting-policy** *acct-policy-id*
no accounting-policy

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

app-profile

Syntax **app-profile** *app-profile-name*
no app-profile

Context config>service>vpls>spoke-sdp

Description This command configures the application profile name.

Parameters *app-profile-name* — Specifies an existing application profile name configured in the **config>app-assure>group>policy** context.

collect-stats

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

Context config>service>vpls>spoke-sdp
config>service>vpls>mesh-sdp
config>service>vpls>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 no collect-stats

VPLS Template Commands

template

Syntax **template**

Context config>service

Description This is the node for service templates.

vpls-template

Syntax **vpls-template** *name/id* **create**
[no] vpls-template *name/id*

Context config>service>template

Description This command is used to create a vpls-template to be used to auto-instantiate a range of VPLS services. Only certain existing VPLS attributes specified in the command reference section can be changed in the vpls-template, not in the instantiated VPLS. The following attributes will be automatically set in the instantiated VPLSes (no template configuration necessary) and the operator cannot change these values.

vpn-id: none

description: “Service <svc id> auto-generated by control VPLS <svc-id>”

service-name: “Service <svc id>” (Auto-generated)

shutdown: no shutdown

Following existing attributes can be set by the user in the instantiated VPLSes:

[no] sap

All the other VPLS attributes are not supported.

Parameters *name/id* — Specifies the name in ASCII or the template ID.

Values name: ASCII string

Values ID: [1..2147483647]

vpls-sap-template

Syntax **vpls-sap-template** *name/id* **create**
[no] vpls-sap-template *name/id*

Context config>service>template

Description This is the command used to create a SAP template to be used in a vpls-template. Only certain existing VPLS SAP attributes can be changed in the vpls-sap-template, not in the instantiated VPLS SAP

Following SAP attributes will be set in the instantiated saps (no configuration allowed):

description: ”Sap <sap-id> controlled by MVRP service <svc id>” – auto generated

shutdown: no shutdown

Parameters *name/id* — Specifies the name in ASCII or the template ID.

Values 1..2147483647

mac-move-level

Syntax **mac-move-level** {**primary | secondary**}
no mac-move-level

Context config>service>template>vpls-sap-template

Description When a sap is instantiated using vpls-sap-template, if the MAC move feature is enabled at VPLS level, the command mac-move-level indicates whether the sap should be populated as primary-port, secondary-port or tertiary-port in the instantiated VPLS.

Default no mac-move-level; SAP is populated as a tertiary-port

temp-flooding

Syntax **temp-flooding flood-time**
no temp-flooding

Context config>service>vpls

config>service>template>vpls-template

Description The temporary flooding is designed to minimize failover times by eliminating the time it takes to flush the MAC tables and if MVRP is enabled the time it takes for MVRP registration. Temporary flooding is initiated only upon xSTP TCN reception. During this procedure while the MAC flush takes place the frames received on one of the VPLS SAPs/pseudowires are flooded in a VPLS context which for MVRP case includes also the unregistered MVRP trunk ports. Note that the MAC Flush action is initiated by the STP TCN reception or if MVRP is enabled for the data VPLS, by the reception of a MVRP New message for the SVLAN ID associated with the data VPLS. As soon as the MAC Flush is done, regardless of whether the temp-flooding timer expired or not, traffic will be delivered according to the regular FIB content which may be built from MAC Learning or based on MVRP registrations. This command provides a flood-time value that configures a fixed amount of time, in seconds, during which all traffic is flooded (BUM or known unicast) as a safety mechanism. Once the flood-time expires, traffic will be delivered according to the regular FIB content which may be built from MAC Learning or based on MVRP registrations. The temporary flooding timer should be configured in such a way to allow auxiliary processes like MAC Flush, MMRP and/or MVRP to complete/converge. The temporary flooding behavior applies to regular VPLS, VPLS instantiated with VPLS-template, IVPLS and BVPLS when MMRP is disabled.

The **no** form of the command disables the temporary flooding behavior.

Default no temp-flooding

Parameters *flood-time* — Specifies the flood time, in seconds.

Values 3 — 600

Provider Tunnel Commands

provider-tunnel

Syntax **provider-tunnel**

Context configure>service>vpls

Description This command creates the context to configure the use of a P2MP LSP for forwarding Broadcast, Unicast unknown and Multicast (BUM) packets of a VPLS or B-VPLS instance. The P2MP LSP is referred to as the Provider Multicast Service Interface (PMSI).

inclusive

Syntax **inclusive**

Context configure>service>vpls>provider-tunnel

Description This command creates the context to configure the use of a P2MP LSP as the default tree for forwarding Broadcast, Unicast unknown, and Multicast (BUM) packets of a VPLS or B-VPLs instance. The P2MP LSP is referred to, in this case, as the Inclusive Provider Multicast Service Interface (I-PMSI).

When enabled, this feature relies on BGP Auto-Discovery (BGP-AD) to discover the PE nodes participating in a given VPLS/B-VPLS instance. The AD route contains the information required to signal both the point-to-point (P2P) PWs used for forwarding unicast known Ethernet frames and the RSVP or mLDP P2MP LSP used to forward the BUM frames.

The root node signals the RSVP P2MP LSP based on an LSP template associated with the I-PMSI at configuration time. The leaf node will join automatically the P2MP LSP, which matches the I-PMSI tunnel information discovered via BGP-AD.

With a mLDP I-PMSI, each leaf node will initiate the signaling of the mLDP P2MP LSP upstream using the P2MP FEC information in the I-PMSI tunnel information discovered via BGP-AD.

If IGMP or PIM snooping are configured on the VPLS/B-VPLS instance, multicast packets matching a L2 multicast Forwarding Information Base (FIB) record will also be forwarded over the P2MP LSP.

The user enables the use of an RSVP P2MP LSP as the I-PMSI for forwarding Ethernet BUM and IP multicast packets in a VPLS/B-VPLS instance using the following commands:

config>service>vpls [b-vpls]>provider-tunnel>inclusive>rsvp>lsp-template *p2mp-lsp-template-name*

The user enables the use of an LDP P2MP LSP as the I-PMSI for forwarding Ethernet BUM and IP multicast packets in a VPLS instance using the following command:

config>service>vpls [b-vpls]>bum-forwarding>provider-tunnel>inclusive>mldp

After the user performs a ‘no shutdown’ under the context of the inclusive node and the expiration of a delay timer, BUM packets will be forwarded over an automatically signaled mLDP P2MP LSP or over an automatically signaled instance of the RSVP P2MP LSP specified in the LSP template.

The user can specify if the node is both root and leaf in the VPLS instance:

config>service>vpls [b-vpls]>provider-tunnel>inclusive>root-and-leaf

The **root-and-leaf** command is required otherwise this node will behave as a leaf only node by default. When the node is leaf only for the I-PMSI of type P2MP RSVP LSP, no PMSI Tunnel Attribute is included in BGP-AD route update messages and thus no RSVP P2MP LSP is signaled but the node can join RSVP P2MP LSP rooted at other PE nodes participating in this VPLS/B-VPLS service. Note that the user must still configure a LSP template even if the node is a leaf only. For the I-PMSI of type mLDP, the leaf-only node will join I-PMSI rooted at other nodes it discovered but will not include a PMSI Tunnel Attribute in BGP-AD route update messages. This way a leaf only node will forward packets to other nodes in the VPLS/B-VPLS using the point-to-point spoke-sdp's.

Note that BGP-AD must have been enabled in this VPLS/B-VPLS instance or the execution of the 'no shutdown" command under the context of the inclusive node is failed and the I-PMSI will not come up. Also note that this feature is not supported with BGP-VPLS. As such, if both BGP-VPLS and BGP-AD are enabled, the execution of the 'no shutdown" command under the context of the inclusive node is also failed. Also, if the I-PMSI is enabled the execution of the 'no shutdown" command under BGP-VPLS is failed.

Any change to the parameters of the I-PMSI, such as disabling the P2MP LSP type or changing the LSP template requires that the inclusive node be first shutdown. The LSP template is configured in MPLS.

If the P2MP LSP instance goes down, VPLS/B-VPLS immediately reverts the forwarding of BUM packets to the P2P PWs. The user can however restore at any time the forwarding of BUM packets over the P2P PWs by performing a 'shutdown' under the context of the inclusive node.

This feature is supported with VPLS, H-VPLS, and B-VPLS. It is not supported with I-VPLS and Routed VPLS. It is also not supported with BGP-VPLS.

data-delay-interval

Syntax **data-delay-interval** *seconds*
no data-delay-interval

Context configure>service>vpls>provider-tunnel>inclusive

Description This command configures the I-PMSI data delay timer.

This delay timer is intended to allow time for the RSVP control plane to signal and bring up the S2L sub-LSP to each destination PE participating in the VPLS/B-VPLS service. The delay timer is started as soon as the P2MP LSP instance becomes operationally up after the user performed a 'no shutdown' under the inclusive node, i.e., as soon as the first S2L sub-LSP is up. In general, it is started when the P2MP LSP instance transitions from the operationally down state to the up state.

For a mLDP P2MP LSP, the delay timer is started as soon as the P2MP FEC corresponding to the I-PMSI is resolved and installed at the root node. Note that the user must factor in the value configured in the data-delay-interval at the root node any delay configured in IGP-LDP sync timer (config>router>interface>ldp-sync-timer) on interfaces over the network. This is because the mLDP P2MP LSP may move to a different interface at the expiry of this timer since the routing upstream of the LDP Label Mapping message may change when this timer expires and the interface metric is restored.

At the expiry of this timer, the VPLS/B-VPLS will begin forwarding of BUM packets over the P2MP LSP instance even if not all the S2L paths are up.

The **no** version of this command re-instates the default value for this delay timer.

Parameters *seconds* — The delay time value in seconds.

Values 3—180 seconds

Default 15 seconds

mldp

Syntax [**no**] **mldp**

Context configure>service>vpls>provider-tunnel>inclusive

Description This command creates the context to configure the parameters of an LDP P2MP LSP used for forwarding Broadcast, Unicast unknown and Multicast (BUM) packets of a VPLS or B-VPLs instance.

root-and-leaf

Syntax [**no**] **root-and-leaf**

Context configure>service>vpls>provider-tunnel>inclusive

Description This command configures the node to operate as both root and leaf of the I-PMSI in a given VPLS/B-VPLS instance.

By default, a node will behave as a leaf only node. When the node is leaf only for the I-PMSI of type P2MP RSVP LSP, no PMSI Tunnel Attribute is included in BGP-AD route update messages and thus no RSVP P2MP LSP is signaled but the node can join RSVP P2MP LSP rooted at other PE nodes participating in this VPLS/B-VPLS service. Note that the user must still configure a LSP template even if the node is a leaf only.

For the I-PMSI of type mLDP, the leaf-only node will join I-PMSI rooted at other nodes it discovered but will not include a PMSI Tunnel Attribute in BGP-AD route update messages. This way a leaf only node will forward packets to other nodes in the VPLS/B-VPLS using the point-to-point spoke-sdp's..

The **no** version of this command re-instates the default value.

rsvp

Syntax [**no**] **rsvp**

Context configure>service>vpls>provider-tunnel>inclusive

Description This command creates the context to configure the parameters of an RSVP P2MP LSP used for forwarding Broadcast, Unicast unknown and Multicast (BUM) packets of a VPLS or B-VPLS instance.

lsp-template

Syntax **lsp-template** *p2mp-lsp-template-name*
no lsp-template

Context configure>service>vpls>provider-tunnel>inclusive>rsvp

Description This command specifies the template name of the RSVP P2MP LSP instance to be used by the leaf node or the root-and-leaf node that participates in BGP-AD VPLS. The P2MP LSP is referred to as the Inclusive Provider Multicast Service Interface (I-PMSI).

After the user performs a "no shutdown" under the context of the inclusive node and the delay timer expires, BUM packets will be forwarded over an automatically signaled instance of the RSVP P2MP LSP specified in the LSP template.

The **no** version of this command removes the P2MP LSP template from the I-PMIS configuration.

Parameters *p2mp-lsp-template-name* — The name of the P2MP LSP template. This is a string of 32 characters maximum.

Default None

VPLS SDP Commands

mesh-sdp

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

Context config>service>vpls

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 a valid 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.

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 is 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 bindings 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 **ethernet** and **vlan** 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** and **vlan** keywords are mutually exclusive. When the VC type is not defined then the default is Ethernet for mesh SDP bindings.

spoke-sdp

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

endpoint — Specifies the service endpoint to which this SDP bind is attached. The service ID of the SDP binding must match the service ID of the service endpoint.

no endpoint — removes the association of a spoke SDP with an explicit endpoint name.

control-word

Syntax [**no**] **control word**

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

Description This command enables the use of the control word on pseudowire packets in VPLS and enables the use of the control word individually on each mesh SDP or spoke SDP. By default, the control word is disabled. When the control word is enabled, all VPLS packets, including the BPDU frames, are encapsulated with the control word when sent over the pseudowire. The T-LDP control plane behavior is the same as in the implementation of control word for VLL services. The configuration for the two directions of the Ethernet pseudowire should match. The **no** form of the command reverts the mesh SDP or spoke SDP to the default behavior of not using the control word.

Default no control word

egress

Syntax **egress**

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

Description This command configures the egress SDP context.

qos

Syntax **qos** *network-policy-id* **port-redirect-group** *queue-group-name* **instance** *instance-id*
no qos [*network-policy-id*]

Context configure>service>apipe>spoke-sdp>egress
configure>service>cpipe>spoke-sdp>egress
configure>service>epipe>spoke-sdp>egress
configure>service>fpipe>spoke-sdp>egress
configure>service>ipipe>spoke-sdp>egress
config>service>vpls>spoke-sdp>egress
config>service>vpls>mesh-sdp>egress
config>service>pw-template>egress
config>service>vprn>interface>spoke-sdp>egress
config>service>ies>interface>spoke-sdp>egress

Description This command is used to redirect pseudowire packets to an egress port queue-group for the purpose of shaping.

The egress pseudowire shaping provisioning model allows the mapping of one ore more pseudowires to the same instance of queues, or policers and queues, which are defined in the queue-group template.

Operationally, the provisioning model consists of the following steps:

1. Create an egress queue-group template and configure queues only or policers and queues for each FC that needs to be redirected.
2. Apply the queue-group template to the network egress context of all ports where there exists a network IP interface on which the pseudowire packets can be forwarded. This creates one instance of the template on the egress of the port. One or more instances of the same template can be created.
3. Configure FC-to-policer or FC-to-queue mappings together with the redirect to a queue-group in the egress context of a network QoS policy. No queue-group name is specified in this step, which means the same network QoS policy can redirect different pseudowires to different queue-group templates.
4. Apply this network QoS policy to the egress context of a spoke-SPD inside a service or to the egress context of a pseudowire template and specify the redirect queue-group name.

One or more spoke-SPDs can have their FCs redirected to use queues only or queues and policers in the same queue-group instance.

The following are the constraints and rules of this provisioning model:

1. When a pseudowire FC is redirected to use a queue or a policer and a queue in a queue-group and the queue-group name does not exist, the association is failed at the time the user associates the egress context of a spoke-SPD to the named queue-group. In such a case, the pseudowire packet will be fed directly to the corresponding egress queue for that FC used by the IP network interface on which the pseudowire packet is forwarded. This queue can be a queue-group queue, or the egress shared queue for that FC defined in the network-queue policy applied to the egress of this port. This is the existing implementation and default behavior for a pseudowire packet.
2. When a pseudowire FC is redirected to use a queue or a policer, and a queue in a queue-group and the queue-group name exists, but the policer-id and/or the queue-id is not defined

in the queue-group template, the association is failed at the time the user associates the egress context of a spoke-SPD to the named queue-group. In such a case, the pseudowire packet will be fed directly to the corresponding egress queue for that FC used by the IP network interface the pseudowire packet is forwarded on.

3. When a pseudowire FC is redirected to use a queue, or a policer and a queue in a queue-group, and the queue-group name exists and the policer-id or policer-id plus queue-id exist, it is not required to check that an instance of that queue-group exists in all egress network ports which have network IP interfaces. The handling of this is dealt with in the data path as follows:

   a When a pseudowire packet for that FC is forwarded and an instance of the referenced queue-group name exists on that egress port, the packet is processed by the queue-group policer and will then be fed to the queue-group queue.

   b When a pseudowire packet for that FC is forwarded and an instance of the referenced queue-group name does not exist on that egress port, the pseudowire packet will be fed directly to the corresponding egress shared queue for that FC defined in the network-queue policy applied to the egress of this port.

4. If a network QoS policy is applied to the egress context of a pseudowire, any pseudowire FC, which is not explicitly redirected in the network QoS policy, will have the corresponding packets feed directly the corresponding the egress shared queue for that FC defined in the network-queue policy applied to the egress of this port.

When the queue-group name the pseudowire is redirected to exists and the redirection succeeds, the marking of the packet DEI/dot1.p/DSCP and the tunnel DEI/dot1.p/DSCP/EXP is performed; according to the relevant mappings of the (FC, profile) in the egress context of the network QoS policy applied to the pseudowire. This is true regardless, wether an instance of the queue-group exists or not on the egress port to which the pseudowire packet is forwarded. If the packet profile value changed due to egress child policer CIR profiling, the new profile value is used to mark the packet DEI/dot1.p and the tunnel DEI/dot1.p/EXP, but the DSCP is not modified by the policer operation.

When the queue-group name the pseudowire is redirected does not exist, the redirection command is failed. In this case, the marking of the packet DEI/dot1.p/DSCP and the tunnel DEI/dot1.p/DSCP/EXP fields is performed according to the relevant commands in the egress context of the network QoS policy applied to the network IP interface to which the pseudowire packet is forwarded.

The **no** version of this command removes the redirection of the pseudowire to the queue-group.

Parameters *network-policy-id* — Specifies the network policy identification. The value uniquely identifies the policy on the system.

Values 1 — 65535

queue-redirect-group *queue-group-name* **—** This optional parameter specifies that the *queue-group-name* will be used for all egress forwarding class redirections within the network QoS policy ID. The specified *queue-group-name* must exist as a port egress queue group on the port associated with the IP interface.

egress-instance *instance-id* **—** Specifies the identification of a specific instance of the queue-group.

Values 1 — 16384

ingress

Syntax **ingress**

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

Description This command configures the ingress SDP context.

qos

Syntax **qos** *network-policy-id* **fp-redirect-group** *queue-group-name* **instance** *instance-id*]
no qos

Context configure>service>apipe>spoke-sdp>ingress
configure>service>cpipe>spoke-sdp>ingress
configure>service>epipe>spoke-sdp>ingress
configure>service>fpipe>spoke-sdp>ingress
configure>service>ipipe>spoke-sdp>ingress
config>service>vpls>spoke-sdp>ingress
config>service>vpls>mesh-sdp>ingress
config>service>pw-template>ingress
config>service>vprn>interface>spoke-sdp>ingress
config>service>ies>interface>spoke-sdp>ingress

Description This command is used to redirect pseudowire packets to an ingress forwarding plane queue-group for the purpose of rate-limiting.

The ingress pseudowire rate-limiting feature uses a policer in queue-group provisioning model. This model allows the mapping of one or more pseudowires to the same instance of policers, which are defined in a queue-group template.

Operationally, the provisioning model in the case of the ingress pseudowire shaping feature consists of the following steps:

1. Create an ingress queue-group template and configure policers for each FC that needs to be redirected and optionally, for each traffic type (unicast or multicast).
2. Apply the queue-group template to the network ingress forwarding plane where there exists a network IP interface to which the pseudowire packets can be received. This creates one instance of the template on the ingress of the FP. One or more instances of the same template can be created.
3. Configure FC-to-policer mappings together with the policer redirect to a queue-group in the ingress context of a network QoS policy. No queue-group name is specified in this step, which means the same network QoS policy can redirect different pseudowires to different queue-group templates.
4. Apply this network QoS policy to the ingress context of a spoke-SDP inside a service, or to the ingress context of a pseudowire template, and specify the redirect queue-group name.
5. One or more spoke-SDPs can have their FCs redirected to use policers in the same policer queue-group instance.

The following are the constraints and rules of this provisioning model when used in the ingress pseudowire rate-limiting feature:

1. When a pseudowire FC is redirected to use a policer in a named policer queue-group and the queue-group name does not exist, the association is failed at the time the user associates the ingress context of a spoke-SDP to the named queue-group. In such a case, the pseudowire packet will feed directly the ingress network shared queue for that FC defined in the network-queue policy applied to the ingress of the MDA/FP.
2. When a pseudowire FC is redirected to use a policer in a named policer queue-group and the queue-group name exists but the policer-id is not defined in the queue-group template, the association is failed at the time the user associates the ingress context of a spoke-SPD to the named queue-group. In such a case, the pseudowire packet will feed directly the ingress network shared queue for that FC defined in the network-queue policy applied to the ingress of the MDA/FP.
3. When a pseudowire FC is redirected to use a policer in a named policer queue-group and the queue-group name exists and the policer-id is defined in the queue-group template, it is not required to check that an instance of that queue-group exists in all ingress FPs which have network IP interfaces. The handling of this is dealt with in the data path as follows:
    a When a pseudowire packet for that FC is received and an instance of the referenced queue-group name exists on that FP, the packet is processed by the policer and will then feed the per-FP ingress shared queues referred to as *policer-output-queues*.
    b When a pseudowire packet for that FC is received and an instance of the referenced queue-group name does not exist on that FP, the pseudowire packets will be fed directly into the corresponding ingress network shared queue for that FC defined in the network-queue policy applied to the ingress of the MDA/FP.
4. If a network QoS policy is applied to the ingress context of a pseudowire, any pseudowire FC which is not explicitly redirected in the network QoS policy will have the corresponding packets feed directly the ingress network shared queue for that FC defined in the network-queue policy applied to the ingress of the MDA/FP.
5. If no network QoS policy is applied to the ingress context of the pseudowire, then all packets of the pseudowire will feed:
    a the ingress network shared queue for the packet FC defined in the network-queue policy applied to the ingress of the MDA/FP. This is the default behavior.
    b a queue-group policer followed by the per-FP ingress shared queues referred to as *policer-output-queues* if the ingress context of the network IP interface from which the packet is received is redirected to a queue-group (csc-policing). The only exceptions to this behavior are for packets received from a IES/VPRN spoke interface and from an R-VPLS spoke-SPD, which is forwarded to the R-VPLS IP interface. In these two cases, the ingress network shared queue for the packet FC defined in the network-queue policy applied to the ingress of the MDA/FP is used.

When a pseudowire is redirected to use a policer queue-group, the classification of the packet for the purpose of FC and profile determination is performed according to default classification rule or the QoS filters defined in the ingress context of the network QoS policy applied to the pseudowire. This is true regardless of wether an instance of the named policer queue-group exists on the ingress FP on which the pseudowire packet is received. The user can apply a QoS filter matching the dot1.p in the VLAN tag corresponding to the Ethernet port encapsulation, the EXP in the outer label when the tunnel is an LSP, the DSCP in the IP header if the tunnel encapsulation is GRE, and the DSCP in the

payload IP header if the user enabled the **ler-use-dscp** option and the pseudowire terminates in IES or VPRN service (spoke-interface).

When the policer queue-group name the pseudowire is redirected does not exist, the redirection command is failed. In this case, the packet classification is performed according to default classification rule or the QoS filters defined in the ingress context of the network QoS policy applied to the network IP interface on which the pseudowire packet is received.

The **no** version of this command removes the redirection of the pseudowire to the queue-group.

Parameters *network-policy-id* — Specifies the network policy identification. The value uniquely identifies the policy on the system.

Values 1 — 65535

fp- redirect-group *queue-group-name* **—** Specifies the name of the queue group template up to 32 characters in length.

ingress-instance *instance-id* **—** Specifies the identification of a specific instance of the queue-group.

Values 1 — 16384

mfib-allowed-mda-destinations

Syntax **mfib-allowed-mda-destinations**

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

Description This command enables the context to configure MFIB-allowed MDA destinations.

The allowed-mda-destinations node and the corresponding **mda** command are used on spoke and mesh SDP bindings to provide a list of MDA destinations in the chassis that are allowed as destinations for multicast streams represented by [*,g] and [s,g] multicast flooding records on the VPLS service. The MDA list only applies to IP multicast forwarding when IGMP snooping is enabled on the VPLS service. The MDA list has no effect on normal VPLS flooding such as broadcast, L2 multicast, unknown destinations or non-snooped IP multicast.

At the IGMP snooping level, a spoke or mesh SDP binding is included in the flooding domain for an IP multicast stream when it has either been defined as a multicast router port, received a IGMP query through the binding or has been associated with the multicast stream through an IGMP request by a host over the binding. Due to the dynamic nature of the way that a spoke or mesh SDP binding is associated with one or more egress network IP interfaces, the system treats the binding as appearing on all network ports. This causes all possible network destinations in the switch fabric to be included in the multicast streams flooding domain. The MDA destination list provides a simple mechanism that narrows the IP multicast switch fabric destinations for the spoke or mesh SDP binding.

If no MDAs are defined within the allowed-mda-destinations node, the system operates normally and will forward IP multicast flooded packets associated with the spoke or mesh SDP binding to all switch fabric taps containing network IP interfaces.

The MDA inclusion list should include all MDAs that the SDP binding may attempt to forward through. A simple way to ensure that an MDA that is not included in the list is not being used by the binding is to define the SDP the binding is associated with as MPLS and use an RSVP-TE LSP with a strict egress hop. The MDA associated with the IP interface defined as the strict egress hop should be

present in the inclusion list. If the inclusion list does not currently contain the MDA that the binding is forwarding through, the multicast packets will not reach the destination represented by the binding.

By default, the MDA inclusion list is empty.

If an MDA is removed from the list, the MDA is automatically removed from the flooding domain of any snooped IP multicast streams associated with a destination on the MDA unless the MDA was the last MDA on the inclusion list. Once the inclusion list is empty, all MDAs are eligible for snooped IP multicast flooding for streams associated with the SDP binding.

mda

Syntax [**no**] **mda** *mda-id*

Context config>service>vpls>mesh-sdp>egress>mfib-allowed-mda-destinations
config>service>vpls>spoke-sdp>egress>mfib-allowed-mda-destinations

Description This command specifies an MFIB-allowed MDA destination for an SDP binding configured in the system.

Parameters *mda-id* — Specifies an MFIB-allowed MDA destination.

Values

slot/mda	slot:	1 — 10
	mda:	1 — 2

vc-label

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

Context config>service>vpls>mesh-sdp>egress
config>service>vpls>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** *vc-label*

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

Description This command configures the ingress VC label.

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

Values 2048 — 18431

static-mac

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

Context 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.

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.

Default none

Parameters *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.

transit-policy

Syntax **transit-policy prefix** *prefix-aasub-policy-id*
no transit-policy

Context config>service>vpls>spoke-sdp

Description This command assigns a transit policy id.

The **no** form of the command removes the transit policy ID from the spoke SDP configuration.

Default no transit-policy

Parameters *prefix-aasub-policy-id* — Specifies the transit policy ID.

Values 1 — 65535

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.

SAP Subscriber Management Commands

cpu-protection

Syntax **cpu-protection** *policy-id* [**mac-monitoring**]
no cpu-protection

Context config>service>vpls>sap
config>template>vpls-sap-template

Description This command assigns an existing CPU protection policy to the associated service SAP. The CPU protection policies are configured in the **config>sys>security>cpu-protection>policy** *cpu-protection-policy-id* context.

If no CPU protection policy is assigned to a service SAP, then a the default policy is used to limit the overall-rate.

Default cpu-protection 254 (for access interfaces)

cpu-protection 255 (for network interfaces)

The configuration of no cpu-protection returns the interface/SAP to the default policies as shown above.

Parameters *policy-id* — Specifies an existing CPU protection policy.

Values 1 — 255

mac-monitoring — When specified, the per MAC rate limiting should be performed, using the per-source-rate from the associated cpu-protection policy.

default-msap-policy

Syntax **default-msap-policy** *policy-name*
no default-msap-policy

Context config>service>vpls>sap

Description This command specifies an existing managed SAP policy. Managed SAPs allow the use of policies and a SAP template for the creation of a SAP. Managed SAP policies are created in the **config>subscr-mgmt** context. This command is only applicable to SAPs created as a capture-sap.

Default none

Parameters *msap-policy-name* — Specifies an existing managed SAP policy name up to 32 characters in length.

sub-sla-mgmt

Syntax [**no**] **sub-sla-mgmt**

Context config>service>vpls>sap

Description This command enables the context to configure subscriber management parameters for this SAP.

Default no sub-sla-mgmt

def-inter-dest-id

Syntax **def-inter-dest-id** {**string** *string* | **use-top-q**}
no def-inter-dest-id

Context config>service>vpls>sap>sub-sla-mgmt

Description This command specifies a default destination string for all subscribers associated with the SAP. The command also accepts the **use-top-q** flag that automatically derives the string based on the top most delineating Dot1Q tag from the SAP's encapsulation.

The **no** form of the command removes the default subscriber identification string from the configuration.

no def-sub-id

Default no def-inter-dest-id

Parameters **use-top-q** — Derives the string based on the top most delineating Dot1Q tag from the SAP's encapsulation.

string *string* — Specifies the subscriber identification applicable for a subscriber host.

def-sla-profile

Syntax **def-sla-profile** *default-sla-profile-name*
no def-sla-profile

Context config>service>vpls>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

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.

mac-da-hashing

Syntax [**no**] **mac-da-hashing**

Context config>service>vpls>sap>sub-sla-mgmt

Description This command specifies whether subscriber traffic egressing a LAG SAP has its egress LAG link selected by a function of the MAC destination address instead of the subscriber ID.

This command is only meaningful if subscriber management is enabled and can be configured for this VPLS service.

multi-sub-sap

Syntax **multi-sub-sap** [*subscriber-limit*]
no multi-sub-sap

Context config>service>vpls>sap>sub-sla-mgmt

Description This command configures the maximum number of subscribers for this SAP.

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

Default 1

Parameters *number-of-sub* — Specifies the maximum number of subscribers for this SAP.

Values 2 — 8000

non-sub-traffic

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

Context config>service>vpls>sap>sub-sla-mgmt>single-sub

Description This command configures non-subscriber traffic profiles. It is used in conjunction with the **profiled-traffic-only** command 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 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 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 service 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>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.

single-sub-parameters

Syntax **single-sub-parameters**

Context config>service>vpls>sap>sub-sla-mgmt

Description This command enables the context to configure single subscriber parameters for this SAP.

sub-ident-policy

Syntax **sub-ident-policy** *sub-ident-policy-name*

Context config>service>vpls>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.

VPLS Multicast Commands

fast-leave

Syntax [**no**] **fast-leave**

Context config>service>vpls>sap>igmp-snooping
config>service>vpls>sap>mld-snooping
config>service>vpls>spoke-sdp>igmp-snooping
config>service>vpls>mesh-sdp>igmp-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 OS router will immediately remove a SAP or SDP from the 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>igmp-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.

Values *service-id*: 1 — 2147483648

group

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

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

Description This command identifies filter policy of multicast groups to be applied to this VPLS entity. The sources of the multicast traffic must be a member of the VPLS.

The **no** form of the command removes the policy association from the VPLS configuration.

Default No group policy is specified.

Parameters *policy-name* — The group policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes. Routing policies are configured in the config>router>policy-options context. The router policy must be defined before it can be imported. For details on IGMP policies, see section "Enabling IGMP group membership report filtering" in the OS Router Configuration Guide.

fault-propagation-bmac

Syntax **fault-propagation-bmac** [*mac-name* | *ieee-address*] [**create**]
no fault-propagation-bmac [*mac-name* | *ieee-address*]

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

Description This command configures associated BMAC addresses for fault propagation on a B-VPLS SAP or SDP binding. The statement can appear up to four times in the configuration to support four remote BMAC addresses in the same remote B-VPLS. The configured VPLS must be a B-VPLS.

The **no** form of the command removes the specified MAC name or MAC address from the list of Fault Propagation BMAC addresses associated with the SAP (or SDP).

Parameters *mac-name* — Specifies a (predefined) MAC name to associate with the SAP or SDP, indirectly specifying a Fault Propagation BMAC address. Up to 32 characters in length.

ieee-address — Specifies a MAC address to associate with the SAP or SDP, directly specifying a Fault Propagation BMAC address. The value should be input in either a xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx format.

force-vlan-vc-forwarding

Syntax [**no**] **force-vlan-vc-forwarding**

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

This command forces vc-vlan-type forwarding in the data path for spoke/mesh SDPs which have either vc-type. This comand is not allowed on vlan-vc-type SDPs.

The **no** form of this command sets default behavior.

Default disabled

hash-label

Syntax **hash-label** [**signal-capability**]
no hash-label

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

Description This command enables the use of the hash label on a VLL, VPRN or VPLS service bound to LDP or RSVP SDP as well as to a VPRN service using the autobind mode with the **ldp**, **rsvp-te**, or **mpls** options. This feature is not supported on a service bound to a GRE SDP or for a VPRN service using the autobind mode with the gre option. This feature is also not supported on multicast packets forwarded using RSVP P2MP LPS or mLDP LSP in both the base router instance and in the multicast VPN (mVPN) instance. It is, however, supported when forwarding multicast packets using an IES/VPRN spoke-interface.

When this feature is enabled, the ingress data path is modified such that the result of the hash on the packet header is communicated to the egress data path for use as the value of the label field of the hash label. The egress data path appends the hash label at the bottom of the stack (BoS) and sets the S-bit to one (1).

In order to allow applications where the egress LER infers the presence of the hash label implicitly from the value of the label, the Most Significant Bit (MSB) of the result of the hash is set before copying into the Hash Label. This means that the value of the hash label will always be in the range [524,288 - 1,048,575] and will not overlap with the signaled/static LSP and signaled/static service label ranges. This also guarantees that the hash label will not match a value in the reserved label range.

The (unmodified) result of the hash continues to be used for the purpose of ECMP and LAG spraying of packets locally on the ingress LER. Note, however, that for VLL services, the result of the hash is overwritten and the ECMP and LAG spraying will be based on service-id when ingress SAP shared queuing is not enabled. However, the hash label will still reflect the result of the hash such that an LSR can use it to perform fine grained load balancing of VLL pseudowire packets.

Packets generated in CPM and that are forwarded labeled within the context of a service (for example, OAM packets) must also include a Hash Label at the BoS and set the S-bit accordingly.

The TTL of the hash label is set to a value of 0.

The user enables the signaling of the hash-label capability under a VLL spoke-sdp, a VPLS spoke-sdp or mesh-sdp, or an IES/VPRN spoke interface by adding the **signal-capability** option. In this case, the decision whether to insert the hash label on the user and control plane packets by the local PE is solely determined by the outcome of the signaling process and can override the local PE configuration. The following are the procedures:

- The 7750 SR local PE will insert the flow label interface parameters sub-TLV with F=1 in the pseudowire ID FEC element in the label mapping message for that spoke-sdp or mesh-sdp.
- If the remote PE includes this sub-TLV with F=1 or F=0, then local PE must insert the hash label in the user and control plane packets.
- If remote PE does not include this sub-TLV (for example, it does not support it, or it is supported but the user did not enable the **hash-label** option or the **signal-capability** option), then the local PE establishes the pseudowire but must not insert the hash label in the user and control packets over that spoke-sdp or mesh-sdp. If the remote PE does not support the **signal-capability** option, then there are a couple of possible outcomes:
  - If the **hash-label** option was enabled on the local configuration of the spoke-sdp or mesh-sdp at the remote PE, the pseudowire packets received by the local PE will have the hash label included. These packets must be dropped. The only way to solve this is to disable the signaling capability option on the local node which will result in the insertion of the hash label by both PE nodes.
  - If the **hash-label** option is not supported or was not enabled on the local configuration of the spoke-sdp or mesh-sdp at the remote PE, the pseudowire received by the local PE will not have the hash label included.
- The user can enable or disable the signal-capability option in CLI as needed. When doing so, the 7750 SR must withdraw the label it sent to its peer and send a new label mapping message with the new value of the F bit in the flow label interface parameters sub-TLV of the pseudowire ID FEC element.

The **no** form of this command disables the use of the hash label.

Default no hash-label

Parameters **signal-capability** — Enables the signaling and negotiation of the use of the hash label between the local and remote PE nodes. The **signal-capability** option is not supported on a VPRN spoke-sdp.

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

igmp-host-tracking

Syntax **igmp-host-tracking**

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

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>vpls>igmp-snooping
config>service>vpls>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>vpls>sap>igmp-snooping

Description This command associates an import policy to filter IGMP packets.

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

Default no import

Parameters *policy-name* — Specifies the import policy name.

max-num-groups

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

Context config>service>vpls>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>vpls>sap>igmp-host-traking
config>service>vpls>mesh-sdp>igmp-snooping
config>service>vpls>sap>igmp-host-tracking
config>service>vpls>sap>igmp-snooping
cconfig>service>vpls>spoke-sdp>igmp-snooping

Description This command configures the maximum number of multicast sources allowed per group.

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

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

Values 1 — 1000

max-num-grp-sources

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

Context config>service>vpls>sap>igmp-host-traking
config>service>vpls>mesh-sdp>igmp-snooping
config>service>vpls>sap>igmp-host-tracking
config>service>vpls>sap>igmp-snooping
cconfig>service>vpls>spoke-sdp>igmp-snooping

Description This command defines the maximum number of multicast (S,G)s that can be joined on this SAP or SDP. If the node receives an IGMP join message that would exceed the configured number of (S,G)s, the request is ignored.

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

Default no max-num-grp-sources

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

import

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

Context config>service>vpls>sap>igmp-snooping
config>service>vpls>spoke-sdp>igmp-snooping
config> service>vpls> mesh-sdp>igmp-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 or SDP 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 import policy name. Values can be string up to 32 characters long of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes. These 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>igmp-snooping
config>service>vpls>spoke-sdp>igmp-snooping
config>service>vpls>mesh-sdp>igmp-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

mcac

Syntax **mcac**

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

mc-constraints

Syntax **mc-constraints**

Context config>service>vpls>sap>igmp-snooping>mcac

Description This command enables the context to configure multicast CAC constraints.

Default none

level

Syntax **level** *level-id* **bw** *bandwidth*
no level *level-id*

Context config>service>vpls>sap>igmp-snooping>mcac>mc-constraints

Description This command configures levels and their associated bandwidth for multicast cac policy on this interface.

Parameters *level-id* — Specifies has an entry for each multicast CAC policy constraint level configured on this system.

Values 1 — 8

bandwidth — Specifies the bandwidth in kilobits per second (kbps) for the level.

Values 1 — 2147483647

number-down

Syntax **number-down** *number-lag-port-down*
no number-down

Context config>service>vpls>sap>igmp-snooping>mcac>mc-constraints

Description This command configure the number of ports down along with level for multicast cac policy on this interface.

Default not enabled

max-num-groups

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

Context config>service>vpls>sap>igmp-snooping
config>service>vpls>spoke-sdp>igmp-snooping
config>service>vpls>mesh-sdp>igmp-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 node receives an IGMP join message that would exceed the configured number of groups, the request is ignored.

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

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

max-num-sources

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

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

Description This command defines the maximum number of multicast sources that can be joined on this SAP or SDP. If the node receives an IGMP join message that would exceed the configured number of sources, the request is ignored.

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

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

Values 1 — 1000

mrouter-port

Syntax [**no**] **mrouter-port**

Context config>service>vpls>sap>igmp-snooping
config>service>vpls>spoke-sdp>igmp-snooping
config>service>vpls>mesh-sdp>igmp-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>igmp-snooping
config>service>vpls>mld-snooping
config>service>vpls>sap>igmp-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>igmp-snooping
config>service>vpls>sap>igmp-snooping
config>service>vpls>spoke-sdp>igmp-snooping
config>service>vpls>mesh-sdp>igmp-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 This command configures the IGMP query interval. If the **send-queries** command is enabled, this parameter specifies the interval between two consecutive general queries sent by the system on 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

Values **config>service>vpls>igmp-snooping**: 1 - 65535
config>service>vpls>sap>igmp-snooping: 2 - 1024

query-src-ip

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

Context config>service>vpls>igmp-snooping

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

query-response-interval

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

Context config>service>vpls>sap>igmp-snooping
config>service>vpls>spoke-sdp>igmp-snooping
config>service>vpls>mesh-sdp>igmp-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 IGMP query response interval. If the **send-queries** command is enabled, this parameter specifies the maximum response time advertised in IGMPv2/v3 queries.

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

If send-queries is not enabled on 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

query-src-ip

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

Context config>service>vpls>mld-snooping

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

report-src-ip

Syntax **report-src-ip**-*address*
no report-src-ip

Context config>service>vpls>igmp-snooping

Description This parameter 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>igmp-snooping
config>service>vpls>sap>igmp-snooping
config>service>vpls>spoke-sdp>igmp-snooping
config>service>vpls>mesh-sdp>igmp-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 **config>service>vpls>sap>igmp-snooping**: 2 — 7
config>service>vpls>igmp-snooping: 1 — 255

mrp

Syntax **mrp**

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

Description This command configures Multiple Registration Protocol (MRP) parameters.

mvrp

Syntax **mvrp**

Context config>service>vpls

Description This command configures MVRP parameters.

attribute-table-size

Syntax [**no**] **attribute-table-size** *value*

Context config>service>vpls>mvrp

Description This command controls the number of attributes accepted on a per BVPLS basis. When the limit is reached, no new attributes will be registered.

If a new lower limit (smaller than the current number of attributes) from a local or dynamic IVPLS is being provisioned, a CLI warning will be issued stating that the system is currently beyond the new limit. The value will be accepted, but any creation of new attributes will be blocked under the attribute count drops below the new limit; the software will then start enforcing the new limit.

Default maximum number of attributes

Parameters *value* — Specifies the number of attributes accepted on a per BVPLS basis.

Values 1 — 4095 for MVRP

attribute-table-size

Syntax [**no**] **attribute-table-size** *value*

Context config>service>vpls>mrp
config>service>vpls>mvrp

Description This command controls the number of attributes accepted on a per BVPLS basis. When the limit is reached, no new attributes will be registered.

If a new lower limit (smaller than the current number of attributes) from a local or dynamic IVPLS is being provisioned, a CLI warning will be issued stating that the system is currently beyond the new limit. The value will be accepted, but any creation of new attributes will be blocked under the attribute count drops below the new limit; the software will then start enforcing the new limit.

Default maximum number of attributes

Parameters *value* — Specifies the number of attributes accepted on a per BVPLS basis.

Values

SR-7/SR-12:	1 — 2047
SR-1	1 — 1023

attribute-table-high-wmark

Syntax [**no**] **attribute-table-high-wmark** *high-water-mark*

Context config>service>vpls>mrp
config>service>vpls>mvrp

Description This command specifies the percentage filling level of the MMRP attribute table where logs and traps are sent.

Default 95%

Parameters *high-water-mark* — Specifies the utilization of the MRP attribute table of this service at which a table full alarm will be raised by the agent.

Values 1% — 100%

attribute-table-low-wmark

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

Context config>service>vpls>mrp
config>service>vpls>mvrp

Description This command specifies the MMRP attribute table low matermark as a percentage. When the percentage filling level of the MMRP attribute table drops below the configured value, the corresponding trap is cleared and/or a log entry is added.

Default 90%

Parameters *low-water-mark* — Specifies utilization of the MRP attribute table of this service at which a table full alarm will be cleared by the agent.

Values 1% — 100%

flood-time

Syntax **flood-time** *flood-time*
no flood-time

Context config>service>vpls>mrp
config>service>vpls>mvrp

Description This command configures the amount of time, in seconds, after a status change in the VPLS service during which traffic is flooded. Once that time expires, traffic will be delivered according to the MMRP registrations that exist in the VPLS.

Default 3 seconds

Parameters *flood-time* — Specifies the MRP flood time, in seconds.

Values 3 — 600

join-time

Syntax [**no**] **join-time** *value*

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

Description This command controls the interval between trasmit opportunities that are applied to the Applicant state machine. An instance of this Join Period Timer is required on a per-Port, per-MRP Participant basis. For additional information, refer to IEEE 802.1ak-2007 section 10.7.4.1.

Default 2

Parameters *value* — Specifies the timer value in 10th of seconds for sending join-messages.

Values 1 — 10 tenths of a second

leave-time

Syntax [**no**] **leave-time** *value*

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

Description This command controls the period of time that the Registrar state machine will wait in the leave state before transitioning to the MT state when it is removed. An instance of the timer is required for each state machine that is in the leave state. The Leave Period Timer is set to the value leave-time when it is started.

A registration is normally in “in” state where there is an MFIB entry and traffic is being forwarded. When a “leave all” is performed (periodically around every 10-15 seconds per SAP/SDP binding -

see leave-all-time-below), a node sends a message to its peer indicating a leave all is occurring and puts all of its registrations in leave state.

The peer refreshes its registrations based on the leave all PDU it receives and sends a PDU back to the originating node with the state of all its declarations.

Refer to IEEE 802.1ak-2007 section 10.7.4.2.

Default 30

Parameters *value* — [30-60] tenths of a second

leave-all-time

Syntax [**no**] **leave-all-time** *value*

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

Description This command controls the frequency with which the LeaveAll state machine generates LeaveAll PDUs. The timer is required on a per-Port, per-MRP Participant basis. The Leave All Period Timer is set to a random value, T, in the range LeaveAllTime<T<1.5*leave-all-time when it is started. Refer to IEEE 802.1ak-2007 section 10.7.4.3.

Default 100

Parameters *value* — [60-300] tenths of a second

mrp-policy

Syntax [**no**] **mrp-policy**-*name*

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

Description This command instructs MMRP to use the mrp-policy specified in the command to control which Group BMAC attributes will be declared and registered on the egress SAP/Mesh-SDP/Spoke-SDP. The Group BMACs will be derived from the ISIDs using the procedure used in the PBB solution. The Group MAC = standard OUI with the last 24 bits being the ISID value. If the policy-name refers to a non-existing mrp-policy the command should return error. Changes to a mrp-policy are aloowed and applied to the SAP/SDPs under which the policy is referenced.

Default no mrp-policy is defined

Parameters *policy-name* — Specifies the redirect 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.

periodic-time

Syntax [**no**] **periodic-time** *value*

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

Description This command controls the frequency the PeriodicTramsmission state machine generates periodic events if the Periodic Transmission Timer is enabled. The timer is required on a per-Port basis. The Periodic Transmissing Timer is set to one second when it is started.

Default 10

Parameters *value* — [10-100] tenths of a second

periodic-timer

Syntax [**no**] **periodic-timer**

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

Description This command enables or disables the Periodic Transmission Timer.

Default disabled

multicast-info-policy

Syntax **multicast-info-policy** *policy-name*
no multicast-info-policy

Context config>service>vpls

Description This command specifies the multicast policy name configured on this service.

per-service-hashing

Syntax [**no**] **per-service-hashing**

Context config>service>vpls
config>template>vpls-template

Description This command enables on a per service basis, consistent per-service hashing for Ethernet services over LAG, over Ethernet tunnel (eth-tunnel) using loadsharing protection-type or over CCAG. Specifically, it enables the new hashing procedures for Epipe, VPLS, regular or PBB services.

The following algorithm describes the hash-key used for hashing when the new option is enabled:

- If the packet is PBB encapsulated (contains an I-TAG ethertype) at the ingress side, use the ISID value from the I-TAG
- If the packet is not PBB encapsulated at the ingress side
  - For regular (non-PBB) VPLS and Epipe services, use the related service ID
  - If the packet is originated from an ingress IVPLS or PBB Epipe SAP
    - •If there is an ISID configured use the related ISID value
    - •If there is no ISID yet configured use the related service ID
  - For BVPLS transit traffic use the related flood list id
    - •Transit traffic is the traffic going between BVPLS endpoints
    - •An example of non-PBB transit traffic in BVPLS is the OAM traffic
- The above rules apply regardless of traffic type
  - Unicast, BUM flooded without MMRP or with MMRP, IGMP snooped

The **no** form of this command implies the use of existing hashing options.

Default no per-service-hashing

pim-snooping

Syntax [**no**] **pim-snooping**

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

Context This command enables PIM snooping for the VPLS service. When enabled, it is enabled for all SAPs except default SAPs. A default SAP is a SAP that has a wildcard VLAN ID, such as sap 1/1/1:*.

The **no** form of the command removes the PIM snooping configuration.

max-num-groups

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

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

Description This command configures the maximum groups for PIM snooping.

Parameters *num-groups* — Specifies the maximum groups for PIM snooping.

Values 1 — 16000

The max number of MFIBs is 1000 for a 7750 router in chassis mode A.

The max number of MFIBs is 4000 for a 7750 router in chassis mode B and C.

oper-group

Syntax [**no**] **oper-group** *name*

Context config>service>vpls>sap
config>service>vpls>spoke-sdp
config>service>vpls>bgp>pw-template-binding

Description This command associates the context to which it is configured to the opeational group specified in the ***nam***e. The oper-group ***name*** must be already configured under config>service before its name is referenced in this command.

The **no** form of the command removes the association.

Default no oper-group

Parameters *name* — A character string of maximum 32 ASCII characters identifying the group instance.

monitor-oper-group

Syntax [**no**] **monitor-oper-group** *name*

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

Description This command specifies the operational group to be monitored by the object under which it is configured. The oper-group ***name*** must be already configured under config>service before its name is referenced in this command.

The **no** form of the command removes the association.

Default no oper-group

Parameters *name* — A character string of maximum 32 ASCII characters identifying the group instance.

hold-time

Syntax **hold-time** *seconds*
no hold-time

Context config>service>vpls>pim-snooping

Description This command configures the duration that allows the PIM-snooping switch to snoop all the PIM states in the VPLS. During this duration, multicast traffic is flooded in the VPLS. At the end of this duration, multicast traffic is forwarded using the snooped states.

When PIM snooping is enabled in VPLS, there is a period of time when the PIM snooping switch may not have built complete snooping state. The switch cannot build states until the routers connected to the VPLS refresh their PIM messages.

This parameter is applicable only if PIM snooping is enabled.

Parameters *seconds* — Specifies the PIM snooping hold time, in seconds

Values 0 — 300

Default 90

mode

Syntax **mode** *mode*

Context config>service>vpls>pim-snooping

Description This command sets the PIM snooping mode to proxy or plain snooping.

Parameters *mode* — Specifies PIM snooping mode.

Values snooping, proxy

Default proxy

precedence

Syntax **precedence** *precedence-value*| **primary**
no precedence

Context config>service>vpls>spoke-sdp

Description This command configures the spoke SDP precendence.

Default 4

Parameters *precedence-value* — Specify the spoke SDP precedence.

Values 0 — 4

primary — Specifies that the prececence is primary.

pw-status-signaling

Syntax [**no**] **pw-status-signaling**

Context config>service>vpls>spoke-sdp

Description This command specifies the type of signaling used by this multi-segment pseudowire provider-edge for this service.

When no pw-status-signaling is enabled, a 7x50 will not include the pseudowire status TLV in the initial label mapping message of the pseudowire used for a spoke SDP. This will force both 7x50 PEs to use the pseudowire label withdrawal method for signaling pseudowire status.

If pw-status-signaling is configured, the node will include the use of the pseudowire status TLV in the initial label mapping message for the pseudowire.

propagate-mac-flush

Syntax [**no**] **propagate-mac-flush**

Context config>service>vpls

Description This command specifies whether MAC flush messages received from the given LDP are propagated to all spoke and mesh SDPs within the context of this VPLS service. The propagation will follow the split-horizon principle and any data-path blocking in order to avoid the looping of these messages.

Default no propagate-mac-flush

send-queries

Syntax [**no**] **send-queries**

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

When send-queries is configured, all type of queries generate ourselves are of the configured version. If a report of a version higher than the configured version is received, the report will get dropped and a new wrong version counter will get incremented. If send-queries is not configured, the version command has no effect. The version used will be the version of the querier. This implies that, for example, when we have a v2 querier, we will never send out a v3 group or group-source specific query when a host wants to leave a certain group.

Default no send-queries

source

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

Context config>service>vpls>sap>igmp-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 a multicast group from a specified source. For a multicast group, more than one source address can be specified. Static (s,g) entries cannot be added, if a starg is previously created.

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

Default none

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

starg

Syntax [**no**] **starg**

Context config>service>vpls>sap>igmp-snooping>static>group
config>service>vpls>spoke-sdp>igmp-snooping>static>group
config>service>vpls>mesh-sdp>igmp-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.

The **no** form of the command removes the starg entry from the configuration.

Default no starg

static

Syntax **static**

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

version

Syntax **version** *version*
no version

Context config>service>vpls>sap>igmp-snooping
config>service>vpls>mesh-sdp>igmp-snooping
config>service>vpls>spoke-sdp>igmp-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 version of IGMP which is running on this SAP or SDP. This object can be used to configure a router capable of running either value. For IGMP to function correctly, all routers on a LAN must be configured to run the same version of IGMP on that LAN.

When the **send-query** command is configured, all type of queries generate ourselves are of the configured **version**. If a report of a version higher than the configured version is received, the report gets dropped and a new "wrong version" counter is incremented.

If the **send-query** command is not configured, the **version** command has no effect. The version used on that SAP or SDP will be the version of the querier. This implies that, for example, when there is a v2 querier, a v3 group or group-source specific query when a host wants to leave a certain group will never be sent.

Parameters *version* — Specify the IGMP version.

Values 1, 2, 3

to-sap

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

Context config>service>vpls>sap>igmp-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 CLI Command Descriptions on page 2569 for command syntax.

VPLS DHCP and Anti-Spoofing Commands

anti-spoof

Syntax **anti-spoof** {**ip** | **mac** | **ip-mac**}
no anti-spoof

Context config>service>vpls>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.

app-profile

Syntax **app-profile** *app-profile-name*
no app-profile

Context config>service>vpls>sap

Description This command configures the application profile name.

Parameters *app-profile-name* — Specifies an existing application profile name configured in the **config>app-assure>group>policy** context.

arp-host

Syntax **arp-host**

Context config>service>vpls>sap

Description This command enables the context to configure ARP host parameters.

host-limit

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

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

Description This command configures the maximum number of ARP hosts.

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

Default 1

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

Values 1 — 32767

min-auth-interval

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

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

Description This command configures the minimum authentication interval.

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

Default 15

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

force-l2pt-boundary

Syntax [**no**] **force-l2pt-boundary**

Context config>service>vpls>sap

Description Enabling force-l2pt-boundary will force that all SAPs managed by the given m-vpls instance on the corresponding port will have to have l2pt-termination enabled. This command is applicable only to SAPs created under m-vpls and this regardless the flavor of STP currently being active. It is not applicable to spoke SDPS.

The execution of this command will fail as soon as at least one of the currently managed SAPs (all SAPs falling within the specified managed-vlan-range) does not have l2pt-termination enabled, and this regardless its admin/operational status.

If force-l2pt-boundary is enabled on a given m-vpls SAP, all newly created SAPs falling into the specified managed-vlan-range will have l2pt-termination enabled per default.

Extending or adding new range into a managed-vlan-range declaration will fail as soon as there is at least one SAPs falling into the specified vlan-range does not have l2pt-termination enabled.

Disabling l2pt-termination on currently managed SAPs will fail as soon as the force-l2pt-boundary is enabled under corresponding m-vpls SAP.

frame-relay

Syntax **frame-relay**

Context config>service>vpls>sap

Description This command enables the context to configure frame-relay parameters.

frf-12

Syntax [**no**] **frf-12**

Context config>service>vpls>sap>fr

Description This command enables FRF12 headers. This must be set to disabled for this entry to be added to an MLFR bundle.

The **no** form of the command disables FRF12 headers.

ete-fragment-threshold

Syntax **ete-fragment-threshold** *threshold*
no ete-fragment-threshold

Context config>service>vpls>sap>fr>frf-12

Description This command configures the FRF.12 fragmentation threshold.

The **no** form of the command rremoves the value.

Default 128

Parameters *threshold* — Specifies the maximum length of a fragment to be transmitted.

Values 128 — 512

interleave

Syntax **interleave**
no interleave

Context config>service>vpls>sap>frame-relay>frf.12

Description This command enables interleaving of high priority frames and low-priority frame fragments within a FR SAP using FRF.12 end-to-end fragmentation.

When this option is enabled, only frames of the FR SAP non expedited forwarding class queues are subject to fragmentation. The frames of the FR SAP expedited queues are interleaved, with no fragmentation header, among the fragmented frames. In effect, this provides a behavior like in MLPPP Link Fragment Interleaving (LFI).

When this option is disabled, frames of all the FR SAP forwarding class queues are subject to fragmentation. The fragmentation header is however not included when the frame size is smaller than the user configured fragmentation size. In this mode, the SAP transmits all fragments of a frame before sending the next full or fragmented frame.

The receive direction of the FR SAP supports both modes of operation concurrently, with and without fragment interleaving.

The **no** form of this command restores the default mode of operation.

Default no interleave

scheduling-class

Syntax **scheduling-class** *class-id*
no scheduling-class

Context config>service>vpls>sap>frame-relay

Description This command specifies the scheduling class to use for this SAP. This object is only applicable for a SAP whose bundle type is set to MLFR.

Parameters *class-id* — Specifies the scheduling class.

Values 0 — 3

host-connectivity-verify

Syntax **host-connectivity-verify source-ip** *ip-address* [**source-mac** *ieee-address*] [**interval** *interval*] [**action** {**remove** | **alarm**}]

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

Description This command enables enables subscriber host connectivity verification on a given SAP within a VPLS service. This tool will periodically scan all known hosts (from dhcp-state) and perform a UC ARP request. The subscriber host connectivity verification will maintain state (connected vs. not-connected) for all hosts.

Default no host-connectivity-verify

Parameters **source-ip** *ip-address* — Specify an unused IP address in the same network for generation of subscriber host connectivity verification packets.

source-mac *ieee-address* — Specifies the source MAC address to be used for generation of subscriber host connectivity verification packets.

interval ***interval*** — The interval, in minutes, which specifies the time interval in which all known sources should be verified. The actual rate is then dependent on number of known hosts and interval.

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

action {**remove** | **alarm**} — Defines the action taken on a subscriber host connectivity verification failure for a given host. The **remove** keyword raises an alarm and removes dhcp-state and releases all allocated resources (queues, table entries, etc.). DHCP release will be signaled to corresponding DHCP server. Static host will be never removed. The **alarm** keyword raises an alarm indicating that the host is disconnected.

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 spaces, the string must be specified with beginning and ending 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.

Default 16

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 list's 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 1 — 30

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 removes the egress filter from each member SAP. The **no egress-filter** command specifies that an egress filter (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 any time.

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.

qinq-etype

Syntax **qinq-etype** [**0x0600**..**0xffff**]
no qinq-etype

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

Description This command specifies the Ethertype used for QinQ encapsulation.

Default no qinq-etype

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

qinq-fixed-tag-value

Syntax **qinq-fixed-tag-value** *tag-value*
no qinq-fixed-tag-value

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

Description This command configures the fixed tag value used for QinQ encapsulation.

Default no qinq-fixed-tag-value

Parameters *tag-value* — Specifies the provisioned common value of the fixed 802.1Q tag of all the QinQ SAP's in this egress multicast group.

The value 0 is used to indicate that the actual value of the fixed tag will be defined implicitly by the corresponding tag of the first SAP added to this egress multicast group.

Values 0, 1 — 4094

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 any time.

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

Provider Edge Discovery Policy Commands

pe-discovery-policy

Syntax [**no**] **pe-discovery-policy** *name*

Context config>service

Description This command configures a provider edge discovery policy and parameters.

The **no** form of the command removes the policy from the configureation.

Parameters *name* — Specifies the RADIUS PE discovery policy name, up to 32 characters in length.

password

Syntax **password** *password*
no password

Context config>service>pe-discovery-policy

Description This command configures the PE discovery password that is used when contacting the RADIUS server for VPLS auto-discovery.

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

Default no password

Parameters *password* — Specifies the password, up to 32 characters in length, used when contacting the RADIUS server for VPLS auto-discovery.

polling-interval

Syntax **polling-interval** *minutes*
no polling-interval

Context config>service>pe-discovery-policy

Description This command configures the PE discovery polling interval.

Default 5

Parameters *minutes* — Specifies the polling interval, in minutes, for RADIUS PE discovery.

Values 1 — 30

server

Syntax **server** *server-index* **address** *ip-address* **secret** *key* [**hash** | **hash2**] [**port** *port-num*]
no server *server-index*

Context config>service>pe-discovery-policy

Description This command adds a RADIUS server.

Parameters *server-index* — The index for the RADIUS server. The index determines the sequence in which the servers are queried for authentication requests. Servers are queried in order from lowest to highest index.

Values 1 — 5

address *ip-address* **—** The IP address of the RADIUS server. Two RADIUS servers cannot have the same IP address. An error message is generated if the server address is a duplicate.

secret *key* **—** The secret key to access the RADIUS server. This secret key must match the password on the RADIUS server.

Values Up to 20 characters in length.

hash — Specifies the key is entered in an encrypted form. If the **hash** parameter is not used, the key is assumed to be in a non-encrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** parameter specified.

hash2 — Specifies the key is entered in a more complex encrypted form. If the **hash2** parameter is not used, the less encrypted **hash** form is assumed.

port — Specifies the UDP port number on which to contact the RADIUS server for authentication.

Values 1 — 65535

timeout

Syntax **timeout** *seconds*
no timeout

Context config>service>pe-discovery-policy

Description This command specifies the number of seconds to wait before timing out a RADIUS server.

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

Default 3 seconds

Parameters *seconds* — The number of seconds to wait for a response from a RADIUS server, expressed as a decimal integer.

Values 1 — 90

radius-discovery

Syntax [**no**] **radius-discovery**

Context config>service>vpls

Description This command enables the RADIUS provider edge discovery for this VPLS service.

Default none

pe-discovery-policy

Syntax **pe-discovery-policy** *name*
no pe-discovery-policy

Context config>service>vpls>radius-discovery

Description This command specifies the existing RADIUS PE discovery policy name. The policy must have been configured in the **config>service** context.

Parameters *name* — Specifies the RADIUS PE discovery policy name, up to 32 characters in length.

user-name-format

Syntax **user-name-format** {**vpn-id** *vpn-id* | **router-distinguisher** *rd*}
no pe-discovery-policy

Context config>service>vpls>radius-discovery

Description This command specifies whether the RADIUS user name is a VPN ID or router-distinguisher.

Parameters **vpn-id** *vpn-id* — Indicates the VPN ID of the associated VPLS service.

router-distinguisher *rd* — Sets the identifier attached to routes that distinguishes the VPN it belongs to.

BGP Auto-Discovery Commands

bgp

Syntax **bgp**

Context config>service>vpls

Description This command enables the context to configure the BGP related parameters for both BGP AD and BGP VPLS.

bgp-vpls

Syntax **bgp-vpls**

Context config>service>vpls

Description This command enables the context to configure the BGP-VPLS parameters and addressing.

max-ve-id

Syntax **max-ve-id** *value*
no max-ve-id

Context config>service>vpls>bgp-vpls

Description This command configures the allowed range for the VE-id value: locally configured and received in a NLRI. Configuration of a VE-id higher than the value specified in this command is not allowed.

Also upon reception of a higher VE-id in an NLRI imported in this VPLS instance (RT = configured import RT) the following action must be taken:

- a trap must be generated informing the operator of the mismatch.
- NLRI must be dropped
- no service labels are to be installed for this VE-id
- no new NLRI must be generated if a new offset is required for VE-id.

The **no** form of this command sets the max-ve-id to un-configured. The BGP VPLS status should be administratively down for "no max-ve-id" to be used.

The max-ve-id value can be changed without shutting down bgp-vpls if the newly provisioned value does not conflict with the already configured local VE-ID. If the value of the local-VE-ID is higher than the new max-ve-id value the command is rejected. The operator needs to decrease first the VE-ID before running the command.

The actions taken for other max-ve-id values are described below:

- max-ve-id value higher than all VE-IDs (local and received) is allowed and there are no effects.

- max-ve-id higher than the local VE-ID but smaller than the remote VE-IDs:
  - Provisioning is allowed
  - A warning message will be generated stating that "Higher VE-ID values were received in the BGP VPLS context. Related pseudowires will be removed."
  - The pseudowires associated with the higher VE-IDs will be removed locally.
  - Note that this is a situation that should be corrected by the operator as the pseudowire may be down just at the local PE, consuming unnecessarily core bandwidth. The higher VE-IDs should be removed or lowered.

If the max-ve-id has increased a BGP route refresh is sent to the VPLS community to get the routes which might have been rejected earlier due to max-ve-id check. Default no max-ve-id – max-ve-id is not configured. A max-ve-id value needs to be provisioned for BGP VPLS to be in "no shutdown" state.

Default no max-ve-id

Parameters *value* — Specifies the allowed range of [1-value] for the VE-id. The configured value must be bigger than the existing VE-ids.

Values 1-65535

ve-name

Syntax **ve-name** *name*
no ve-name

Context config>service>vpls>bgp-vpls

Description This command creates or edits a ve-name. Just one ve-name can be created per BGP VPLS instance.

The **no** form of the command removes the configured ve-name from the bgp vpls node. It can be used only when the BGP VPLS status is shutdown. Command "no shutdown" cannot be used if there is no ve-name configured.

Default no ve-name

Parameters *name* — A character string identifying the VPLS Edge instance.

Values 32 ASCII chars max

ve-id

Syntax **ve-id** *ve-id-value*
no ve-id

Context config>service>vpls>bgp-vpls>ve-name

Description This command configures a ve-id. Just one ve-id can be configured per BGP VPLS instance. The VE-ID can be changed without shutting down the VPLS Instance. When the VE-ID changes, BGP is withdrawing its own previously advertised routes and sending a route-refresh to all the peers which

would result in reception of all the remote routes again. The old pseudowires are removed and new ones are instantiated for the new VE-ID value.

The **no** form of the command removes the configured ve-id. It can be used just when the BGP VPLS status is shutdown. Command "no shutdown" cannot be used if there is no ve-id configured.

Default no ve-id

Parameters *value* — A two bytes identifier that represents the local instance in a VPLS and is advertised through the BGP NLRI. Must be lower or equal with the max-ve-id.

Values 1-65535

shutdown

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

Context config>service>vpls>bgp-vpls

Description This command administratively enables/disables the local BGP VPLS instance. On de-activation an MP-UNREACH-NLRI must be sent for the local NLRI.

The **no** form of the command enables the BGP VPLS addressing and the related BGP advertisement. The associated BGP VPLS MP-REACH-NLRI will be advertised in an update message and the corresponding received NLRIs must be considered to instantiate the data plane. RT, RD usage: same like in the BGP AD solution, if the values are not configured here, the value of the VPLS-id from under the bgp-ad node is used. If VPLS-id value is not configured either the MH site cannot be activated – i.e. no shutdown returns an error. Same applies if a pseudowire template is not specified under the bgp node.

Default shutdown

bgp-ad

Syntax [**no**] **bgp-ad**

Context config>service>vpls

Description This command configures BGP auto-discovery.

pw-template-binding

Syntax **pw-template-binding** *policy-id* [**split-horizon-group** *group-name*] [**import-rt** {*ext-community*, ...(up to 5 max)}]
no pw-template-bind *policy-id*

Context config>service>vpls>bgp-ad
config>service>vpls>bgp

Description This command binds the advertisements received with the route target (RT) that matches the configured list (either the generic or the specified import) to a specific pw-template. If the RT list is not present the pw-template is used for all of them.

The pw-template-binding applies to both BGP-AD and BGP-VPLS if these features are enabled in the VPLS.

For BGP VPLS the following additional rules govern the use of pseudowire-template:

- On transmission the settings for the L2-Info extended community in the BGP Update are derived from the pseudowire template attributes. If multiple pseudowire templates (with or without import-rt) are specified for the same VPLS instance the first pw-template entry will be used.
- On reception the values of the parameters in the L2-Info extended community of the BGP Update are compared with the settings from the corresponding pw-template. The following steps are used to determine the local pw-template:
  - The RT values are matched to determine the pw-template.
  - If multiple pw-templates matches are found from the previous steps, the first configured pw-template entry will be considered.
  - If the values used for Layer 2 MTU or C Flag do not match the pseudowire setup fails.

The tools perform commands can be used to control the application of changes in pw-template for both BGP-AD and BGP-VPLS.

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

Default none

Parameters *policy-id* — Specifies an existing policy ID.

Values 1 — 2147483647

split-horizon-group *group-name* **—** The specified group-name overrides the split horizon group template settings.

import-rt *ext-comm* — Specify communities allowed to be accepted from remote PE neighbors. An extended BGP community in the **type**:**x**:**y** format. The value **x** can be an integer or IP address. The **type** can be the target or origin. **x** and **y** are 16-bit integers.

Values target: {*ip-addr:comm-val*| *2byte-asnumber:ext-comm-val*|*4byte-asnumber:comm-val*}

ip-addr	a.b.c.d
comm-val	0 — 65535
2byte-asnumber	0 — 65535
ext-comm-val	0 — 4294967295
4byte-asnumber	0 — 4294967295

oper-group

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

Context config>service>vpls>sap
config>service>vpls>spoke-sdp
config>service>vpls>bgp>pw-template-binding

Description This command associates the context to which it is configured to the operational group specified in the *group-name*. The **oper-group** *group-name* must be already configured under **config>service** context before its name is referenced in this command.

The **no** form of the command removes the association.

Parameters *group-name* — Specifies a character string of maximum 32 ASCII characters identifying the group instance.

route-target

Syntax **route-target** {*ext-community*|{[**export** *ext-community*][**import** *ext-community*]}}
no route-target

Context config>service>vpls>bgp-ad
config>service>vpls>bgp

Description This command configures the route target (RT) component that will be signaled in the related MP-BGP attribute to be used for BGP auto-discovery, BGP VPLS and BGP Multi-Homing if these features are configured in this VPLS service.

If this command is not used, the RT is built automatically using the VPLS ID. The ext-comm can have the same two formats as the VPLS ID, a two-octet AS-specific extended community, IPv4 specific extended community.

The following rules apply:

- if BGP AD VPLS-id is configured & no RT is configured under BGP node - RT = VPLS-ID
- if BGP AD VPLS-id is not configured then an RT value must be configured under BGP node (this is the case when only BGP VPLS is configured)
- if BGP AD VPLS-id is configured and an RT value is also configured under BGP node, the configured RT value prevails

Parameters **export** *ext-community* — Specify communities allowed to be sent to remote PE neighbors.

import *ext-community* — Specify communities allowed to be accepted from remote PE neighbors.

vpls-id

Syntax **vpls-id** *vpls-id*

Context config>service>vpls>bgp-ad

Description This command configures the VPLS ID component that will be signaled in one of the extended community attributes (*ext-comm*).

Values and format (6 bytes, other 2 bytes of type-subtype will be automatically generated)

Parameters *vpls-id* — Specifies a globally unique VPLS ID for BGP auto-discovery in this VPLS service.

Values

vpls-id :	<ip-addr:comm-val>|<as-number:ext-comm-val>	
ip-addr	a.b.c.d	
comm-val	0 — 65535	

as-number	1..65535
ext-comm-val	0..4294967295

vsi-export

Syntax **vsi-export** *policy-name* [*policy-name*...(up to 5 max)]
no vsi-export

Context config>service>vpls>bgp-ad
config>service>vpls>bgp

Description This command specifies the name of the VSI export policies to be used for BGP auto-discovery, BGP VPLS and BGP Multi-Homing if these features are configured in this VPLS service. If multiple policy names are configured, the policies are evaluated in the order they are specified. The first policy that matches is applied.

The policy name list is handled by the SNMP agent as a single entity.

vsi-id

Syntax **vsi-id**

Context config>service>vpls>bgp-ad

Description This command enables the context to configure the Virtual Switch Instance Identifier (VSI-ID).

prefix

Syntax **prefix** *low-order-vsi-id*
no prefix

Context config>service>vpls>bgp-ad>vsi-id

Description This command specifies the low-order 4 bytes used to compose the Virtual Switch Instance Identifier (VSI-ID) to use for NLRI in BGP auto-discovery in this VPLS service.

If no value is set, the system IP address will be used.

Default no prefix

Parameters *low-order-vsi-id* — Specifies a unique VSI ID.

Values 0— 4294967295

route-distinguisher

Syntax **route-distinguisher** [*ip-addr:comm-val* | *as-number:ext-comm-val*]
no route-distinguisher

Context config>service>vpls>bgp-ad>vsi-id
config>service>vpls>bgp

Description This command configures the Route Distinguisher (RD) component that will be signaled in the MP-BGP NLRI for L2VPN AFI. This value will be used for BGP-AD, BGP VPLS and BGP Multi-Homing NLRI if these features are configured.

If this command is not configured, the RD is automatically built using the BGP-AD VPLS ID. The following rules apply:

- if BGP AD VPLS-id is configured & no RD is configured under BGP node - RD = VPLS-ID
- if BGP AD VPLS-id is not configured then an RD value must be configured under BGP node (this is the case when only BGP VPLS is configured)
- if BGP AD VPLS-id is configured and an RD value is also configured under BGP node, the configured RD value prevails

Values and format (6 bytes, other 2 bytes of type will be automatically generated)

Parameters *ip-addr:comm-val* — Specifies the IP address.

Values

ip-addr	a.b.c.d
comm-val	0 — 65535

as-number:ext-comm-val — Specifies the AS number and the

Values

as-number	1 — 65535
ext-comm-val	0 — 4294967295

vsi-import

Syntax **vsi-import** *policy-name* [*policy-name*...(up to 5 max)]
no vsi-import

Context config>service>vpls>bgp-ad>vsi-id
config>service>vpls>bgp

Description This command specifies the name of the VSI import policies to be used for BGP auto-discovery, BGP VPLS and BGP Multi-Homing if these features are configured in this VPLS service. If multiple policy names are configured, the policies are evaluated in the order they are specified. The first policy that matches is applied.

The policy name list is handled by the SNMP agent as a single entity.