VLL Service Configuration Commands

Generic Commands

shutdown

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

Context config>service>apipe
config>service>apipe>sap
config>service>apipe>spoke-sdp
config>service>cpipe
config>service>cpipe>sap
config>service>cpipe>spoke-sdp
config>service>epipe
config>service>epipe>bgp-vpws
config>service>epipe>sap
config>service>epipe>spoke-sdp
config>service>epipe>sap>eth-cfm>mep
config>service>epipe>spoke-sdp>eth-cfm>mep
config>service>fpipe
config>service>fpipe>sap
config>service>fpipe>spoke-sdp
config>service>ipipe
config>service>ipipe>sap
config>service>ipipe>spoke-sdp

Description This command administratively disables an entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics.

The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they may be deleted.

Services are created in the administratively down (**shutdown**) state. When a **no shutdown** command is entered, the service becomes administratively up and then tries to enter the operationally up state. Default administrative states for services and service entities is described below in Special Cases.

The **no** form of this command places the entity into an administratively enabled state.

Special Cases

Service Admin State — Bindings to an SDP within the service will be put into the out-of-service state when the service is shutdown. While the service is shutdown, all customer packets are dropped and counted as discards for billing and debugging purposes.

Service Operational State — A service is regarded as operational providing that at least one SAP and one SDP are operational or if two SAP's 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.

description

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

Context
config>service>apipe
config>service>apipe>sap
config>service>apipe>endpoint
config>service>cpipe
config>service>cpipe>endpoint
config>service>cpipe>sap
config>service>epipe
config>service>epipe>sap
config>service>epipe>spoke-sdp
config>service>epipe>endpoint
config>service>fpipe
config>service>fpipe>sap
config>service>fpipe>endpoint
config>service>ipipe
config>service>ipipe>sap
config>service>ipipe>endpoint

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.

Service Commands

apipe

Syntax **apipe** *service-id* [**customer** *customer-id*] [**vpn** *vpn-id*] [**vc-type** {**atm-vcc** | **atm-sdu** | **atm-vpc** | **atm-cell**}] [**vc-switching**]
no apipe *service-id*

Context config>service

Description The Apipe service provides a point-to-point Layer 2 VPN connection to a remote SAP or to another local SAP. An Apipe can connect an ATM or Frame Relay endpoint either locally or over a PSN to a remote endpoint of the same type or of a different type and perform interworking between the two access technologies.

Parameters *service-id* — The unique service identification number or string identifying the service in the service domain. This ID must be unique to this service and may not be used for any other service of any type. The *service-id* must be the same number used for every 7750 SR, 7450 ESS and 7710 SR 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)

vc-type — Keyword that specifies a 15 bit value that defines the type of the VC signaled to the peer. Its values are defined in *draft-ietf-pwe3-iana-allocation* and it defines both the signaled VC type as well as the resulting datapath encapsulation over the Apipe.

Values atm-vcc, atm-sdu, atm-vpc, atm-cell

Default atm-sdu

vc-switching — Specifies if the pseudowire switching signalling is used for the spoke SDPs configured in this service.

cpipe

Syntax **cpipe** *service-id* [**customer** *customer-id*] [**vpn** *vpn-id*] [**vc-type** {**satop-e1** | **satop-t1** | **satop-e3** | **satop-t3** [**vc-switching**] | **cesopsn** | **cesopsn-cas**}] [**vc-switching**] [**create**]

no cpipe *service-id*

Context config>service

Description This command configures a Circuit Emulation Services instance. 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.

By default, no services exist until they are explicitly created with this command.

The **no** form of this command deletes the service instance with the specified *service-id*. The service cannot be deleted until the service has been shutdown.

Parameters *service-id* — The unique service identification number or string 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: Specifies an existing service name up to 64 characters in length.

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

Values 1 — 2147483647

vpn *vpn-id* — Specifies the VPN ID number which allows you to identify virtual private networks (VPNs) by a VPN ID. If this parameter is not specified, the VPN ID uses the same service ID number.

Values 1 — 2147483647

Default **null** (0)

vc-type — The vc-type defines the type of unstructured or structured circuit emulation service to be configured.

Values **satop-e1**: unstructured E1 circuit emulation service
satop-t1: unstructured DS1 circuit emulation service
satop-e3: unstructured E3 circuit emulation service
satop-t3: unstructured DS3 circuit emulation service
cesopsn: basic structured n*64 kbps circuit emulation service
cesopsn-cas: structured n*64 kbps circuit emulation service with signaling

vc-switching — Specifies if the pseudowire switching signalling is used for the spoke SDPs configured in this service.

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

epipe

Syntax **epipe** *service-id* **customer** *customer-id* [**vpn** *vpn-id*] [**vc-switching**] [**create**]
epipe *service-id*
no epipe *service-id*

Context config>service

This command configures an Epipe service instance. This command is used to configure a point-to-point epipe service. An Epipe connects two endpoints defined as Service Access Points (SAPs). Both SAPs may be defined in one 7750 SR or they may be defined in separate devices connected over the service provider network. When the endpoint SAPs are separated by the service provider network, the far end SAP is generalized into a Service Distribution Point (SDP). This SDP describes a destination and the encapsulation method used to reach it.

No MAC learning or filtering is provided on an Epipe.

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.

By default, no epipe services exist until they are explicitly created with this command.

The **no** form of this command deletes the epipe service instance with the specified *service-id*. The service cannot be deleted until the service has been shutdown.

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

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

customer *customer-id* **—** Specifies 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 ID. If this parameter is not specified, the VPN ID uses the same service ID number.

Values 1 — 2147483647

Default **null** (0)

vc-switching — Specifies if the pseudowire switching signalling is used for the spoke SDPs configured in this service.

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

fpipe

Syntax **fpipe** *service-id* [**customer** *customer-id*] [**vpn** *vpn-id*] [**vc-type** {**fr-dlci**}] [**vc-switching**]
no fpipe *service-id*

Context config>service

Description This command configures an Fpipe service. An Fpipe provides a point-to-point L2 VPN connection to a remote SAP or to another local SAP. An Fpipe connects only Frame Relay endpoints either locally or over a PSN to a remote endpoint of the same type.

Parameters *service-id* — The unique service identification number or string identifying the service in the service domain. This ID must be unique to this service and may not be used for any other service of any type. The *service-id* must be the same number used for every 7750 SR, 7450 ESS and 7710 SR 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)

vc-type — Specifies a 15 bit value that defines the type of the VC signaled to the peer. Its values are defined in *draft-ietf-pwe3-iana-allocation* and it defines both the signaled VC type as well as the resulting datapath encapsulation over the apipe.

Values fr-dlci

vc-switching — Specifies if the pseudowire switching signalling is used for the spoke SDPs configured in this service.

ipipe

Syntax **ipipe** *service-id* [**customer** *customer-id*] [**create**] [**vpn** *vpn-id*] [**vc-switching**]
no ipipe *service-id*

Context config>service

Description This command configures an IP-Pipe service.

Parameters *service-id* — The unique service identification number or string identifying the service in the service domain. This ID must be unique to this service and may not be used for any other service of any type. The *service-id* must be the same number used for every 7750 SR, 7450 ESS and 7710 SR 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)

vc-switching — Specifies if the pseudowire switching signalling is used for the spoke SDPs configured in this service.

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

VLL Global Commands

bgp

Syntax **bgp**

Context config>service>epipe

Description This command enables the context to configure the BGP related parameters BGP used for Multi-Homing and BGP VPWS.

The **no** form of this command removes the string from the configuration.

pw-template-binding

Syntax **pw-template-binding** *policy-id* [**import-rt** {*ext-community*,.(upto 5 max)}]
no pw-template-binding *policy-id*

Context config>service>epipe>bgp

Description This command binds the advertisements received with the route targets (RT) that match the configured list (either the generic or the specified import) to a specific pw-template. If the RT list is not present, or if multiple matches are found, the numerically lowest pw-template is used.

The pw-template-binding applies to BGP-VPWS when enabled in the Epipe.

For p the following additional rules govern the use of pseudowire-template:

- On transmission, the settings for the L2-Info extended community in the BGP updates are derived from the pseudowire template attributes. If multiple pseudowire template bindings (with or without import-rt) are specified for the same VPWS 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 updates are compared with the settings from the corresponding pseudowire template bindings. The following steps are used to determine the local pw-template:
  - The RT values are matched to determine the pw-template.
  - If multiple pw-template-binding matches are found from the previous step, the first (numerically lowest) configured pw-template entry will be considered.
  - If the value used for Layer 2 MTU (unless the value zero is received) does not match the pseudowire is created but with the oper state down.
  - If the values used for the C (control word) or S (sequenced delivery) flags are not zero the pseudowire is not created.

The **tools perform** commands can be used to control the application of changes in pw-template for BGP-VPWS.

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

Parameters *policy-id* — Specifies an existing policy ID.

Values 1 — 2147483647

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.

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

route-distinguisher

Syntax **route-distinguisher** [**ip-addr:***comm-val* | **as-number:***ext-comm-val*]
no route-distinguisher

Context config>service>epipe>bgp

Description This command configures the Route Distinguisher (RD) component thatis signaled in the MPBGP NLRI for L2VPN AFI. This value is used for BGP Multi-Homing and BGP-VPWS.

An RD value must be configured under BGP node.

Format: Six 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:	

as-number:ext-comm-val — Specifies the AS number.

Values

as-number	1 — 65535
ext-comm-val	0 — 4294967295

route-target

Syntax **route-target** {*ext-community*|{[**export** *ext-community*][**import** *ext-community*]}}
no route-target

Context config>service>epipe>bgp

Description This command configures the route target (RT) component that is signaled in the related MPBGP attribute to be used for BGP Multi-Homing and BGP-VPWS when configured in the Epipe service. The ext-comm can have two formats:

- A two-octet AS-specific extended community, IPv4 specific extended community.
- An RT value must be configured under BGP node when BGP Epipe is configured.

Parameters *export ext-community* — Specifies communities allowed to be sent to remote PE neighbors.

import ext-community — Specifies communities allowed to be accepted from remote PE neighbors.

bgp-vpws

Syntax [**no**] **bgp-vpws**

Context config>service>epipe

Description This command enables the context to configure BGP-VPWS parameters and addressing.

Default no bgp-vpws

remote-ve-name

Syntax [**no**] **remote-ve-name** *name*

Context config>service>epipe>bgp-vpws

Description This command creates or edits a remote-ve-name. A single remote-ve-name can be created per BGP VPWS instance if the service is single-homed or uses a single pseudowire to connect to a pair of dual-homed systems. When the service requires active/standby pseudowires to be created to remote dual-homed systems then two remote-ve-names must be configured.

This context defines the remote PE to which a pseudowire will be signaled.
remote-ve-name commands can be added even if bgp-vpws is not shutdown.

The **no** form of the command removes the configured remote-ve-name from the bgp vpws node. It can be used when the BGP VPWS status is either shutdown or "no shutdown".

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

ve-id

Syntax **ve-id** *value*
no ve-id

Context config>service>epipe>bgp-vpws>ve-name
config>service>epipe>bgp-vpws>remote-ve-name

Description This command configures a ve-id for either the local VPWS instance when configured under the ve-name, or for the remote VPWS instance when configured under the remote-ve-name.

A single ve-id can be configured per ve-name or remote-ve-name. The ve-id can be changed without shutting down the VPWS instance. When the ve-name ve-id changes, BGP withdraws the previously advertised route and sends a route-refresh to all the peers which would result in reception of all the remote routes again. The old PWs are removed and new ones are instantiated for the new ve-id value.

When the remote-ve-name ve-id changes, BGP withdraws the previously advertised route and send a new update matching the new ve-id. The old pseudowires are removed and new ones are instantiated for the new ve-id value.

NLRIs received whose advertised ve-id does not match the list of ve-ids configured under the remote ve-id will not have a spoke-SDP binding auto-created but will remain in the BGP routing table but not in the L2 route table. A change in the locally configured ve-ids may result in auto-sdp-bindings either being deleted or created, based on the new matching results.

Each ve-id configured within a service must be unique.

The **no** form of the command removes the configured ve-id. It can be used just when the BGP VPWS 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 or remote VPWS instance and is advertised through the BGP NLRI.

Values 1 — 65535

ve-name

[**no**] **ve-name** *name*

Context config>service>epipe>bgp-vpws

Description This command configures the name of the local VPWS instance in this service.

The **no** form of the command removes the ve-name.

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

shutdown

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

Context config>service>epipe>bgp-vpws

Description This command administratively enables/disables the local BGP VPWS instance. On de-activation an MP-UNREACH-NLRI is sent for the local NLRI.

The **no** form of the command enables the BGP VPWS addressing and the related BGP advertisement. The associated BGP VPWS MP-REACH-NLRI will be advertised in an update message and the corresponding received NLRIs must be considered to instantiate the data plane.

Default shutdown

site

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

Context config>service>epipe

Description This command configures a Epipe 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 Epipe service.

boot-timer

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

Context config>service>epipe>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 — 600

sap

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

Context config>service>epipe>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.

site-activation-timer

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

Context config>service>epipe>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 is 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>epipe>site

Description This command configures the identifier for the site in this service. It must match between services but it is local to the service.

Parameters *value* — Specifies the site identifier.

Values 1 — 65535

site-preference

Syntax **site-preference** *preference-value*
site-preference {**primary**|**backup**}
no site-preference

Context config>service>epipe>site

Description This command defines the value to advertise in the VPLS preference field of the BGP VPWS and BGP Multi-homing NLRI extended community. This value can be changed without having to shutdown the site itself. The site-preference is only applicable to VPWS services.

When not configured, the default is zero, indicating that the VPLS preference is not in use.

Default no site-preference, value=0

Parameters *preference-value* — Specifies the preference value to advertise in the NLRI L2 extended community for this site.

Values 1 — 65535

Parameters **primary —** Sets the site-preference to 65535.

backup — Sets the site-preference to 1.

ce-address-discovery

Syntax [**no**] **ce-address-discovery** [**ipv6**]

Context config>service>ipipe

Description This command specifies whether the service will automatically discover the CE IP addresses.

When enabled, the addresses will be automatically discovered on SAPs that support address discovery, and on the spoke SDPs. When enabled, addresses configuration on the Ipipe SAP and spoke SDPs will not be allowed.

If disabled, CE IP addresses must be manually configured for the SAPs to become operationally up.

Default no ce-address-discovery

Parameters **ipv6 —** The **ipv6** keyword enables IPv6 CE address discovery support on the Ipipe so that both IPv4 and IPv6 address discovery are supported. If the **ipv6** keyword is not included, then only IPv4 address discovery is supported and IPv6 packets are dropped. This feature requires IOM2 or better. It requires chassis mode C or above. If any Ipipe services require IPv6 support, then all network ports on the node must be configured on 7750 IOM-3-XPs.

stack-capability-signaling

Syntax [**no**] **stack-capability-signaling**

Context config>service>ipipe

Description This command enables stack capability signaling in the initial label mapping message of the ipipe PW to indicate that IPv6 is supported.

When enabled, the 7750 includes the stack capability TLV with the IPv6 stack bit set according to the **ce-address-discovery ipv6** keyword, and also checks the value of the stack-capability TLV received from the far end.

This command must be blocked if no **ce-address-discovery** is specified, or the **ipv6** keyword is not included with the **ce-address-discovery** command.

This command if only applicable to the ipipe service and must be blocked for all other services.

This command has no effect if both SAPs on the ipipe service are local to the node.

This feature requires IOM2 or better. It requires chassis mode C or above. If any Ipipe services require IPv6 support, then all network ports on the node must be configured on 7750 IOM-3-XPs.

Default no stack-capability-signaling

endpoint

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

Context config>service>apipe
config>service>cpipe
config>service>fpipe

config>service>epipe
config>service>ipipe

Description This command configures a service endpoint.

Parameters *endpoint-name* — Specifies an endpoint name.

per-service-hashing

Syntax **[no] per-service-hashing**

Context config>service>epipe

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

tunnel

Syntax **tunnel** *service-id* **backbone-dest-mac** *ieee-address* **isid** *ISID*
no tunnel

Context config>service>epipe>pbb

Description This command configures a Provider Backbone Bridging (PBB) tunnel with Backbone VPLS (B-VPLS) service information.

Parameters *service-id* — Specifies the B-VPLS service for the PBB tunnel associated with this service.

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

backbone-dest-mac *ieee-address* — Specifies the backbone destination MAC-address for PBB packets.

isid *ISID* — Specifies a 24 bit service instance identifier for the PBB tunnel associated with this service. As part of the PBB frames, it is used at the destination PE as a demultiplexor field.

Values 0 — 16777215

active-hold-delay

Syntax **active-hold-delay** *active-hold-delay*
no active-hold-delay

Context config>service>cpipe>endpoint
config>service>apipe>endpoint
config>service>epipe>endpoint
config>service>fpipe>endpoint
config>service>ipipe>endpoint

Description This command specifies that the node will delay sending the change in the T-LDP status bits for the VLL endpoint when the MC-LAG transitions the LAG subgroup which hosts the SAP for this VLL endpoint from **active** to **standby** or when any object in the endpoint. For example, SAP, ICB, or regular spoke SDP, transitions from up to down operational state.

By default, when the MC-LAG transitioned the LAG subgroup which hosts the SAP for this VLL endpoint from **active** to **standby**, the node sends immediately new T-LDP status bits indicating the new value of "standby" over the spoke SDPs which are on the mate-endpoint of the VLL. The same applies when any object in the endpoint changes an operational state from up to down.

There is no delay applied to the VLL endpoint status bit advertisement when the MC-LAG transitions the LAG subgroup which hosts the SAP from **standby** to **active** or when any object in the endpoint transitions to an operationally up state.

Default 0 — A value of zero means that when the MC-LAG transitioned the LAG subgroup which hosts the SAP for this VLL endpoint from **active** to **standby**, the node sends immediately new T-LDP status bits indicating the new value of **standby** over the spoke SDPs which are on the mate-endpoint of the VLL. The same applies when any object in the endpoint changes an operational state from up to down.

Parameters **active-hold-delay** — Specifies the active hold delay in 100s of milliseconds.

Values 0 — 60

revert-time

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

Context config>service>apipe>endpoint

config>service>fpipe>endpoint
config>service>cpipe>endpoint
config>service>epipe>endpoint
config>service>ipipe>endpoint

Description This command configures the time to wait before reverting back to the primary spoke SDP defined on this service endpoint, after having failed over to a backup spoke SDP.

Parameters *revert-time* — Specify the time, in seconds, to wait before reverting to the primary SDP.

Values 0 — 600

Values 0

infinite — Causes the endpoint to be non-revertive.

standby-signaling-master

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

Context config>service>vll>endpoint

Description When this command is enabled, the pseudowire standby bit (value 0x00000020) will be sent to T-LDP peer for each spoke-sdp of the endpoint that is selected as a standby.

This command is mutually exclusive with a VLL mate SAP created on a mc-lag/mc-aps or ICB. It is also mutually exclusive with vc-switching.

Default standby-signaling-master

standby-signaling-slave

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

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

Description When this command is enabled, the node will block the transmit forwarding direction of a spoke SDP based on the pseudowire standby bit received from a T-LDP peer.

This command is present at the endpoint level as well as the spoke-SDP level. If the spoke SDP is part of an explicit-endpoint, it will not be possible to change this setting at the spoke-sdp level. An existing spoke SDP can be made part of the explicit endpoint only if the settings do not conflict. A newly created spoke SDP, which is part of a given explicit-endpoint, will inherit this setting from the endpoint configuration.

This command is mutually exclusive with an endpoint that is part of an mc-lag, mc-aps or an ICB.

If the command is disabled, the node assumes the existing independent mode of behavior for the forwarding on the spoke SDP.

Default disabled

interworking

Syntax **interworking** {**frf-5**}
no interworking

Context config>service>apipe

Description This command specifies the interworking function that should be applied for packets that ingress/egress SAPs that are part of an Apipe service.

Interworking is applicable only when the two endpoints (i.e., the two SAPs or the SAP and the spoke-sdp) are of different types. Also, there are limitations on the combinations of SAP type, vc-type, and interworking values as shown in the following table.

SAP Type	Allowed VC-Type Value	Allowed Interworking Value
ATM VC	atm-vcc, atm-sdu	none
	fr-dlci	Not Supported
FR DLCI	fr-dlci	none
	atm-sdu	frf-5

Default **none** (Interworking must be configured before adding a Frame-Relay SAP to an Apipe service.)

Parameters **frf-5 —** Specify Frame Relay to ATM Network Interworking (FRF.5).

service-name

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

Context config>service>apipe
config>service>cpipe
config>service>fpipe
config>service>ipipe
config>service>epipe

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 SR OS platforms.

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

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

service-mtu

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

Context config>service>epipe
config>service>ipipe
config>service>apipe
config>service>cpipe
config>service>fpipe

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.

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.

Because this connects a Layer 2 to a Layer 3 service, adjust either the service-mtu under the Epipe service. The MTU that is advertised from the Epipe side is service-mtu minus EtherHeaderSize.

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

By default if no service-mtu is configured it is (1514 - 14) = 1500.

Default apipe, fpipe: 1508

ipipe: 1500

epipe: 1514

The following table displays MTU values for specific VC types.

SAP 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 (Q-in-Q with preserved bottom Qtag)	1518	1504

octets — The size of the MTU in octets, expressed as a decimal integer, between 1 — 9194.

signaled-vc-type-override

Syntax **signaled-vc-type-override atm-vcc**
no signaled-vc-type-override

Context <root>

Description This command overrides the pseudowire type signaled to type 0x0009 N:1 VCC cell within an Apipe VLL service of vc-type atm-cell. Normally, this service vc-type signals a pseudowire of type 0x0003 ATM Transparent Cell.

This command is not allowed in an Apipe VLL of vc-type value atm-cell if a configured ATM SAP is not using a connection profile. Conversely, if the signaling override command is enabled, only an ATM SAP with a connection profile assigned will be allowed.

The **override** command is not allowed on Apipe VLL service of vc-type value other than atm-cell. It is also not allowed on a VLL service with the vc-switching option enabled since signaling of the PW FEC in a Multi-Segment PW (MS-PW) is controlled by the T-PE nodes. Thus for this feature to be used on a MS-PW, it is required to configure an Apipe service of vc-type atm-cell at the T-PE nodes with the signaled-vc-type-override enabled, and to configure a Apipe VLL service of vc-type atm-vcc at the S-PE node with the vc-switching option enabled.

The **no** form of this command returns the Apipe VLL service to signal its default pseudowire type

Default none

Parameters **atm-vcc** — Specifies the pseudowire type to be signaled in the pseudowire establishment.

connection-profile

Syntax **connection-profile** *conn-prof-id* [**create**]
no connection-profile *conn-prof-id*

Context <root>

Description This command creates a profile for the user to configure the list of discrete VPI/VCI values to be assigned to an ATM SAP of an Apipe VLL of **vc-type atm-cell.**

A connection profile can only be applied to a SAP which is part of an Apipe VLL service of **vc-type atm-cell**. The ATM SAP can be on a regular port or APS port.

A maximum of 8000 connection profiles can be created on the system.

The **no** form of this command deletes the profile from the configuration.

Default none

Parameters *conn-prof-id* — Specifies the profile number.

Values 1 — 8000

member

Syntax **member** *encap-value* [**create**]

no member *encap-value*

Context config>connection-profile

Description This command allows the adding of discrete VPI/VCI values to an ATM connection profile for assignment to an ATM SAP of an Apipe VLL of vc-type atm-cell.

Up to a maximum of 16 discrete VPI/VCI values can be configured in a connection profile. The user can modify the content of a profile which triggers a re-evaluation of all the ATM SAPs which are currently using the profile.

The **no** form of this command deletes the member from the configuration..

Default none

Parameters *encap-value* — Specifies the VPI and VCI values of this connection profile member.

Values vpi: NNI: 0 — 4095; UNI: 0 — 255
vci: 1, 2, 5 — 65535

VLL SAP Commands

sap

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

Context config>service>apipe
config>service>cpipe
config>service>fpipe
config>service>ipipe
config>service>epipe

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 device. 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. 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 following are supported:

- ATM VPI/VCI on an ATM port for vc-type atm-vcc and atm-sdu
- ATM VPI on an ATM port for vc-type atm-vpc
- ATM virtual trunk - a range of VPIs on an ATM port for vc-type atm-cell
- ATM port for vc-type atm-cell
- ATM connection profile for vc-type atm-cell
- Frame Relay DLCI on a port for vc-type atm-sdu
- ATM SAP carries the IPv4 packet using RFC 2684, VC-Mux or LLC/SNAP routed PDU encapsulation for an Ipipe service
- Frame Relay SAP RFC 2427, routed PDU encapsulation for an Ipipe service
- Ethernet SAP RFC 1332, PPP IPCP encapsulation of an IPv4 packet for an Ipipe service
- Ethernet SAP HDLC SAP uses the routed IPv4 encapsulation for an Ipipe service

- ATM - Frame Relay, PPP/IPCP - PPP/IPCP
- Frame Relay-Frame Relay, ATM - ATM
- Ethernet-Ethernet
- cHDLC-cHDLC
- Ethernet SAPs support null, dot1q, and qinq
- An ATM SAP can be part of an IMA bundle.
- A PPP SAP can be part of an MLPPP bundle.
- A FR SAP can be part of a MLFR bundle.

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 SAP can be defined with Ethernet ports, SONET/SDH or TDM channels. At most, only one sdp-id can be bound to an VLL service. Since a VLL is a point-to-point service, it can have, at most, two end points. The two end points can be one SAP and one SDP or two SAPs. Up to 49 SDPs can be associated with a service in a single router. Each SDP must have a unique router destination or an error will be generated.

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).

Two Frame Relay SAPs cannot be configured on an Apipe service. The limitation is for an Apipe service in local mode, which has two SAPs associated with the service, as opposed to a configuration with a SAP and a SDP in remote case, the only combination of the type of SAPs allowed is either two ATM SAPs or an ATM SAP and a Frame Relay SAP. The CLI prevents adding two Frame Relay SAPs under an Apipe service.

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

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

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

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

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

endpoint — Adds a SAP endpoint association.

no endpoint — removes the association of a SAP or a spoke-sdp with an explicit endpoint name.

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

Configuration Example

```
*A:bksim2801>config>service>apipe>sap$
===============================================================
ATM PVCs, Port 1/1/1
===============================================================
VPI/VCI     Owner     Type     Ing.TD   Egr.TD  Adm  OAM     Opr
---------------------------------------------------------------
2/102       SAP       PVC      1        1       up   ETE-AIS dn
10/100      SAP       PVC      1        1       up   ETE-AIS dn
===============================================================
*A:bksim2801#
```

sap

Syntax [**no**] **sap eth-tunnel-***tunnel-id*[:*eth-tunnel-sap-id*] [**create**]

Context config>service>epipe
config>service>ipipe
config>service>vpls

Description This command configures an Ethernet tunnel SAP.

An Ethernet tunnel control SAP has the format eth-tunnel-*tunnel-id* and is not configured with an Ethernet tunnel SAP ID. No Ethernet tunnel tags can be configured under a control SAP since the control SAP uses the control tags configured under the Ethernet tunnel port. This means that at least one member port and control tag must be configured under the Ethernet tunnel port before this command is executed. The control SAP is needed for carrying G.8031 and 802.1ag protocol traffic. This SAP can also carry user data traffic.

An Ethernet tunnel same-fate SAP has the format eth-tunnel-*tunnel-id:eth-tunnel-sap-id*. Same-fate SAPs carry only user data traffic. Multiple same-fate SAPs can be configured on one Ethernet tunnel port and share the fate of that port, provided the SAPs are properly configured with corresponding tags.

Ethernet tunnel SAPs are supported under VPLS, Epipe and Ipipe services only.

Default no sap

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

Values 1 — 1024

eth-tunnel-sap-id — Specifies a SAP ID of a same-fate SAP.

Values 0 — 4094

lag-link-map-profile

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

Context config>service>epipe>sap
config>service>ipipe>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's/network interface's 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 the LAG the SAP/network interface exists on.

agg-rate-limit

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

Context config>service>epipe>sap>ingress

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

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

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

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

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

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

Values 1 — 40000000, max

agg-rate-limit

Syntax **agg-rate-limit** *agg-rate* [**queue-frame-based-accounting**]
agg-rate-limit *agg-rate* [**queue-frame-based-accounting**] (Epipe services)
no agg-rate-limit

Context config>service>cpipe>sap>egress
config>service>ipipe>sap>egress

config>service>fpipe>sap>egress
config>service>epipe>sap>egress

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

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

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

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

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

Parameters *agg-rate* — Defines the rate, in kilobits-per-second, that the maximum aggregate rate 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. Note that this parameter is configurable in Epipe VLL services.

Default Frame based accounting is disabled by default.

policer-control-override

Syntax **policer-control-override** [**create**]
no policer-control-override

Context config>service>apipe>sap>egress
config>service>apipe>sap>ingress
config>service>fpipe>sap>egress
config>service>fpipe>sap>ingress
config>service>epipe>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>apipe>sap>egress>policer-control-override
config>service>apipe>sap>ingress>policer-control-override
config>service>fpipe>sap>egress>policer-control-override
config>service>fpipe>sap>ingress>policer-control-override
config>service>epipe>sap>egress>policer-control-override

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>apipe>sap>egress>policer-control-override
config>service>apipe>sap>ingress>policer-control-override
config>service>fpipe>sap>egress>policer-control-override
config>service>fpipe>sap>ingress>policer-control-override
config>service>epipe>sap>egress>policer-control-override

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>apipe>sap>egress>policer-control-override>priority-mbs-threshold

config>service>apipe>sap>ingress>policer-control-override>priority-mbs-threshold
config>service>fpipe>sap>egress>policer-control-override>priority-mbs-threshold
config>service>fpipe>sap>ingress>policer-control-override>priority-mbs-threshold
config>service>epipe>sap>egress>policer-control-override>priority-mbs-threshold

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>apipe>sap>egress>policer-control-override>priority-mbs-threshold
config>service>apipe>sap>ingress>policer-control-override>priority-mbs-threshold
config>service>fpipe>sap>egress>policer-control-override>priority-mbs-threshold
config>service>fpipe>sap>ingress>policer-control-override>priority-mbs-threshold
config>service>epipe>sap>egress>policer-control-override>priority-mbs-thresholds

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>apipe>sap>egress>policer-control-override>priority-mbs-threshold>priority
config>service>apipe>sap>ingress>policer-control-override>priority-mbs-threshold>priority
config>service>fpipe>sap>egress>policer-control-override>priority-mbs-threshold>priority
config>service>fpipe>sap>ingress>policer-control-override>priority-mbs-threshold>priority
config>service>epipe>sap>egress>policer-control-override>priority-mbs-threshold>priority
config>service>epipe>sap>ingress>policer-control-override>priority-mbs-threshold>priority

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>apipe>sap>egress
config>service>apipe>sap>ingress
config>service>fpipe>sap>egress
config>service>fpipe>sap>ingress
config>service>epipe>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.

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 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>apipe>sap>egress
config>service>apipe>sap>ingress
config>service>fpipe>sap>egress
config>service>fpipe>sap>ingress
config>service>epipe>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>apipe>sap>egress>policer-override
config>service>apipe>sap>ingress>policer-override
config>service>fpipe>sap>egress>policer-override
config>service>fpipe>sap>ingress>policer-override
config>service>epipe>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-in-kbytes*
no cbs

Context config>service>apipe>sap>egress>policer-override>policer
config>service>apipe>sap>ingress>policer-override>policer
config>service>epipe>sap>egress>policer-override>policer
config>service>epipe>sap>egress>queue-override>queue

Description This command can be used to override specific attributes of the specified queue's CBS parameters.

It is permissible, and possibly desirable, to oversubscribe the total CBS reserved buffers for a given access port egress buffer pool. Oversubscription may be desirable due to the potential large number of service queues and the economy of statistical multiplexing the individual queue's CBS 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>apipe>sap>egress>policer-override
config>service>apipe>sap>ingress>policer-override>policer
config>service>epipe>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.

kilobytes — When kilobytes is defined, the value given for size is interpreted as the queue's MBS value given in kilobytes.

packet-byte-offset

Syntax **packet-byte-offset** {**add** *add-bytes* | **subtract** *sub-bytes*}

Context config>service>apipe>sap>egress>policer-override>policer
config>service>apipe>sap>ingress>policer-override>policer
config>service>epipe>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

percent-rate

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

Context config>service>apipe>sap>egress>policer-override>policer
config>service>apipe>sap>ingress>policer-override>policer

config>service>epipe>sap>egress>policer-override>policer

Description This command configures the percent rates (CIR and PIR) override.

Parameters *pir-rate* — The pir-percent parameter is used to express the policer's PIR as a percentage of the policers's parent arbiter rate.

Values Percentage ranging from 0.01 to 100.00. The default is 100.00.

cir *cir-rate* **—** Configures the administrative CIR specified by the user.

Values 0 — 20000000, max

percent-rate

Syntax **percent-rate** *pir-percent* [**cir** *cir-percent*] [**port-limit**|**local-limit**]
no percent-rate

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

Description The percent-rate command within the SAP ingress and egress QOS policy enables supports for a queue's PIR and CIR rate to be configured as a percentage of the egress port's line rate or of its parent scheduler's rate.

When the rates are expressed as a port-limit, the actual rates used per instance of the queue will vary based on the port speed. For example, when the same QOS policy is used on a 1-Gigabit and a 10-Gigabit Ethernet port, the queue's rates will be 10 times greater on the 10 Gigabit port due to the difference in port speeds. This enables the same QOS policy to be used on SAPs on different ports without needing to use SAP based queue overrides to modify a queue's rate to get the same relative performance from the queue.

If the port's speed changes after the queue is created, the queue's PIR and CIR rates will be recalculated based on the defined percentage value.

Values When the rates are expressed as a local-limit, the actual rates used per instance of the queue are relative to the queue's parent scheduler rate. This enables the same QOS policy to be used on SAPs with different parent scheduler rates without needing to use SAP based queue overrides to modify a queue's rate to get the same relative performance from the queue.

If the parent scheduler rate changes after the queue is created, the queue's PIR and CIR rates will be recalculated based on the defined percentage value.

Queue rate overrides can only be specified in the form as configured in the QoS policy (a SAP override can only be specified as a percent-rate if the associated QoS policy was also defined as percent-rate). Likewise, a SAP override can only be specified as a rate (kbps) if the associated QoS policy was also defined as a rate. Queue-overrides are relative to the limit type specified in the QOS policy.

When no percent-rate is defined within a SAP ingress or egress queue-override, the queue reverts to the defined shaping and CIR rates within the SAP ingress and egress QOS policy associated with the queue.

Parameters *percent-of-line-rate* — The percent-of-line-rate parameter is used to express the queue's shaping rate as a percentage of line rate. The line rate associated with the queue's port may dynamically

change due toconfiguration or auto-negotiation. The line rate may also be affected by an egress port scheduler defined max-rate.

pir-percent — The pir-percent parameter is used to express the queue's PIR as a percentage dependant on the use of the port-limit or local-limit.

Values Percentage ranging from 0.01 to 100.00. The default is 100.00.

pir-percent — The pir-percent parameter is used to express the queue's PIR as a percentage dependant on the use of the port-limit or local-limit.

cir *cir-percen* **—** The cir keyword is optional and when defined the required cir-percent CIR parameter expresses the queue's CIR as a percentage dependant on the use of the port-limit or local-limit.

Percentage ranging from 0.00 to 100.00. The default is 100.00

rate

Syntax **rate** {*rate* | **max**} [**cir** {**max** | *rate*}]

Context config>service>apipe>sap>egress>policer-override>policer
config>service>apipe>sap>ingress>policer-override>policer
config>service>epipe>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** st*at-mode*
no stat-mode

Context config>service>apipe>sap>egress>policer-override>policer
config>service>apipe>sap>ingress>policer-override>policer

config>service>epipe>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.

ce-address

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

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

Description This command specifies the IP address of the CE device associated with an Ipipe SAP or spoke SDP. In the case of a SAP, it is the address of the CE device directly attached to the SAP. For a spoke SDP, it is the address of the CE device reachable through that spoke SDP (for example, attached to the SAP on the remote node). The address must be a host address (no subnet addresses are accepted) as there must be only one CE device attached to an Ipipe SAP. The CE address specified at one end of an Ipipe will be used in processing ARP messages at the other endpoint, as the router acts as a proxy for ARP messages.

Parameters. *ip-address* — specifies the IP address of the CE device associated with an Ipipe SAP.

qinq-mark-top-only

Syntax [**no**] **qinq-mark-top-only**

Context config>service>cpipe>sap>egress
config>service>apipe>sap>egress
config>service>epipe>sap>egress
config>service>fpipe>sap>egress
config>service>apipe>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 the enabled, only the P-bits/DEI bit in the top Q-tag are marked.

Default no qinq-mark-top-only

multi-service-site

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

Context config>service>ipipe>sap
config>service>apipe>sap
config>service>cpipe>sap
config>service>fpipe>sap
config>service>epipe>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.

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.

ring-node

Syntax **ring-node** *ring-node-name*
no ring-node

Context config>service>epipe>sap

Description This command configures a multi-chassis ring-node for this SAP.

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

Default none

tod-suite

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

Context config>service>apipe>sap
config>service>cpipe>sap
config>service>fpipe>sap
config>service>ipipe>sap
config>service>epipe>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.

transit-policy

Syntax **transit-policy prefix** *prefix-aasub-policy-id*
no transit-policy

Context config>service>epipe>sap

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

use-broadcast-mac

Syntax [**no**] **use-broadcast-mac**

Context config>service>ipipe>sap

Description This command enables the user of a of broadcast MAC on SAP.

An Ipipe VLL service with the **ce-address-discovery** command enabled forwards unicast IP packets using the broadcast MAC address until the ARP cache is populated with a valid entry for the CE IP and MAC addresses.

The **no** form of this command enables the user of a of broadcast MAC on SAP.

Default no use-broadcast-mac

mac

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

Context config>service>ipipe>sap

Description This command assigns a specific MAC address to an Ipipe SAP.

The **no** form of this command returns the MAC address of the SAP to the default value.

Default The physical MAC address associated with the Ethernet interface where the SAP is configured.

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-refresh

Syntax **mac-refresh** *refresh interval*
no mac-refresh

Context config>service>ipipe>sap

Description This command specifies the interval between ARP requests sent on this Ipipe SAP. When the SAP is first enabled, an ARP request will be sent to the attached CE device and the received MAC address will be used in addressing unicast traffic to the CE. Although this MAC address will not expire while the Ipipe SAP is enabled and operational, it is verified by sending periodic ARP requests at the specified interval.

The **no** form of this command restores mac-refresh to the default value.

Default 14400

Parameters *refresh interval* — Specifies the interval, in seconds, between ARP requests sent on this Ipipe SAP.

Values 0 — 65535

accounting-policy

Syntax **accounting-policy** *acct-policy-id*
no accounting-policy

Context config>service>apipe>sap
config>service>cpipe>sap
config>service>epipe>sap
config>service>epipe>spoke-sdp
config>service>fpipe>sap
config>service>ipipe

Description This command creates the accounting policy context that can be applied to a SAP.

An accounting policy must be defined before it can be associated with a SAP. If the *policy-id* does not exist, an error message is generated.

A maximum of one accounting policy can be associated with a SAP 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, 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>epipe>sap
config>service>epipe>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.

bandwidth

Syntax **bandwidth** *bandwidth*
no bandwidth

Context config>service>epipe>spoke-sdp
config>service>fpipe>spoke-sdp
config>service>apipe>spoke-sdp
config>service>ipipe>spoke-sdp
config>service>cpipe>spoke-sdp

Description This command specifies the bandwidth to be used for VLL bandwidth accounting by the VLL CAC feature.

The service manager keeps track of the available bandwidth for each SDP. The maximum value is the sum of the bandwidths of all constituent LSPs in the SDP. The SDP available bandwidth is adjusted by the user configured booking factor.

If an LSP consists of a primary and many secondary standby LSPs, then the bandwidth used in the maximum SDP available bandwidth is that of the active path. Any change to and LSP active path bandwidth will update the maximumn SDP available bandwidth. Note however that a change to any constituent LSP bandwidth due to re-signaling of the primary LSP path or the activation of a secondary path which causes overbooking of the maximum SDP available bandwidth causes a warning and a trap to be issued but no further action is taken. The activation of a bypass or detour LSP in the path of the primary LSP does not change the maximum SDP available bandwidth.

When the user binds a VLL service to this SDP, an amount of bandwidth equal to bandwidth is subtracted from the SDP available bandwidth adjusted by the booking factor. When the user deletes this VLL service binding from this SDP, an amount of bandwidth equal to bandwidth is added back into the SDP available bandwidth.

If the total SDP available bandwidth when adding this VLL service is about to overbook, a warning is issued and the binding is rejected. This means that the spoke-sdp bandwidth does not update the maximum SDP available bandwidth. In this case, the spoke-sdp is put in operational down state and a status message of “pseudowire not forwarding” is sent to the remote SR-Series PE node. A trap is also generated. The service manager will not put the spoke-sdp into operational UP state until the user performs a shutdown/no-shutdown of the spoke-sdp and the bandwidth check succeeds. Thus, the service manager will not automatically audit spoke-sdp’s subsequently to their creation to check if bandwidth is available.

If the VLL service contains an endpoint with multiple redundant spoke-sdp's, each spoke-sdp will have its bandwidth checked against the available bandwidth of the corresponding SDP.

If the VLL service performs a pseudowire switching (VC switching) function, each spoke-sdp is separately checked for bandwidth against the corresponding SDP.

Note this feature does not alter the way service packets are sprayed over multiple RSVP LSPs, which are part of the same SDP. In other words, by default load balancing of service packets occurs over the SDP LSP's based on service-id, or based on a hash of the packet header if ingress SAP shared queuing is enabled. In both cases, the VLL bandwidth is not checked against the selected LSP(s) available bandwidth but on the total SDP available bandwidth. Thus, if there is a single LSP per SDP, these two match.

If class-forwarding is enabled on the SDP, VLL service packets are forwarded to the SDP LSP which the packet forwarding class maps to, or if this is down to the default LSP. However, the VLL bandwidth is not checked against the selected LSP available bandwidth but on the total SDP available bandwidth. If there is a single LSP per SDP, these two match.

If a non-zero bandwidth is specified for a VLL service and attempts to bind the service to an LDP or a GRE SDP, a warning is issued that CAC failed but the VLL is established. A trap is also generated.

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

Values 0 — 100000000, max in units of kilobits/sec.

Default 0

block-on-peer-fault

Syntax [**no**] **block-on-peer-fault**

Context config>service>epipe>spoke-sdp

Description When enabled, this command blocks the transmit direction of a PW when any of the following PW status codes is received from the far end PE:

0x00000001	Pseudowire Not Forwarding
0x00000002	Local Attachment Circuit (ingress) Receive Fault
0x00000004	Local Attachment Circuit (egress) Transmit Fault
0x00000008	Local PSN-facing PW (ingress) Receive Fault
0x00000010	Local PSN-facing PW (egress) Transmit Fault

The transmit direction is unblocked when the following PW status code is received:

0x00000000	Pseudowire forwarding (clear all failures)

This command is mutually exclusive with **no pw-status-signaling**, and **standby-signaling-slave**. It is not applicable to spoke SDPs forming part of an MC-LAG or spoke SDPs in an endpoint.

Default no block-on-peer-fault

collect-stats

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

Context config>service>cpipe>sap
config>service>cpipe>spoke-sdp
config>service>epipe>spoke-sdp
config>service>apipe>sap
config>service>fpipe>sap
config>service>epipe>sap

Description This command enables accounting and statistical data collection for either the SAP, 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

cpu-protection

Syntax **cpu-protection** *policy-id* [**mac-monitoring**] | [**eth-cfm-monitoring** [**aggregate**][**car**]]
no cpu-protection

Context config>service>apipe>sap
config>service>epipe>spoke-sdp
config>service>epipe>sap

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.

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.

If **no** CPU protection policy is assigned to a service SAP then a the default policy is used to limit the overall-rate.

Parameters *policy-id* — Specifies an existing CPU protection policy.

Values 1 — 255

mac-monitoring — This keyword enables MAC monitoring.

eth-cfm-monitoring — This keyword enables Ethernet Connectivity Fault Management monitoring.

aggregate — This keyword applies the rate limit to the sum of the per peer packet rates.

car — (Committed Access Rate) This keyword causes Eth-CFM packets to be ignored when enforcing the overall-rate.

dist-cpu-protection

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

Context config>service>epipe>sap
config>service>apipe>sap
config>service>cpipe>sap
config>service>fpipe>sap
config>service>ipipe>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 an msap-policy)

Default no dist-cup-protection

ethernet

Syntax **ethernet**

Context config>service>epipe>sap

Description Use this command to configure Ethernet properties in this SAP.

llf

Syntax [**no**] **llf**

Context config>service>apipe>sap>atm
config>service>epipe>sap>ethernet

Description This command enables Link Loss Forwarding (LLF) on an Ethernet port or an ATM port. This feature provides an end-to-end OAM fault notification for Ethernet VLL serviceand for ATM VLL service of vc-type atm-cell. It brings down the Ethernet port (Ethernet LLF) or sends a SONET/SDH Path AIS (ATM LLF) towards the attached CE when there is a local fault on the Pseudowire or service, or a remote fault on the SAP or pseudowire, signaled with label withdrawal or T-LDP status bits. It ceases when the fault disappears.

The Ethernet port must be configured for null encapsulation.

The ATM port must be configured as a SAP on an apipe service of vc-type atm-cell. The ATM port must also be configured on the following MDAs:

1-port OC12/STM4 ASAP MDA. At OC3/STM1 port level

4-port ATM MDA at OC12/STM4 or OC3/STM1 port level

16-port ATM MDA at OC3/STM1 port level

The ATM port must be configured as a SAP on an apipe service of vc-type atm-cell. The ATM port must also be configured on the following MDAs:

1-port OC12/STM4 ASAP MDA. At OC3/STM1 port level

4-port ATM MDA at OC12/STM4 or OC3/STM1 port level

16-port ATM MDA at OC3/STM1 port level

Circuit Emulation Commands

cem

Syntax **cem**

Context config>service>cpipe>sap
config>service>epipe>sap

Description This command enables the context to specify circuit emulation (CEM) properties.

local-ecid

Syntax **local-ecid** *emulated circuit identifier*
no local-ecid

Context config>service>epipe>sap>cem

Description This command defines the Emulated Circuit Identifiers (ECID) to be used for the local (source) end of the circuit emulation service.

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

Default 65535

Parameters *emulated circuit identifier* — Specifies the value to be used as the local (source) ECID for the circuit emulation service. On CES packet reception, the ECID in the packet will be compared to the configured local-ecid value. These must match for the packet payload to be used for the TDM circuit. The remote-ecid value is inserted into the MEF-8 CES packet to be transmitted.

Values 0 — 1048575

packet

Syntax **packet jitter-buffer** *milliseconds* [**payload-size** *bytes*]
packet payload-size *bytes*
no packet *bytes*

Context config>service>cpipe>sap
config>service>epipe>sap>cem

Description This command specifies the jitter buffer size, in milliseconds, and payload size, in bytes.

Default The default value depends on the CEM SAP endpoint type, and if applicable, the number of timeslots:

Endpoint Type	Timeslots	Default Jitter Buffer (in ms)
unstructuredE1	n/a	5
unstructuredT1	n/a	5
unstructuredE3	n/a	5
unstructuredT3	n/a	5
nxDS0 (E1/T1)		32
	N = 1	16
	N = 2..4	8
	N = 5..15	5
nxDS0WithCas (E1)	N	8
nxDS0WithCas (T1)	N	12

Parameters *milliseconds* — specifies the jitter buffer size in milliseconds (ms).

Configuring the payload size and jitter buffer to values that result in less than 2 packet buffers or greater than 32 packet buffers is not allowed. Setting the jitter butter value to 0 sets it back to the default value.

Values 1 — 250

payload-size *bytes* — Specifies the payload size (in bytes) of packets transmitted to the packet service network (PSN) by the CEM SAP. This determines the size of the data that will be transmitted over the service. If the size of the data received is not consistent with the payload size then the packet is considered malformed.

Default The default value depends on the CEM SAP endpoint type, and if applicable, the number of timeslots:

Endpoint Type	Timeslots	Default Payload Size (in bytes)
unstructuredE1	n/a	256
unstructuredT1	n/a	192
unstructuredE3	n/a	1024
unstructuredT3	n/a	1024
nxDS0 (E1/T1)	N = 1	64
	N = 2..4	N x 32
	N = 5..15	N x 16
	N >= 16	N x 8
nxDS0WithCas (E1)	N	N x 16
nxDS0WithCas (T1)	N	N x 24

For all endpoint types except for nxDS0WithCas, the valid payload size range is from the default to 2048 bytes.

For nxDS0WithCas, the payload size divide by the number of timeslots must be an integer factor of the number of frames per trunk multi-frame (for example, 16 for E1 trunk and 24 for T1 trunk).

For 1xDS0, the payload size must be a multiple of 2.

For NxDS0, where N > 1, the payload size must be a multiple of the number of timeslots.

For unstructuredE1, unstructuredT1, unstructuredE3 and unstructuredT3, the payload size must be a multiple of 32 bytes.

Configuring the payload size and jitter buffer to values that result in less than 2 packet buffers or greater than 32 packet buffer is not allowed.

Setting the payload size to 0 sets it back to the default value.

Values 0, 16 — 2048

remote-ecid

Syntax **remote-ecid** *emulated circuit identifier*
no remote-ecid

Context config>service>epipe>sap>cem

Description This command defines the Emulated Circuit Identifiers (ECID) to be used for the remote (destination) end of the circuit emulation service.

Parameters *emulated circuit identifier* — Specifies the value to be used as the remote (destination) ECID for the circuit emulation service. Upon CES packet reception, the ECID in the packet will be compared to the configured local-ecid value. These must match for the packet payload to be used for the TDM circuit. The remote-ecid value is inserted into the MEF-8 CES packet to be transmitted.

remote-mac

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

Context config>service>epipe>sap>cem

Description This command defines the destination IEEE MAC address to be used to reach the remote end of the circuit emulation service.

Default 00:00:00:00:00:00

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. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

report-alarm

Syntax [**no**] **report-alarm** [**stray**] [**malformed**] [**pktloss**] [**overrun**] [**underrun**] [**rpktloss**] [**rfault**] [**rrdi**]

Context config>service>epipe>sap>cem

Description This command indicates the type of CEM SAP alarm.

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

Default On: stray, malformed, pktloss and overrun

Off: rpktloss, rfault, rrdi

Parameters **stray —** Reports the reception of packets not destined for this CES circuit.

malformed — Reports the reception of packet not properly formatted as CES packets.

pktloss — Reports the lack of reception of CES packets.

overrun — Reports reports the reception of too many CES packets resulting in a overrun of the receive jitter buffer.

underrun — Reportsreports the reception of too few CES packets resulting in a overrun of the receive jitter buffer.

rpktloss — Reports hat the remote peer is currently in packet loss status.

rfault — Reports that the remote TDM interface is currently not in service.

rrdi — Reports that the remote TDM interface is currently in RDI status.

rtp-header

Syntax [**no**] **rtp-header**

Context config>service>epipe>sap>cem
config>service>cpipe>sap>cem

Description This command specifies whether an RTP header is used when packets are transmitted to the packet service network (PSN) by the CEM SAP. This mode must be enabled for differential-timed DS1/E1s. It can optionally be enabled for other DS1/E1s for interoperability purposes.

Default no rtp-header

ETH-CFM Service Commands

eth-cfm

Syntax **eth-cfm**

Context config>service>epipe>spoke-sdp
config>service>epipe
config>service>epipe>sap
config>service>ipipe>sap

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

ais-enable

Syntax [**no**] **ais-enable**

Context config>service>epipe>sap>eth-cfm
config>service>epipe>sap>eth-cfm>mep
config>service>epipe>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>epipe>sap>eth-cfm>mep
config>service>epipe>spoke-sdp>eth-cfm>aid-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>epipe>sap>eth-cfm>mep
config>service>epipe>spoke-sdp>eth-cfm>aid-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>epipe>sap>eth-cfm>mep
config>service>epipe>spoke-sdp>eth-cfm>aid-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

eth-tunnel

Syntax **eth-tunnel**

Context config>service>epipe>sap
config>service>ipipe>sap

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

path

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

Context config>service>epipe>sap>eth-tunnel
config>service>ipipe>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**}]
no mep *mep-id* **domain** *md-index* **association** *ma-index* **primary-valn-enable**
[**vlan** *vlan-id*]

Context config>service>epipe>sap>eth-cfm
config>service>ies>sub-if>eth-cfm
config>service>epipe>spoke-sdp>eth-cfm
config>service>ipipe>sap>eth-cfm

Description This command provisions the maintenance endpoint (MEP).

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

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

Values 1 — 81921

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. The UP direction is not supported for all Fpipe services. For example, Ipipe does not support the direction of UP for MEPs.

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

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

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

ccm-enable

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

Context config>service>epipe>spoke-sdp>eth-cfm>mep
config>service>epipe>sap>eth-cfm>mep
config>service>ipipe>sap>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>epipe>spoke-sdp>eth-cfm>mep
config>service>epipe>sap>eth-cfm>mep
config>service>ipipe>sap>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

ccm-padding-size

Syntax **ccm-padding-size** *ccm-padding*
no ccm-padding-size *ccm-padding*

Context config>service>epipe>sap>eth-cfm>mep
config>service>ipipe>sap>eth-cfm>mep
config>service>epipe>sdp> 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
config>service>vpls>sap>eth-cfm>mep
config>service>vpls>spoke-sdp>eth-cfm>mep
config>service>vpls>mesh-sdp>eth-cfm>mep
config>service>ies>if>sap>eth-cfm>mep>
config>service>ies>if>spoke-sdp>eth-cfm>mep
config>service>ies>sub-if>grp-if>sap>eth-cfm>mep
config>service>vprn>if>sap>eth-cfm>mep
config>service>vprn>if>spoke-sdp>eth-cfm>mep

config>service>vprn>sub-if>grp-if>sap>eth-cfm>mep
config>port>ethernet>eth-cfm>mep
config>lag>eth-cfm>eth-cfm>mep
config>router>if>eth-cfm>mep

Description Set the byte size of the optional Data TLV to be included in the ETH-CC PDU. This will increase the size of the ETH-CC PDU by the configured value. The base size of the ETH-CC PDU, including the Interface Status TLV and Port Status TLV, is 83 bytes not including the Layer Two encapsulation. CCM padding is not supported when the CCM-Interval is less than one second.

Default [no] ccm-padding-size

Parameters *ccm-padding* — specifies the byte size of the Optional Data TLV

Values 3 — 1500

ccm-tlv-ignore

Syntax **ccm-tlv-ignore** [**interface-status**][**port-status**]
no ccm-tlv-ignore

Context config>port>ethernet>eth-cfm>mep
config>lag>eth-cfm>mep
config>router>interface>eth-cfm>mep

Description This command allows the receiving MEP to ignore the specified TLVs in CCM PDU. Ignored TLVs will be reported as absent and will have no impact on the MEP state machine.

The **no** form of the command means the receiving MEP will process all recognized TLVs in the CCM PDU.

Default no ccm-tlv-ignore

Parameters **interface-status —** ignores the interface status TLV on reception.

port-status — ignores the port status TVL on reception.

eth-test-enable

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

Context config>service>epipe>spoke-sdp>eth-cfm>mep
config>service>epipe>sap>eth-cfm>mep
config>service>ipipe>sap>eth-cfm>mep

Description For this 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 performed 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 indicates the problem.

bit-error-threshold

Syntax **bit-error-threshold** *errors*
no bit-error-threshold

Context config>service>epipe>sap>eth-cfm>mep>eth-test-enable

Description This command is used to specify the threshold value of bit errors.

test-pattern

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

Context config>service>epipe>spoke-sdp>eth-cfm>mep>eth-test-enable
config>service>epipe>sap>eth-cfm>mep>eth-test-enable
config>service>ipipe>sap>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.

Default all-zeros

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.

fault-propagation-enable

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

Context config>service>epipe>sap>eth-cfm>mep
config>service>epipe>spoke-sdp>eth-cfm>mep
config>service>ipipe>sap>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>epipe>spoke-sdp>eth-cfm>mep
config>service>epipe>sap>eth-cfm>mep
cconfig>service>ipipe>sap>eth-cfm>mep

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>epipe>spoke-sdp>eth-cfm>mep
config>service>epipe>sap>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 SDP).

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

Values 6-byte unicast mac-address (xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx) of the MEP. 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

mip

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

Context config>service>epipe>sap>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 —** provides a method for manually configuring the MIP MAC.

mac-address — Specifies the MAC address of the MIP.

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

tunnel-fault

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

Context config>service>epipe>eth-cfm
config>service>epipe>sap>eth-cfm
config>service>ipipe>eth-cfm
config>service>ipipe>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 **ais-enable** command under **config>service>epipe>sap>eth-cfm>ais-enable** context 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)

Service Filter and QoS Policy Commands

egress

Syntax **egress**

Context config>service>apipe>sap
config>service>cpipe>sap
config>service>cpipe>spoke-sdp
config>service>epipe>spoke-sdp
config>service>fpipe>sap
config>service>ipipe>sap
config>service>epipe>sap

Description This command enables the context to configure egress SAP parameters.

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

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

Description This command forces vc-vlan-type forwarding in the data path for spoke and mesh SDPs which have either vc-type. This comand is not allowed on vlan-vc-type SDPs.

The **no** version of this command sets default behavior.

Default Per default this feature is disabled

ingress

Syntax **ingress**

Context config>service>apipe>sap
config>service>cpipe>sap
config>service>cpipe>spoke-sdp
config>service>epipe>spoke-sdp
config>service>fpipe>sap
config>service>ipipe>sap
config>service>epipe>sap
config>service>epipe>sap

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

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

filter

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

Context config>service>epipe>sap>egress
config>service>epipe>sap>ingress
config>service>epipe>spoke-sdp>egress
config>service>epipe>spoke-sdp>ingress
config>service>ipipe>spoke-sdp>egress
config>service>ipipe>sap>ingress
config>service>ipipe>sap>egress
config>service>ipipe>spoke-sdp>ingress
config>service>epipe>sap>egress
config>service>epipe>sap>ingress

Description This command associates an IP 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 matching criteria. Only one filter can 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.

IP filters apply only to RFC 2427-routed IP packets. Frames that do not contain IP packets will not be subject to the filter and will always be passed, even if the filter's default action is to drop.

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**.

Note that IPv6 filters are not supported on a Layer 2 SAP that is configured with QoS MAC criteria. Also, MAC filters are not supported on a Layer 2 SAP that is configured with QoS IPv6 criteria.

Special Cases **Epipe —** Both MAC and IP filters are supported on an Epipe 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>epipe>sap>egress
config>service>ipipe>sap>egress

Description This command configures HSMDA egress and ingress queue overrides.

packet-byte-offset

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

Context config>service>epipe>sap>egress>hsmda-queue-over
config>service>ipipe>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). For example, this 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 20 bytes may be added to the packet and up to 43 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 mentioned above, the variable accounting size offered by the packet-byte-offset command is targeted at the queue and queue group level. When the queue group represents the last-mile bandwidth constraints for a subscriber, the offset allows the HSMDA queue group to provide an accurate accounting to prevent overrun and underrun conditions for the subscriber. The accounting size of the packet is ignored by the secondary shapers, the scheduling priority level shapers and the scheduler maximum rate. The actual on-the-wire frame size is used for these functions to allow an accurate representation of the behavior of the subscriber's packets on an Ethernet aggregation network.

The packet-byte-offset value can be overridden for the HSMDA queue at the SAP or subscriber profile level.

The **no** form of the command removes any accounting size changes to packets handled by the queue. The command does not effect overrides that may exist on SAPs or subscriber profiles associated with the queue.

Parameters **add** *add-bytes* **—** The **add** keyword is mutually exclusive with the subtract keyword. Either the add or subtract keyword must be specified. The add keyword is used to indicate that the following byte value should be added to the packet for queue and queue group level accounting functions. The corresponding byte value must be specified when executing the packet-byte-offset command.

Values 0 — 31

subtract *sub-bytes* **—** The **subtract** keyword is mutually exclusive with the add keyword. Either the add or subtract keyword must be specified. The subtract keyword is used to indicate that the following byte value should be subtracted from the packet for queue and queue group level accounting functions. The corresponding byte value must be specified when executing the packet-byte-offset command.

Values 1 — 32

queue

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

Context config>service>epipe>sap>egress>hsmda-queue-over
config>service>ipipe>sap>egress>hsmda-queue-over

Description This command, within the QoS policy hsmda-queue context, is a container for the configuration parameters controlling the behavior of an HSMDA queue. Unlike the standard QoS policy queue command, this command is not used to actually create or dynamically assign the queue to the object

which the policy is applied. The queue identified by queue-id always exists on the SAP or subscriber context whether the command is executed or not. In the case of HSMDA SAPs and subscribers, all eight queues exist at the moment the system allocates an HSMDA queue group to the object (both ingress and egress).

Best-Effort, Expedited and Auto-Expedite Queue Behavior Based on Queue-ID

With standard service queues, the scheduling behavior relative to other queues is based on two items, the queues Best-Effort or Expedited nature and the dynamic rate of the queue relative to the defined CIR. HSMDA queues are handled differently. The create time auto-expedite and explicit expedite and best-effort qualifiers have been eliminated and instead the scheduling behavior is based solely on the queues identifier. Queues with a queue-id equal to 1 are placed in scheduling class 1. Queues with queue-id 2 are placed in scheduling class 2. And so on up to scheduling class 8. Each scheduling class is either mapped directly to a strict scheduling priority level based on the class ID, or the class may be placed into a weighted scheduling class group providing byte fair weighted round robin scheduling between the members of the group. Two weighted groups are supported and each may contain up to three consecutive scheduling classes. The weighed group assumes its highest member classís inherent strict scheduling level for scheduling purposes. Strict priority level 8 has the highest priority while strict level 1 has the lowest. When grouping of scheduling classes is defined, some of the strict levels will not be in use.

Single Type of HSMDA Queues

Another difference between HSMDA queues and standard service queues is the lack of Multipoint queues. At ingress, an HSMDA SAP or subscriber does not require Multipoint queues since all forwarding types (broadcast, multicast, unicast and unknown) forward to a single destination ñ the ingress forwarding plane on the IOM. Instead of a possible eight queues per forwarding type (for a total of up to 32) within the SAP ingress QoS policy, the hsmda-queues node supports a maximum of eight queues.

Every HSMDA Queue Supports Profile Mode Implicitly

Unlike standard service queues, the HSMDA queues do not need to be placed into the special mode profile at create time in order to support ingress color aware policing. Each queue may handle in-profile, out-of-profile and profile undefined packets simultaneously. As with standard queues, the explicit profile of a packet is dependant on ingress sub-forwarding class to which the packet is mapped.

The **no** form of the command restores the defined queue-id to its default parameters. All HSMDA queues having the queue-id and associated with the QoS policy are re-initialized to default parameters.

Parameters *queue-id* — Specifies the HSMDA queue to use for packets in this forwarding class. This mapping is used when the SAP is on a HSMDA MDA.

Values 1 — 8

rate

Syntax **rate** *pir-rate*
no rate

Context config>service>epipe>sap>egress>hsmda-queue-over
config>service>ipipe>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>epipe>sap>egress>hsmda-queue-over>queue
config>service>ipipe>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>epipe>sap>egress>hsmda-queue-over
config>service>ipipe>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.

slope-policy

Syntax **slope-policy** *hsmda-slope-policy-name*
no slope-policy

Context config>service>epipe>sap>egress>hsmda-queue-over
config>service>ipipe>sap>egress>hsmda-queue-over

Description This command assigns an HSMDA slope policy to the SAP. The policy may be assigned to an ingress or egress HSMDA queue. The policy contains the Maximum Buffer Size (MBS) that will be applied to the queue and the high and low priority RED slope definitions. The function of the MBS and RED slopes is to provide congestion control for an HSMDA queue. The MBS parameter defines the

maximum depth a queue may reach when accepting packets. The low and high priority RED slopes provides for random early detection of congestion and slope based discards based on queue depth.

An HSMDA slope policy can be applied to queues defined in the SAP ingress and SAP egress QoS policy HSMDA queues context. Once an HSMDA slope policy is applied to a SAP QoS policy queue, it cannot be deleted. Any edits to the policy are updated to all HSMDA queues indirectly associated with the policy.

Default HSMDA Slope Policy

An HSMDA slope policy named "default" always exists on the system and does not need to be created. The default policy is automatically applied to all HSMDA queues unless another HSMDA slope policy is specified for the queue. The default policy cannot be modified or deleted. Attempting to execute the **no hsmda-slope-policy** *default* command results in an error.

The **no** form of the command removes the specified HSMDA slope policy from the configuration. If the HSMDA slope policy is currently associated with an HSMDA queue, the command will fail.

Parameters *hsmda-slope-policy-name* — Specifies a HSMDA slope policy up to 32 characters in length. The HSMDA slope policy must be exist prior to applying the policy name to an HSMDA queue.

secondary-shaper

Syntax **secondary-shaper** *secondary-shaper-name*
no secondary-shaper

Context config>service>epipe>sap>egress>hsmda-queue-over
config>service>ipipe>sap>egress>hsmda-queue-over

Description This command configures an HSMDA egress secondary shaper.

Parameters *secondary-shaper-name* — Specifies a secondary shaper name up to 32 characters in length.

filter

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

Context config>service>fpipe>sap>egress
config>service>fpipe>sap>ingress
config>service>cpipe>spoke-sdp>egress
config>service>cpipe>spoke-sdp>ingress
config>service>fpipe>spoke-sdp>egress
config>service>fpipe>spoke-sdp>ingress
config>service>ipipe>spoke-sdp>egress
config>service>ipipe>sap>ingress
config>service>ipipe>sap>egress
config>service>ipipe>spoke-sdp>ingress

Description This command associates a 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 matching criteria. Only one filter can be applied to a SAP at a time.

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

IP filters apply only to RFC 2427-routed IP packets. Frames that do not contain IP packets will not be subject to the filter and will always be passed, even if the filter's default action is to drop.

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

Parameters **ip** *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

qos

Syntax **qos** *policy-id* [**shared-queuing**] [**fp-redirect-group** *queue-group-name* **instance** *instance-id*]
no qos

Context config>service>apipe>sap>ingress
config>service>fpipe>sap>ingress
config>service>ipipe>sap>ingress
config>service>epipe>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. If the policy-id does not exist, an error will be returned.

The **qos** command, when used under the ingress context, is used to associate ingress 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.

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 *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. The shared-queueing keyword specifies the shared queue policy will be 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 specifies that this queue-id is for multipoint forwarded traffic only. This queue-id can only be explicitly mapped to the forwarding class multicast, broadcast, or unknown unicast ingress traffic. Attempting to map forwarding class unicast traffic to a multipoint queue generates an error; no changes are made to the current unicast traffic queue mapping.

A queue must be created as multipoint. The **multipoint** designator cannot be defined after the queue is created. If an attempt is made to modify the command to include the **multipoint** keyword, an error is generated and the command will not execute.

The **multipoint** keyword can be entered in the command line on a pre-existing multipoint queue to edit queue-id parameters.

Default Present (the queue is created as non-multipoint).

Values Multipoint or not present.

fp-redirect-group — This keyword can only be used on SAP ingress and associates a SAP ingress with 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.

queue-group-name — Specifies the name of the queue group to be instance 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 forwarding plane queue group, created under *config>card>fp>ingress>access*.

instance *instance-id* **—** Specifies the instance of the named queue group on the IOM/IMM/XMA ingress forwarding plane.

qos

Syntax **qos** *policy-id* [**port-redirect-group** *queue-group-name* **instance** *instance-id*]
no qos

Context config>service>apipe>sap>egress
config>service>cpipe>sap>egress
config>service>fpipe>sap>egress
config>service>ipipe>sap>egress
config>service>epipe>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, when used under the egress context, is used to associate 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.

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 *policy-id* — The egress policy ID to associate with SAP on egress. The policy ID must already exist.

Values 1 — 65535

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>apipe>sap>egress
config>service>apipe>sap>ingress
config>service>cpipe>sap>egress
config>service>cpipe>sap>ingress
config>service>fpipe>sap>egress
config>service>fpipe>sap>ingress
config>service>ipipe>sap>egress
config>service>ipipe>sap>ingress
config>service>epipe>sap>egress
config>service>epipe>sap>ingress

Description This command enables the context to configure override values for the specified SAP egress or ingress QoS queue. These values override the corresponding ones specified in the associated SAP egress or ingress QoS policy. If the policy was created as a template policy, this command overrides the parameter and its description and queue parameters in the policy.

queue

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

Context config>service>apipe>sap>egress>queue-override
config>service>apipe>sap>ingress>queue-override
config>service>cpipe>sap>egress>queue-override
config>service>cpipe>sap>ingress>queue-override
config>service>fpipe>sap>egress>queue-override
config>service>fpipe>sap>ingress>queue-override
config>service>ipipe>sap>egress>queue-override
config>service>ipipe>sap>ingress>queue-override
config>service>epipe>sap>egress>queue-override
config>service>epipe>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** *adaptation-rule*}] [**cir** *adaptation-rule*}]
no adaptation-rule

Context config>service>apipe>sap>egress>queue-override>queue
config>service>apipe>sap>ingress>queue-override>queue
config>service>fpipe>sap>egress>queue-override>queue
config>service>fpipe>sap>ingress>queue-override>queue
config>service>ipipe>sap>egress>queue-override>queue
config>service>ipipe>sap>ingress>queue-override>queue
config>service>epipe>sap>egress>queue-override>queue
config>service>epipe>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>apipe>sap>egress>queue-override>queue
config>service>cpipe>sap>egress>queue-override>queue
config>service>epipe>sap>egress>queue-override>queue

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

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

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

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

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

  As a special case, when a queue or associated intermediate scheduler is configured with a CIR-weight equal to 0, the system automatically sets the queue's frame based within-cir offered-load to 0, preventing it from receiving bandwidth during the port scheduler's within-cir pass.
- Frame based PIR — The frame based PIR is calculated by multiplying the packet to frame factor with the queue's configured PIR and then adding the result to that PIR. If the queue PIR is set to 7500 octets and the packet to frame factor equals 0.1, the frame based PIR would be 7500 x 1.1 or 8250 octets.
- Frame based within-pir offered-load — The frame based within-pir offered-load is the portion of the frame based offered-load considered to be within the frame based PIR. The frame based within-pir offered-load is the lesser of the frame based offered-load and the frame based PIR. If the frame based offered-load equaled 11000 octets and the frame based PIR equaled 8250 octets, the frame based within-pir offered-load would be limited to 8250 octets. If the frame based offered-load equaled 7000 octets and the frame based PIR equaled 8250 octets, the frame based within-pir offered load would equal 7000 octets.

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

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

The **no** form of this command restores the average frame overhead parameter for the queue to the default value of 0 percent. When set to 0, the system uses the packet based queue statistics for calculating port scheduler priority bandwidth allocation. If the no avg-frame-overhead command is

executed in a queue-override queue id context, the avg-frame-overhead setting for the queue within the sap-egress QoS policy takes effect.

Default 0

Parameters *percent* — This parameter sets the average amount of packet-to-frame encapsulation overhead expected for the queue. This value is not used by the system for egress Ethernet queues.

Values 0.00 — 100.00

burst-limit

Syntax **burst-limit** {**default** | *size* [**byte** | **kilobyte**]}
no burst-limit

Context config>service>apipe>sap>egress>queue-override>queue
config>service>apipe>sap>ingress>queue-override>queue
config>service>cpipe>sap>egress>queue-override>queue
config>service>cpipe>sap>ingress>queue-override>queue
config>service>fpipe>sap>egress>queue-override>queue
config>service>fpipe>sap>ingress>queue-override>queue
config>service>ipipe>sap>egress>queue-override>queue
config>service>ipipe>sap>ingress>queue-override>queue
config>service>epipe>sap>egress>queue-override>queue
config>service>epipe>sap>ingress>queue-override>queue

Description The `queue burst-limit` command is used to define an explicit shaping burst size for a queue. The configured size defines the shaping leaky bucket threshold level that indicates the maximum burst over the queue's shaping rate.

The `burst-limit` command is supported under the sap-ingress and sap-egress QoS policy queues. The command is also supported under the ingress and egress queue-group-templates queues.

The **no** form of this command is used to restore the default burst limit to the specified queue. This is equivalent to specifying burst-limit default within the QoS policies or queue group templates. When specified within a queue-override queue context, any current burst limit override for the queue will be removed and the queue's burst limit will be controlled by its defining policy or template.

Parameters **default —** The default parameter is mutually exclusive to specifying an explicit size value. When burst-limit default is executed, the queue is returned to the system default value.

size — When a numeric value is specified (size), the system interprets the value as an explicit burst limit size. The value is expressed as an integer and by default is interpreted as the burst limit in Kilobytes. If the value is intended to be interpreted in bytes, the byte qualifier must be added following size.

Values 1 to 14,000 (14,000 or 14,000,000 depending on bytes or kilobytes)

Default No default for size, use the default keyword to specify default burst limit

byte — The **bytes** qualifier is used to specify that the value given for size must be interpreted as the burst limit in bytes. The byte qualifier is optional and mutually exclusive with the kilobytes qualifier.

kilobyte — The **kilobyte** qualifier is used to specify that the value given for size must be interpreted as the burst limit in Kilobytes. The kilobyte qualifier is optional and mutually exclusive with the bytes qualifier. If neither bytes nor kilobytes is specified, the default qualifier is kilobytes.

cbs

Syntax **cbs** *size-in-kbytes*
no cbs

Context config>service>apipe>sap>egress>queue-override>queue
config>service>apipe>sap>ingress>queue-override>queue
config>service>cpipe>sap>egress>queue-override>queue
config>service>cpipe>sap>ingress>queue-override>queue
config>service>fpipe>sap>egress>queue-override>queue
config>service>fpipe>sap>ingress>queue-override>queue
config>service>ipipe>sap>egress>queue-override>queue
config>service>ipipe>sap>ingress>queue-override>queue
config>service>epipe>sap>egress>queue-override>queue
config>service>epipe>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 to 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, default

high-prio-only

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

Context config>service>apipe>sap>egress>queue-override>queue
config>service>apipe>sap>ingress>queue-override>queue
config>service>cpipe>sap>egress>queue-override>queue
config>service>cpipe>sap>ingress>queue-override>queue
config>service>fpipe>sap>egress>queue-override>queue
config>service>fpipe>sap>ingress>queue-override>queue
config>service>ipipe>sap>egress>queue-override>queue
config>service>ipipe>sap>ingress>queue-override>queue
config>service>epipe>sap>egress>queue-override>queue
config>service>epipe>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* [**bytes**|**kilobytes**]
no mbs

Context config>service>ipipe>sap>ingress>hsmda-queue-override>queue
config>service>epipe>sap>ingress>hsmda-queue-override>queue
config>service>apipe>sap>ingress>queue-override>queue
config>service>cpipe>sap>ingress>queue-override>queue
config>service>fpipe>sap>ingress>queue-override>queue
config>service>ipipe>sap>ingress>queue-override>queue
config>service>epipe>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 packet's RED slope will not force the discard of the packet. Setting proper CBS parameters and controlling CBS oversubscription is one major safeguard to queue starvation (when a queue does not receive its fair share of buffers). Another is properly setting the RED slope parameters for the needs of services on this port or channel. If the CBS value is larger than the MBS value, an error will occur, preventing the MBS change.

The defined high-prio-only value cannot be greater than the MBS size of the queue. Attempting to change the MBS to a value smaller than the high priority reserve will generate an error and fail execution. Attempting to set the high-prio-only value larger than the current MBS size will also result in an error and fail execution.

The **no** form of this command returns the MBS size 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

parent

Syntax **parent** {[**weight** *weight*] [**cir-weight** *cir-weight*]}
no parent

Context config>service>epipe>sap>egress>queue-override>queue
config>service>epipe>sap>ingress>queue-override>queue
config>service>apipe>sap>egress>queue-override>queue
config>service>apipe>sap>ingress>queue-override>queue

Description This command defines an optional parent scheduler that further governs the available bandwidth given the queue aside from the queue's PIR setting. When multiple schedulers and/or queues share a child status with the parent scheduler, the **weight** or **level** parameters define how this queue contends with the other children for the parent's bandwidth.

Checks are not performed to see if a *scheduler-name* exists when the parent command is defined on the queue. Scheduler names are configured in the **config>qos>scheduler-policy>tier** *level* context. Multiple schedulers can exist with the *scheduler-name* and the association pertains to a scheduler that should exist on the egress SAP as the policy is applied and the queue created. When the queue is created on the egress SAP, the existence of the *scheduler-name* is dependent on a scheduler policy containing the *scheduler-name* being directly or indirectly applied (through a multi-service customer site) to the egress SAP. If the *scheduler-name* does not exist, the queue is placed in the orphaned operational state. The queue will accept packets but will not be bandwidth limited by a virtual scheduler or the scheduler hierarchy applied to the SAP. The orphaned state must generate a log entry and a trap message. The SAP which the queue belongs to must also depict an orphan queue status. The orphaned state of the queue is automatically cleared when the *scheduler-name* becomes available on the egress SAP.

The parent scheduler can be made unavailable due to the removal of a scheduler policy or scheduler. When an existing parent scheduler is removed or inoperative, the queue enters the orphaned state mentioned above and automatically return to normal operation when the parent scheduler is available

again.

When a parent scheduler is defined without specifying weight or strict parameters, the default bandwidth access method is weight with a value of 1.

The **no** form of the command removes a child association with a parent scheduler. If a parent association does not currently exist, the command has no effect and returns without an error. Once a parent association has been removed, the former child queue attempts to operate based on its configured rate parameter. Removing the parent association on the queue within the policy takes effect immediately on all queues using the SAP egress QoS policy.

Parameters **weight** *weight* — These optional keywords are mutually exclusive to the keyword **level**. *weight* defines the relative weight of this queue in comparison to other child schedulers and queues while vying for bandwidth on the parent *scheduler-name*. Any queues or schedulers defined as weighted receive no parental bandwidth until all strict queues and schedulers on the parent have reached their maximum bandwidth or are idle. In this manner, weighted children are considered to be the lowest priority.

All **weight** values from all weighted active queues and schedulers with a common parent scheduler are added together. Then, each individual active weight is divided by the total, deriving the percentage of remaining bandwidth provided to the queue or scheduler after the strict children are serviced. A weight is considered to be active when the pertaining queue or scheduler has not reached its maximum rate and still has packets to transmit. All child queues and schedulers with a weight of 0 are considered to have the lowest priority level and are not serviced until all strict and non-zero weighted queues and schedulers are operating at the maximum bandwidth or are idle.

Values 0 — 100

Default 1

cir-weight ***cir-weight*** **—** Defines the weight the queue or scheduler will use at the within-cir port priority level (defined by the cir-level parameter). The weight is specified as an integer value from 0 to 100 with 100 being the highest weight. When the cir-weight parameter is set to a value of 0 (the default value), the queue or scheduler does not receive bandwidth during the port schedulers within-cir pass and the cir-level parameter is ignored. If the cir-weight parameter is 1 or greater, the cir-level parameter comes into play.

Values 0 — 100

percent-rate

Syntax **percent-rate** *pir-percent* [**cir** *cir-percent*]

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

Description The **percent-rate** command supports a queue's shaping rate and CIR rate as a percentage of the egress port's line rate. When the rates are expressed as a percentage within the template, the actual rate used per instance of the queue group queue-id will vary based on the port speed. For example, when the same template is used to create a queue group on a 1-Gigabit and a 10-Gigabit Ethernet port, the queue's rates will be 10 times greater on the 10 Gigabit port due to the difference in port speeds. This enables the same template to be used on multiple ports without needing to use port based queue overrides to modify a queue's rate to get the same relative performance from the queue.

If the port's speed changes after the queue is created, the queue's shaping and CIR rates will be recalculated based on the defined percentage value.

The rate and percent-rate commands override one another. If the current rate for a queue is defined using the percent-rate command and the rate command is executed, the percent-rate values are deleted. In a similar fashion, the percent-rate command causes any rate command values to be deleted. A queue's rate may dynamically be changed back and forth from a percentage to an explicit rate at anytime.

An egress port queue group queue rate override may be expressed as either a percentage or an explicit rate independent on how the queue's template rate is expressed.

The **no** form of this command returns the queue to its default shaping rate and cir rate. When **no percent-rate** is defined within a port egress queue group queue override, the queue reverts to the defined shaping and CIR rates within the egress queue group template associated with the queue.

Parameters *pir-percent* — The *percent-of-line-rate* parameter is used to express the queue's shaping rate as a percentage of line rate. The line rate associated with the queue's port may dynamically change due to configuration or auto-negotiation. The line rate may also be affected by an egress port scheduler defined max-rate.

Values Percentage ranging from 0.01 to 100.00. The default is 100.00.

cir *cir-percent* **—** The **cir** keyword is optional and when defined the required *percent-of-line-rate* CIR parameter expresses the queue's committed scheduling rate as a percentage of line rate. The line rate associated with the queue's port may dynamically change due to configuration or auto-negotiation. The line rate may also be affected by an egress port scheduler defined max-rate.

rate

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

Context config>service>apipe>sap>egress>queue-override>queue
config>service>apipe>sap>ingress>queue-override>queue
config>service>cpipe>sap>egress>queue-override>queue
config>service>cpipe>sap>ingress>queue-override>queue
config>service>fpipe>sap>egress>queue-override>queue
config>service>fpipe>sap>ingress>queue-override>queue
config>service>ipipe>sap>egress>queue-override>queue
config>service>ipipe>sap>ingress>queue-override>queue
config>service>epipe>sap>egress>queue-override>queue
config>service>epipe>sap>ingress>queue-override>queue

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

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

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

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

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

Values 0 — 100000000, **max, sum**

Default 0

scheduler-override

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

Context config>service>apipe>sap>egress
config>service>apipe>sap>ingress
config>service>cpipe>sap>egress
config>service>cpipe>sap>ingress
config>service>fpipe>sap>egress
config>service>fpipe>sap>ingress
config>service>ipipe>sap>egress
config>service>ipipe>sap>ingress
config>service>epipe>sap>egress
config>service>epipe>sap>ingress

Description This command specifies the set of attributes whose values have been overridden by 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 [**no**] **scheduler** *scheduler-name*

Context config>service>apipe>sap>egress>sched-override
config>service>apipe>sap>ingress>sched-override
config>service>cpipe>sap>egress>sched-override
config>service>cpipe>sap>ingress>sched-override
config>service>fpipe>sap>egress>sched-override
config>service>fpipe>sap>ingress>sched-override
config>service>ipipe>sap>egress>sched-override
config>service>ipipe>sap>ingress>sched-override
config>service>epipe>sap>egress>sched-override
config>service>epipe>sap>ingress>sched-override

Description This command can be used to override specific attributes of the specified scheduler name. A scheduler defines bandwidth controls that limit each child (other schedulers and queues) associated with the scheduler. Scheduler objects are created within the hierarchical tiers of the policy. It is assumed that each scheduler created will have queues or other schedulers defined as child associations. The scheduler can be a child (take bandwidth from a scheduler in a higher tier, except for schedulers created in tier 1). 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>apipe>sap>egress>sched-override>scheduler
config>service>apipe>sap>ingress>sched-override>scheduler
config>service>cpipe>sap>egress>sched-override>scheduler
config>service>cpipe>sap>ingress>sched-override>scheduler
config>service>fpipe>sap>egress>sched-override>scheduler
config>service>fpipe>sap>ingress>sched-override>scheduler
config>service>ipipe>sap>egress>sched-override>scheduler
config>service>ipipe>sap>ingress>sched-override>scheduler
config>service>epipe>sap>egress>sched-override>scheduler
config>service>epipe>sap>ingress>sched-override>scheduler

Description This command can be used to override specific attributes of the specified scheduler rate. The **rate** command defines the maximum bandwidth that the scheduler can offer its child queues or schedulers. The maximum rate is limited to the amount of bandwidth the scheduler can receive from its parent scheduler. If the scheduler has no parent, the maximum rate is assumed to be the amount available to the scheduler. When a parent is associated with the scheduler, the CIR parameter provides the amount of bandwidth to be considered during the parent scheduler's 'within CIR' distribution phase.

The actual operating rate of the scheduler is limited by bandwidth constraints other than its maximum rate. The scheduler's parent scheduler may not have the available bandwidth to meet the scheduler's needs or the bandwidth available to the parent scheduler could be allocated to other child schedulers or child queues on the parent based on higher priority. The children of the scheduler may not need the maximum rate available to the scheduler due to insufficient offered load or limits to their own maximum rates.

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

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

Parameters *pir-rate* — The **pir** parameter accepts a step multiplier value that specifies the multiplier used to determine the PIR rate at which the queue will operate. A value of 0 to 100000000 or the keyword **max** or **sum** 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 to 250 or the keyword max is accepted. Any other value will result in an error without modifying the current CIR rate.

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

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

Values 0 — 10000000, max, sum

Default sum

scheduler-policy

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

Context config>service>apipe>sap>ingress
config>service>apipe>sap>egress
config>service>cpipe>sap>ingress
config>service>cpipe>sap>egress
config>service>fpipe>sap>ingress
config>service>fpipe>sap>egress
config>service>ipipe>sap>ingress
config>service>ipipe>sap>egress
config>service>epipe>sap>ingress
config>service>epipe>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.

vlan-translation

Syntax **vlan-translation** {*vlan-id* | **copy-outer**}
no vlan-translation

Context config>service>epipe>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 "copy-outer" keyword, the outer vlan-id will be copied to inner position on QinQ encapsulated ports. The feature is not supported on default-dot1q saps (1/1/1:* and 1/1/1:0), nor on TopQ saps.

The **no** version of the 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 that the preserved vlan-id will be overwritten with this value.

Values 0 — 4094

outer-copy — Keyword specifies to use the outer VLAN ID.

match-qinq-dot1p

Syntax **match-qinq-dot1p** {**top** | **bottom**}
no match-qinq-dot1p de

Context config>service>ipipe>sap>ingress
config>service>epipe>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
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

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	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
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

Egress SAP Type	Ingress Packet Preserved Dot1P State	Marked (or Remarked) PBits
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**.

VLL Frame Relay Commands

frame-relay

Syntax **frame-relay**

Context config>service>apipe>sap
config>service>fpipe>sap
config>service>ipipe>sap
config>service>epipe>sap

Description This command enables the context to configure Frame Relay parameters.

frf-12

Syntax [**no**] **frf-12**

Context config>service>fpipe>sap>frame-relay
config>service>ipipe>sap>frame-relay
config>service>epipe>sap>frame-relay

Description This command enables the use of FRF12 headers.

The **no** form of the command disables the use of FRF12 headers.

ete-fragment-threshold

Syntax **ete-fragment-threshold** *threshold*
no ete-fragment-threshold

Context config>service>fpipe>sap>frame-relay>frf-12
config>service>ipipe>sap>frame-relay>frf-12
config>service>epipe>sap>frame-relay>frf-12

Description This command specifies the maximum length of a fragment to be transmitted.

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

Parameters *threshold* — The maximum length of a fragment to be transmitted.

Values 128 — 512

Default 0

interleave

Syntax [**no**] **interleave**

Context config>service>epipe>sap>frame-relay>frf.12
config>service>ipipe>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*

Description config>service>apipe>sap>frame-relay
config>service>fpipe>sap>frame-relay
config>service>ipipe>sap>frame-relay
config>service>epipe>sap>frame-relay

Description This command specifies the scheduling class to use for this SAP.

Parameters *class-id* — Specifies the scheduling class to use for this sap.

Values 0 — 3

Default 0

VLL SDP Commands

spoke-sdp

Syntax **spoke-sdp** *sdp-id*[**:***vc-id*] [**vc-type** {**ether** | **vlan**}] [**no-endpoint**]
spoke-sdp *sdp-id*[**:***vc-id*] [**vc-type** {**ether** | **vlan**}] **endpoint** *endpoint-name* [**icb**]
no spoke-sdp *sdp-id*[**:***vc-id*]

Context config>service>cpipe
config>service>epipe

Description This command binds a service to an existing Service Distribution Point (SDP). A spoke SDP is treated like the equivalent of a traditional bridge "port" where flooded traffic received on the spoke SDP is replicated on all other "ports" (other spoke and mesh SDPs or SAPs) and not transmitted on the port it was received.

The SDP has an operational state which determines the operational state of the SDP within the service. For example, if the SDP is administratively or operationally down, the SDP for the service will be down.

The SDP must already be defined in the **config>service>sdp** context in order to associate an SDP with an Epipe,VPLS, VPRN, VPRN service. If the **sdp** *sdp-id* is not already configured, an error message is generated. If the *sdp-id* does exist, a binding between that *sdp-id* and the service is created. SDPs must be explicitly associated and bound to a service. If an SDP is not bound to a service, no far-end 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 **Epipe —** At most, only one *sdp-id* can be bound to an Epipe service. Since an Epipe is a point-to-point service, it can have, at most, two end points. The two end points can be one SAP and one SDP or two SAPs. Vc-switching VLLs are an exception. If the VLL is a "vc-switching" VLL, then the two endpoints must both be SDPs.

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

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.
- The VC type value for a VPLS service is defined as 0x000B.

Values ethernet

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.

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.

no endpoint — Removes the association of a spoke SDP with an explicit endpoint name.

endpoint *endpoint-name* **—** Specifies the name of the service endpoint.

icb — Configures the spoke SDP as an inter-chassis backup SDP binding.

spoke-sdp

Syntax **spoke-sdp** *sdp-id*[**:***vc-id*] [**no-endpoint**]
spoke-sdp *sdp-id*[:*vc-id*] **endpoint** *endpoint-name* [**icb**]
no spoke-sdp *sdp-id*[**:***vc-id*]

Context config>service>apipe
config>service>cpipe
config>service>fpipe
config>service>ipipe

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 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 SR/ESS 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.

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

vc-id — The virtual circuit identifier.

Values 1 — 4294967295

no endpoint — Adds or removes a spoke SDP association.

endpoint *endpoint-name* **—** Specifies the name of the service endpoint.

icb — Configures the spoke SDP as an inter-chassis backup SDP binding.

hash-label

Syntax **hash-label** [**signal-capability**]
no hash-label

Context config>service>epipe>spoke-sdp
config>service>fpipe>spoke-sdp
config>service>ipipe>spoke-sdp
config>service>pw-template
config>service>vprn
config>service>vprn>interface>spoke-sdp
config>service>ies>interface>spoke-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 PW 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 PW 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 PW 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 PW 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 PW 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 PW 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.

cell-concatenation

Syntax **cell-concatenation**

Context config>service>apipe>spoke-sdp

Description This command enables the context to provide access to the various options that control the termination of ATM cell concatenation into an MPLS frame. Several options can be configured simultaneously. The concatenation process for a given MPLS packet ends when the first concatenation termination condition is met. The concatenation parameters apply only to ATM N:1 cell mode VLL.

aal5-frame-aware

Syntax [**no**] **aal5-frame-aware**

Context config>service>apipe>spoke-sdp>cell-concat

Description This command enables the configuration of AAL5 end-of-message (EOM) to be an indication to complete the cell concatenation operation.

The **no** form of the command resets the configuration to ignore the AAL5 EOM as an indication to complete the cell concatenation.

clp-change

Syntax [**no**] **clp-change**

Context config>service>apipe>spoke-sdp>cell-concat

Description This command enables the configuration of CLP change to be an indication to complete the cell concatenation operation.

The **no** form of the command resets the configuration to ignore the CLP change as an indication to complete the cell concatenation.

max-cells

Syntax **max-cells** *cell-count*
no max-cells [*cell-count*]

Context config>service>apipe>spoke-sdp>cell-concat

Description This command enables the configuration of the maximum number of ATM cells to accumulate into an MPLS packet. The remote peer will also signal the maximum number of concatenated cells it is willing to accept in an MPLS packet. When the lesser of (the configured value and the signaled value) number of cells is reached, the MPLS packet is queued for transmission onto the pseudowire. It is ensured that the MPLS packet MTU conforms to the configured service MTU.

The **no** form of this command sets max-cells to the value '1' indicating that no concatenation will be performed.

Parameters *cell-count* — Specify the maximum number of ATM cells to be accumulated into an MPLS packet before queueing the packet for transmission onto the pseudowire.

Values 1 — 128

Default 1

max-delay

Syntax **max-delay** *delay-time*

no max-delay [*delay-time*]

Context config>service>apipe>spoke-sdp>cell-concat

Description This command enables the configuration of the maximum amount of time to wait while performing ATM cell concatenation into an MPLS packet before transmitting the MPLS packet. This places an upper bound on the amount of delay introduced by the concatenation process. When this amount of time is reached from when the first ATM cell for this MPLS packet was received, the MPLS packet is queued for transmission onto the pseudowire.

The **no** form of this command resets max-delay to its default value.

Parameters *delay-time* — Specify the maximum amount of time, in hundreds of microseconds, to wait before transmitting the MPLS packet with whatever ATM cells have been received. For example, to bound the delay to 1 ms the user would configure 10 (hundreds of microseconds). The delay-time is rounded up to one of the following values 1, 5, 10, 50, 100, 200, 300 and 400.

Values 1 — 400

Default 400

control-word

Syntax [**no**] **control-word**

Context config>service>apipe>spoke-sdp
config>service>cpipe>spoke-sdp
config>service>epipe>spoke-sdp
config>service>fpipe>spoke-sdp
config>service>ipipe>spoke-sdp

Description The control word command provides the option to add a control word as part of the packet encapsulation for pseudowire types for which the control word is optional. These are Ethernet pseudowires (Epipe). ATM N:1 cell mode pseudowires (apipe vc-types atm-vcc and atm-vpc) and VT pseudowire (apipe vc-type atm-cell).
The configuration for the two directions of the pseudowire must match because the control word negotiation procedures described in Section 6.2 of RFC 4447 are not supported. The C-bit in the pseudowire FEC sent in the label mapping message is set to 1 when the control word is enabled. Otherwise, it is set to 0.

The service will only come up if the same C-bit value is signaled in both directions. If a spoke-sdp is configured to use the control word but the node receives a label mapping message with a C-bit clear, the node releases the label with the an “Illegal C-bit” status code as per Section 6.1 of RFC 4447. As soon as the user also enabled the control the remote peer, the remote peer will withdraw its original label and will send a label mapping with the C-bit set to 1 and the VLL service will be up in both nodes.

pw-path-id

Syntax [**no**] **pw-path-id**

Context config>service>epipe>spoke-sdp
config>service>cpipe>spoke-sdp
config>service>vpls>spoke-sdp
config>service>ies>spoke-sdp
config>service>vprn>spoke-sdp

Description This command enables the context to configure an MPLS-TP Pseudowire Path Identifier for a spoke-sdp. All elements of the PW path ID must be configured in order to enable a spoke-sdp with a PW path ID.

For an IES or VPRN spoke-sdp, the pw-path-id is only valid for ethernet spoke-sdps.

The **pw-path-id** only configurable if all of the following is true:

- The system is using network chassis mode D
- SDP signaling is off
- control-word is enabled (control-word is disabled by default)
- the service type is epipe, vpls, cpipe, or IES/VPRN interface
- mate SDP signaling is off for vc-switched services

The **no** form of the command deletes the PW path ID.

Default no pw-path-id

agi

Syntax **agi** *agi*
no agi

Context config>service>epipe>spoke-sdp>pw-path-id
config>service>cpipe>spoke-sdp>pw-path-id
config>service>vpls>spoke-sdp>pw-path-id
config>service>ies>spoke-sdp>pw-path-id
config>service>vprn>spoke-sdp>pw-path-id

Description This command configures the attachment group identifier for an MPLS-TP PW.

Parameters *agi* — Specifies the attachment group identifier.

Values 0 — 4294967295

saii-type2

Syntax **saii-type2** *global-id:node-id:ac-id*
no saii-type2

Context config>service>epipe>spoke-sdp>pw-path-id
config>service>cpipe>spoke-sdp>pw-path-id
config>service>vpls>spoke-sdp>pw-path-id
config>service>ies>spoke-sdp>pw-path-id
config>service>vprn>spoke-sdp>pw-path-id

Description This command configures the source individual attachement identifier (SAII) for an MPLS-TP spoke-sdp. If this is configured on a spoke-sdp for which vc-switching is also configured i.e. it is at an S-PE, then the values must match those of the taii-type2 of the mate spoke-sdp.

Parameters *global-id* — Specifies the global ID at the source PE or T-PE for the MPLS-TP PW for a spoke-SDP.

Values 0 — 4294967295

node-id — Specifies the node ID at the source PE or T-PE for the MPLS-TP PW for a spoke-SDP.

Values a.b.c.d or 0 — 4294967295

ac-id — Specifies the attachment circuit ID at the source PE or T-PE for the MPLS-TP PW for a spoke-SDP. If this node is the source of the PW, then the AC ID must be set to a locally unique value.

Values 1 — 4294967295

taii-type2

Syntax **taii-type2** *global-id:node-id:ac-id*
no taii-type2

Context config>service>epipe>spoke-sdp>pw-path-id
config>service>cpipe>spoke-sdp>pw-path-id
config>service>vpls>spoke-sdp>pw-path-id
config>service>ies>spoke-sdp>pw-path-id
config>service>vprn>spoke-sdp>pw-path-id

Description This command configures the source individual attachement identifier (SAII) for an MPLS-TP spoke-sdp. If this is configured on a spoke-sdp for which vc-switching is also configured i.e. it is at an S-PE, then the values must match those of the taii-type2 of the mate spoke-sdp.

Parameters *global-id* — Specifies the global ID at the target PE or T-PE for the MPLS-TP PW for a spoke-SDP.

Values 0 — 4294967295

node-id — Specifies the node ID at the target PE or T-PE for the MPLS-TP PW for a spoke-SDP.

Values a.b.c.d or 0 — 4294967295

ac-id — Specifies the attachment circuit ID at the target PE or T-PE for the MPLS-TP PW for a spoke-SDP. If this node is the source of the PW, then the AC ID must be set to a locally unique value.

Values 1 — 4294967295

control-channel-status

Syntax [**no**] **control-channel-status**

Context config>service>epipe>spoke-sdp
config>service>cpipe>spoke-sdp
config>service>vpls>spoke-sdp

config>service>ies>spoke-sdp
config>service>vprn>spoke-sdp

Description This command enables the configuration of static pseudowire status signaling on a spoke-sdp for which signaling for its SDP is set to OFF.

A control-channel-status no shutdown is allowed only if all of the following is true:

- The system is using network chassis mode D
- SDP signaling is off
- The control-word is enabled (control-word by default is disabled)
- The service type is epipe, apipe, vpls, cpipe, or IES/VPRN
- Mate sdp signaling is off (in vc-switched services)
- pw-path-id is configured for this spoke

The **no** form of this command renoves control channel status signaling form a spoke-sdp. It can only be removed if control channel status is shutdown.

Default no control-channel-status

refresh-timer

Syntax **refresh-timer** *value*
no refresh-timer

Context config>service>epipe>spoke-sdp>control-channel-status
config>service>cpipe>spoke-sdp>control-channel-status
config>service>vpls>spoke-sdp>control-channel-status
config>service>ies>spoke-sdp>control-channel-status
config>service>vprn>spoke-sdp>control-channel-status

Description This command configures the refresh timer for control channel status signaling packets. By default, no referesh packets are sent.

Default no refresh-timer

Parameters *value* — Specifies the refresh timer value.

Values 10 — 65535 seconds

Default 0 (off)

control-word

Syntax [**no**] **control-word**

Context config>service>ies>spoke-sdp
config>service>vprn>spoke-sdp

Description This command enables/disables the PW control word on spoke-sdps terminated on an IES or VPRN interface. The control word must be enabled to allow MPLS-TP OAM on the spoke-sdp

It is only valid for MPLS-TP spoke-sdps.

Default no control-word

egress

Syntax **egress**
config>service>apipe>spoke-sdp
config>service>cpipe>spoke-sdp
config>service>fpipe>spoke-sdp
config>service>ipipe>spoke-sdp

Description This command configures the egress SDP context.

hash-label

Syntax **hash-label** [**signal-capability**]
no hash label]

Context
config>service>epipe>spoke-sdp
config>service>fpipe>spoke-sdp
config>service>ipipe>spoke-sdp

Description This command enables the use of the hash label on a VLL, VPLS, or VPRN service bound to LDP or RSVP SDP as well as to a VPRN service using the autobind mode with the 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..

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 1 to indicate that.

In order to allow for applications whereby 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 that are generated in CPM and 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 **no** form of this command disables the use of the hash label.

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 PW 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 PW 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 PW 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 PW 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 PW 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.

ignore-oper-down

Syntax **ignore-oper-down**
[no] ignore-oper-down

Context config>service>epipe>sap>

Description ePipe service will not transition to Oper State: Down when a SAP fails and when this optional command configured under that specific SAP. Only a single SAP in an ePipe may have this optional command included.

Default no ignore-oper-down

ingress

Syntax **ingress**

Context config>service>fpipe>spoke-sdp
config>service>apipe>spoke-sdp
config>service>cpipe>spoke-sdp

Description This command configures the ingress SDP context.

filter

Syntax **filter** [**ip** *ip-filter-id*]
no filter

Context config>service>fpipe>spoke-sdp>egress
config>service>fpipe>spoke-sdp>ingress

Description This command associates an IP filter policy with an ingress or egress Service Distribution Point (SDP). Filter policies control the forwarding and dropping of packets based on IP matching criteria. Only one filter can be applied to a spoke SDP at a time.

The **filter** command is used to associate a filter policy with a specified *ip-filter-id* with an ingress or egress spoke SDP. The *ip-filter-id* must already be defined before the **filter** command is executed. If the filter policy does not exist, the operation will fail and an error message returned.
IP filters apply only to RFC 2427-routed IP packets. Frames that do not contain IP packets will not be subject to the filter and will always be passed, even if the filter's default action is to drop.

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

Parameters **ip —** Keyword indicating the filter policy is an IP filter.

ip-filter-id — The filter name acts as the ID for the IP filter policy. The filter ID must already exist within the created IP filters.

Values 1 — 65535

qos

Syntax **qos** *network-policy-id* **port-redirect-group** *queue-group-name* [**instance** *instance-id*]
no qos

Context config>service>apipe>spoke-sdp>egress
config>service>cpipe>spoke-sdp>egress
config>service>epipe>spoke-sdp>egress
config>service>fpipe>spoke-sdp>egress
config>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 PW packets to an egress port queue-group for the purpose of shaping.

The egress PW shaping provisioning model allows the mapping of one or more PWs to the same instance of queues, or policers and queues, that 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 that the PW packets can be forwarded on. 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 PWs to different queue-group templates.
4. Apply this network QoS policy to the egress context of a spoke-sdp inside a service, or to the egress context of a PW template and specify the redirect queue-group name.

One or more spoke-sdps 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 PW 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-sdp to the named queue-group. In such a case, the PW packet will be fed directly to the corresponding egress queue for that FC used by the IP network interface the PW packet is forwarded on. 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 PW packet.
2. When a PW 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-sdp to the named queue-group. In such a case, the PW packet will be fed directly to the corresponding egress queue for that FC used by the IP network interface the PW packet is forwarded on.
3. When a PW 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 that have network IP interfaces. The handling of this is dealt with in the data path as follows:
   – When a PW 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.

– When a PW packet for that FC is forwarded and an instance of the referenced queue-group name does not exist on that egress port, the PW 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 PW, any PW FC that 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 PW is redirected to exists and the redirection succeeds, the marking of the packet's DEI/dot1.p/DSCP and the tunnel's 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 PW. This is true regardless if an instance of the queue-group exists or not on the egress port the PW packet is forwarded to. If the packet's profile value changed due to egress child policer CIR profiling, the new profile value is used to mark the packet's DEI/dot1.p and the tunnel's DEI/dot1.p/EXP, but the DSCP is not modified by the policer's operation.

When the queue-group name the PW is redirected does not exist, the redirection command is failed. In this case, the marking of the packet's DEI/dot1.p/DSCP and the tunnel's 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 the PW packet is forwarded to.

The **no** version of this command removes the redirection of the PW 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

port-redirect-group *queue-group-name* **—** Specifies the name of the queue group template up to 32 characters in length.

instance *instance-id* **—** Specifies the optional identification of a specific instance of the queue-group.

Values 1—40960

qos

Syntax **qos** *network-policy-id* **fp-redirect-group** *queue-group-name* **instance** *instance-id*
no qos

Context config>service>apipe>spoke-sdp>ingress
config>service>cpipe>spoke-sdp>ingress
config>service>epipe>spoke-sdp>ingress
config>service>fpipe>spoke-sdp>ingress
config>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 PW packets to an ingress forwarding plane queue-group for the purpose of rate-limiting.

The ingress PW rate-limiting feature uses a policer in queue-group provisioning model. This model allows the mapping of one or more PWs to the same instance of policers that are defined in a queue-group template.

Operationally, the provisioning model in the case of the ingress PW 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 that the PW packets can be received on. 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 PWs 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 PW template and specify the redirect queue-group name.

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 PW rate-limiting feature:

1. When a PW 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 PW 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 PW 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-sdp to the named queue-group. In such a case, the PW 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 PW 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 that have network IP interfaces. The handling of this is dealt within the data path as follows:
   - When a PW 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”.
   - When a PW packet for that FC is received and an instance of the referenced queue-group name does not exist on that FP, the PW 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 PW, any PW FC that is not explicitly redirected in the network QoS policy will have the corresponding packets feed directly into

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 PW, then all packets of the PW will feed:
   - the ingress network shared queue for the packet's FC defined in the network-queue policy applied to the ingress of the MDA/FP. This is the default behavior.
   - 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 an IES/VPRN spoke interface and from an R-VPLS spoke-sdp that is forwarded to the R-VPLS IP interface. In these two cases, the ingress network shared queue for the packet's FC defined in the network-queue policy applied to the ingress of the MDA/FP is used.

When a PW 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 the default classification rule or the QoS filters defined in the ingress context of the network QoS policy applied to the PW. This is true regardless if an instance of the named policer queue-group exists on the ingress FP the PW packet is received on. 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's IP header if the user enabled the ler-use-dscp option and the PW terminates in IES or VPRN service (spoke-interface).

When the policer queue-group name the PW is redirected does not exist, the redirection command is failed. In this case, the packet classification is performed according to the default classification rule or the QoS filters defined in the ingress context of the network QoS policy applied to the network IP interface the PW packet is received on.

The **no** version of this command removes the redirection of the PW to the queue-group.

Parameters *network-policy-id* — Specifies the network policy indentifiction. 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.

instance *instance-id* — Specifies the identification of a specific instance of the queue-group.

Values 1—16384

vc-label

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

Context config>service>fpipe>spoke-sdp>egress
config>service>apipe>spoke-sdp>egress
config>service>cpipe>spoke-sdp>egress
config>service>ipipe>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>apipe>spoke-sdp>ingress
config>service>cpipe>spoke-sdp>ingress
config>service>fpipe>spoke-sdp>ingress
config>service>ipipe>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

monitor-oper-group

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

Context config>service>epipe>spoke-sdp
config>service>epipe>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.

Default none

Parameters *group-name* — Specifies an oper group name.

oper-group

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

Context config>service>epipe>sap

Description This command configures the operational group identifier.

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

Default none

Parameters *group-name* — Specifies the Operational-Group identifier up to 32 characters in length.

precedence

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

Context config>service>apipe>spoke-sdp
config>service>cpipe>spoke-sdp
config>service>fpipe>spoke-sdp
config>service>ipipe>spoke-sdp
config>service>epipe>spoke-sdp

Description This command specifies the precedence of the SDP binding when there are multiple SDP bindings attached to one service endpoint. The value of zero can only be assigned to one SDP bind making it the primary SDP bind. When an SDP binding goes down, the next highest precedence SDP binding will begin to forward traffic.

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

Default 4

Parameters *precedence-value* — Specifies the spoke SDP precedence.

Values 1 — 4

primary — Specifies to make this the primary spoke SDP.

pw-status-signaling

Syntax [**no**] **pw-status-signaling**

Context config>service>epipe>spoke-sdp

Description This command enables pseudowire status signaling for this spoke SDP binding.

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

Default pw-status-signaling

vc-label

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

Context config>service>cpipe>spoke-sdp>egress
config>service>epipe>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>cpipe>spoke-sdp>ingress
config>service>epipe>spoke-sdp>ingress

Description This command configures the ingress VC label.

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

Values 2048 — 18431

vlan-vc-tag

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

Context config>service>epipe>spoke-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.

spoke-sdp-fec

Syntax **spoke-sdp-fec**
spoke-sdp-fec *spoke-sdp-fec-id* [**fec** *fec-type*] [**aii-type** *aii-type*] [**create**]
spoke-sdp-fec *spoke-sdp-fec-id* **no-endpoint**
spoke-sdp-fec *spoke-sdp-fec-id* [**fec** *fec-type*] [**aii-type** *aii-type*] [**create**] **endpoint** *name* [**icb**]

Context config>service>epipe

Description This command binds a service to an existing Service Distribution Point (SDP), using a dynamic MS-PW.

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.

When using dynamic MS-PWs, the particular SDP to bind-to is automatically selected based on the Target Attachment Individual Identifier (TAII) and the path to use, specified under spoke-SDP FEC. The selected SDP will terminate on the first hop S-PE of the MS-PW. Therefore, an SDP must already be defined in the config>service>sdp context that reaches the first hop 7x50 of the MS-PW. The 7x50 will in order to associate an SDP with a service. If an SDP to that 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.

It differs from the spoke-sdp command in that the spoke-sdp command creates a spoke SDP binding that uses a PW with the PW ID FEC. However, the spoke-sdp-fec command enables PWs with other FEC types to be used. In Release 9.0, only the Generalised ID FEC (FEC129) may be specified using this command.

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 none

Parameters *spoke-sdp-fec-id* — An unsigned integer value identifying the spoke-SDP.

Values 1 —4294967295

fec *fec-type* — An unsigned integer value for the type of the FEC used by the MS-PW.

Values 129 — 130

aii-type *aii-type* — An unsigned integer value for the Attachment Individual Identifier (AII) type used to identify the MS-PW endpoints.

Values 1 — 2

endpoint *endpoint-name* — Specifies the name of the service endpoint

no endpoint — Adds or removes a spoke SDP association.

icb — Configures the spoke-SDP as an inter-chassis backup SDP binding.

auto-config

Syntax [**no**] **auto-config**

Context config>service>epipe>spoke-sdp-fec

Description his command enables single sided automatic endpoint configuration of the spoke-SDP. The 7x50 acts as the passive T-PE for signaling this MS-PW.

Automatic Endpoint Configuration allows the configuration of a spoke-SDP endpoint without specifying the TAII associated with that spoke-SDP. It allows a single-sided provisioning model where an incoming label mapping message with a TAII that matches the SAII of that spoke-SDP to be automatically bound to that endpoint. In this mode, the far end T-PE actively initiates MS-PW signaling and will send the initial label mapping message using T-LDP, while the 7x50 T-PE for which auto-config is specified will act as the passive T-PE.

The **auto-config** command is blocked in CLI if signaling active has been enabled for this spoke-SDP. It it is only applicable to spoke SDPs configured under the Epipe, IES and VPRN interface context.

The **no** form of the command means that the 7x50 T-PE either acts as the active T-PE (if signaling active is configured) or automatically determines which 7x50 will initate MS-PW signaling based on the prefix values configured in the SAII and TAII of the spoke-SDP. If the SAII has the greater prefix value, then the 7x50 will initiate MS-PW signaling without waiting for a label mapping message from the far end. However, if the TAII has the greater value prefix, then the 7x50 will assume that the far end T-PE will intiate MS-PW signaling and will wait for that label mapping message beforeresponding with a T-LDP label mapping message for the MS-PW in the reverse direction.

Default no auto-config

path

Syntax **path** *name*
no path

Context config>service>epipe>spoke-sdp-fec

Description This command specifies the explicit path, containing a list of S-PE hops, that should be used for this spoke SDP. The path-name should correspond to the name of an explicit path configured in the **config>service>pw-routing** context.

If no path is configured, then each next-hop of the MS-PW used by the spoke-SDP will be chosen locally at each T-PE and S-PE.

Default no path

Parameters *path-name* — The name of the explicit path to be used, as configured under config>service>pw-routing.

precedence

Syntax **precedence** *prec-value*
precedence primary
no precedence

Context config>service>epipe>spoke-sdp-fec

Description This command specifies the precedence of the SDP binding when there are multiple SDP bindings attached to one service endpoint. The value of zero can only be assigned to one SDP bind making it the primary SDP bind. When an SDP binding goes down, the next highest precedence SDP binding will begin to forward traffic.

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

Default 42

Parameters *precedence-value* — Specifies the spoke SDP precedence.

Values 1 — 4

primary — Specifies to make this the primary spoke SDP.

pw-template-bind

Syntax **pw-template-bind** *policy-id*
no pw-template-bind

Context config>service>epipe>spoke-sdp-fec

Description This command binds includes the parameters included in a specific PW Template to a spoke SDP.

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

Default none

Parameters *policy-id* — Specifies the existing policy ID

Values 1 — 2147483647

retry-count

Syntax **retry-count** *retry-count*
no retry-count

Context config>service>epipe>spoke-sdp-fec

Description This optional command specifies the number of attempts software should make to re-establish the spoke-SDP after it has failed. After each successful attempt, the counter is reset to zero.

When the specified number is reached, no more attempts are made and the spoke-sdp is put into the shutdown state.

Use the no shutdown command to bring up the path after the retry limit is exceeded.

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

Default 30

Parameters *retry-count* — The maximum number of retries before putting the spoke-sdp into the shutdown state.

Values 10 — 10000

retry-timer

Syntax **retry-timer** *retry-timer*
no retry-timer

Context config>service>epipe>spoke-sdp-fec

Description This command specifies a retry-timer for the spoke-SDP. This is a configurable exponential back-off timer that determines the interval between retries to re-establish a spoke-SDP if it fails and a label withdraw message is received with the status code “AII unreachable”.

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

Default 30

Parameters *retry-timer* — The initial retry-timer value in seconds.

Values 10 — 480

saii-type2

Syntax **saii-type2** *global-id:prefix:ac-id*
no saii-type2

Description This command configures the source attachment individual identifier for the spoke-sdp. This is only applicable to FEC129 AII type 2.

Parameters *global-id* — A Global ID of this 7x50 T-PE. This value must correspont to one of the global_id values configured for a local-prefix under **config>service>pw-routing>local-prefix** context.

Values 1 — 4294967295

prefix — The prefix on this 7x50 T-PE that the spoke-sdp SDP is associated with.This value must correspond to one of the prefixes configured under **config>service>pw-routing>local-prefix context.**

Values an IPv4-formatted address a.b.c.d or 1 — 4294967295

ac-id — An unsigned integer representing a locally unique identifier for the spoke-SDP.

Values 1 — 4294967295

signaling

Syntax **signaling** *signaling*

Context config>service>epipe>spoke-sdp-fec

Description This command enables a user to configure this 7x50 as the active pr passive T-PE for signaling this MS-PW, or to automatically select whether this T-PE is active or passive based on the prefix. In an active role, this endpoint initiates MS-PW signaling without waiting for a T-LDP label mapping message to arrive from the far end T-PE. In a passive role, it will wait for the initial label mapping message from the far end before sending a label maping for this end of the PW. In auto maode, if the SAII has the greater prefix value, then the 7x50 will initiate MS-PW signaling without waiting for a label mapping message from the far end. However, if the TAII has the greater value prefix, then the 7x50 will assume that the far end T-PE will intiate MS-PW signaling and will wait for that label mapping message beforeresponding with a T-LDP label mapping message for the MS-PW in the reverse direction.

The **no** form of the command means that the 7x50 T-PE automatically selects the which 7x50 will initate MS-PW signaling based on the prefix values configured in the SAII and TAII of the spoke-SDP, as described above.

Default auto

Parameters *signaling* — Configures this 7x50 as the active T-PE for signaling this MS-PW.

Values auto, master

standby-signaling-slave

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

Context config>service>epipe>spoke-sdp-fec

taii-type2

Syntax **taii-type2** *global-id:prefix:ac-id*
no taii-type2

Context config>service>epipe>spoke-sdp-fec

Description taii-type2 configures the target attachment individual identifier for the spoke-sdp. This is only applicable to FEC129 AII type 2.

This command is blocked in CLI if this end of the spoke-SDP is configured for single-sided auto configuration (using the **auto-config** command).

Parameters *global-id* — A Global ID of this 7x50 T-PE. This value must correspont to one of the global_id values configured for a local-prefix under **config>service>pw-routing>local-prefix** context.

Values 1 — 4294967295

prefix — The prefix on this 7x50 T-PE that the spoke-sdp SDP is associated with.This value must correspond to one of the prefixes configured under **config>service>pw-routing>local-prefix context.**

Values an IPv4-formatted address a.b.c.d or 1 — 4294967295

ac-id — An unsigned integer representing a locally unique identifier for the spoke-SDP.

Values 1 — 4294967295

ATM Commands

atm

Syntax **atm**

Context config>service>epipe>sap
config>service>apipe>sap
config>service>ipipe>sap
config>service>epipe>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>epipe>sap
config>service>epipe>sap>atm
config>service>apipe>sap>atm
config>service>fpipe>sap

This command configures egress ATM attributes for the SAP.

ingress

Syntax **ingress**

Context config>service>epipe>sap
config>service>epipe>sap>atm
config>service>epipe>sap
config>service>apipe>sap>atm

Description This command configures ingress ATM attributes for the SAP.

encapsulation

Syntax **encapsulation** *atm-encap-type*

Context config>service>epipe>sap>atm
config>service>ipipe>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

traffic-desc

Syntax **traffic-desc** *traffic-desc-profile-id*
no traffic-desc

Context config>service>epipe>sap
config>service>apipe>sap>atm>egress
config>service>apipe>sap>atm>ingress
config>service>epipe>sap>atm>egress
config>service>epipe>sap>atm>ingress

Description This command assigns an ATM traffic descriptor profile to a given context (for example, a SAP).

When configured under the ingress context, the specified traffic descriptor profile defines the traffic contract in the forward direction. When configured under the egress context, the specified traffic descriptor profile defines the traffic contract in the backward direction.

The **no** form of the command reverts the traffic descriptor to the default traffic descriptor profile.

Default The default traffic descriptor (trafficDescProfileId. = 1) is associated with newly created PVCC-delimited SAPs.

Parameters *traffic-desc-profile-id* — Specify a defined traffic descriptor profile (see the QoS atm-td-profile command).

OAM Commands

oam

Syntax **oam**

Context config>service>epipe>sap
config>service>apipe>sap>atm

Description This command enables the context to configure OAM functionality for a PVCC delimiting a SAP.

- The ATM-capable MDAs support end-to-end and segment OAM functionality (AIS, RDI, Loop-back) over both F5 (VC) and end-to-end F4 (VP) OAM:
- ITU-T Recommendation I.610 - B-ISDN Operation and Maintenance version 11/95
- GR-1248-CORE - Generic Requirements for Operations of ATM N3 June 1996
- GR-1113-CORE - Bellcore, Asynchronous Transfer Mode (ATM) (AAL) Protocols Generic Requirements, Issue 1, July 1994

alarm-cells

Syntax [**no**] **alarm-cells**

Context config>service>epipe>sap>oam
config>service>epipe>sap>oam
config>service>apipe>sap>atm>oam

Description This command configures AIS/RDI fault management on a PVCC. Fault management allows PVCC terminations 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, a PVCC's operational status is no longer affected by a 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 due to alarm-cell processing) and RDI cells are not generated as result of the PVCC going into AIS or RDI state. The PVCC's OAM status, however, will record OAM faults as described above.

Default Enabled for PVCCs delimiting IES SAPs

terminate

Syntax [**no**] **terminate**

Context config>service>apipe>sap>atm>oam

Description This command specifies whether this SAP will act as an OAM termination point. ATM SAPs can be configured to tunnel or terminate OAM cells.

When configured to not terminate (the default is **no terminate**), the SAP will pass OAM cells through the VLL without inspecting them. The SAP will respond to OAM loopback requests that are directed to the local node by transmitting a loopback reply. Other loopback requests are transparently tunneled through the pseudowire. In this mode, it is possible to launch a loopback request towards the directly-attached ATM equipment and see the results of the reply.

When configured to terminate, the SAP will respond to AIS by transmitting RDI and will signal the change of operational status to the other endpoint (for example, through LDP status notifications). The SAP will respond to OAM loopback requests by transmitting a loopback reply. In this mode, it is possible to launch a loopback request towards the directly-attached ATM equipment and see the results of the reply.

For Apipe services, the user has the option of enabling or disabling this option for VC types atm-vcc and atm-sdu since these service types maintain the ATM layer and/or the AAL5 layer across the VLL. It is not supported on atm-vpc and atm-cell apipe vc types since the VLL must pass the VC level (F5) OAM cells.

The **terminate** option for OAM is the only and default mode of operation supported for an ATM SAP which is part of Epipe, Ipipe, VPLS, and IES/VPRN. This is because the ATM and AAL5 layers are terminated.

For Apipe services, the user has the option of enabling or disabling this option for vc types atm-vcc and atm-sdu since these service types maintain the ATM layer and/or the AAL5 layer across the VLL. It is not supported on atm-vpc and atm-cell Apipe vc types since the VLL must pass the VC level (F5).

The **terminate** option for OAM is the only and default mode of operation supported for an ATM SAP which is part of Epipe, Ipipe, VPLS, and IES/VPRN. This is because the ATM and AAL5 layers are terminated.

Default no terminate

Cpipe Commands

endpoint

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

Context config>service>cpipe

Description This command configures a service endpoint.

Parameters *endpoint-name* — Specifies an endpoint name.

active-hold-delay

Syntax **active-hold-delay** *active-hold-delay*
no active-hold-delay

Context config>service>cpipe>endpoint

Description This command specifies that the node will delay sending the change in the T-LDP status bits for the service endpoint when the MC-LAG transitions the LAG subgroup which hosts the SAP for this VLL endpoint from "active" to "standby" or when any object in the endpoint. For example., SAP, ICB, or regular spoke SDP, transitions from up to down operational state.

By default, when the MC-LAG transitioned the LAG subgroup which hosts the SAP for this VLL endpoint from "active" to "standby", the node sends immediately new T-LDP status bits indicating the new value of "standby" over the spoke SDPs which are on the mate-endpoint of the VLL. The same applies when any object in the endpoint changes an operational state from up to down.

There is no delay applied to the VLL endpoint status bit advertisement when the MC-LAG transitions the LAG subgroup which hosts the SAP from "standby" to "active" or when any object in the endpoint transitions to an operationally up state.

Default 0 — A value of zero means that when the MC-LAG transitioned the LAG subgroup which hosts the SAP for this VLL endpoint from "active" to "standby", the node sends immediately new T-LDP status bits indicating the new value of "standby" over the spoke SDPs which are on the mate-endpoint of the VLL. The same applies when any object in the endpoint changes an operational state from up to down.

Parameters *active-hold-delay* — Specifies the active hold delay in 100s of milliseconds.

Values 0 — 60

revert-time

Syntax **revert-time** *revert-time*
no revert-time

Context config>service>cpipe>endpoint

Description This command configures the time to wait before reverting back to the primary spoke SDP defined on this service endpoint, after having failed over to a backup spoke SDP.

Parameters *revert-time* — Specify the time, in seconds, to wait before reverting to the primary SDP.

Values 0 — 600

infinite — Causes the endpoint to be non-revertive.

CES SAP Commands

sap

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

Context config>service>cpipe

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 service router. Each SAP must be unique.

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

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

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

If a port is shutdown, all SAPs on that port become operationally down. When a service is shutdown, SAPs for the service are not displayed as operationally down although all traffic traversing the service will be discarded.

The operational state of a SAP is relative to the operational state of the port on which the SAP is defined.

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

Default No SAPs are defined.

Special Cases A SAP can be defined with Ethernet ports, SONET/SDH or TDM channels. At most, only one sdp-id can be bound to an VLL service. Since a VLL is a point-to-point service, it can have, at most, two end points. The two end points can be one SAP and one SDP or two SAPs. Up to 49 SDPs can be associated with a service in a single router. Each SDP must have a unique router destination or an error will be generated.

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. 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.

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

If the card in the slot has Media Dependent Adapters (MDAs) installed, the *port-id* must be in the slot_number/MDA_number/port_number format. For example 61/2/3 specifies port 3 on MDA 2 in slot 61.

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

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

bundle-id — Specifies the multilink bundle to be associated with this IP interface. The **bundle** keyword must be entered at the beginning of the parameter.
The command syntax must be configured as follows:

bundle-id: **bundle-***type-slot-id*/*mda-slot.bundle-num*
bundle-id value range: 1 — 128

For example:

```
*A:ALA-12>config#  port bundle-ppp-5/1.1
*A:ALA-12>config>port# multilink-bundle
```

bgprp-id — Specifies the bundle protection group ID to be associated with this IP interface. The **bpgrp** keyword must be entered at the beginning of the parameter.

The command syntax must be configured as follows:

bpgrp-id: bpgrp-*type-bpgrp-num*
type: ima
bpgrp-num value range: 1 — 1280

For example:

```
*A:ALA-12>config# port bpgrp-ima-1
*A:ALA-12>config>service>cpipe$ sap bpgrp-ima-1
```

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

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

endpoint — Adds a SAP endpoint association.

no endpoint — Removes the association of a SAP or a spoke-sdp with an explicit endpoint name.

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

cem

Syntax **cem**

Context config>service>cpipe>sap

Description This command enables the context to specify circuit emulation (CEM) properties.

packet

Syntax **packet jitter-buffer** *milliseconds* [**payload-size** *bytes*]
packet payload-size *bytes*
no packet

Context config>service>cpipe>sap

Description This command specifies the jitter buffer size, in milliseconds, and payload size, in bytes.

Default The default value depends on the CEM SAP endpoint type, and if applicable, the number of timeslots:

Endpoint Type	Timeslots	Default Jitter Buffer (in ms)
unstructuredE1	n/a	5
unstructuredT1	n/a	5
unstructuredE3	n/a	5
unstructuredT3	n/a	5

Endpoint Type	Timeslots	Default Jitter Buffer (in ms)
nxDS0 (E1/T1)	N = 1	32
	N = 2..4	16
	N = 5..15	8
	N >= 16	5
nxDS0WithCas (E1)	N	8
nxDS0WithCas (T1)	N	12

Parameters *milliseconds* — specifies the jitter buffer size in milliseconds (ms).

Configuring the payload size and jitter buffer to values that result in less than 2 packet buffers or greater than 32 packet buffers is not allowed.

Setting the jitter butter value to 0 sets it back to the default value.

Values 1 — 250

payload-size *bytes* — Specifies the payload size (in bytes) of packets transmitted to the packet service network (PSN) by the CEM SAP. This determines the size of the data that will be transmitted over the service. If the size of the data received is not consistent with the payload size, then the packet is considered malformed.

Default The default value depends on the CEM SAP endpoint type, and if applicable, the number of timeslots:

Endpoint Type	Timeslots	Default Payload Size (in bytes)
unstructuredE1	n/a	256
unstructuredT1	n/a	192
unstructuredE3	n/a	1024
unstructuredT3	n/a	1024
nxDS0 (E1/T1)	N = 1	64
	N = 2..4	N x 32
	N = 5..15	N x 16
	N >= 16	N x 8
nxDS0WithCas (E1)	N	N x 16
nxDS0WithCas (T1)	N	N x 24

For all endpoint types except for nxDS0WithCas, the valid payload size range is from the default to 2048 bytes.

For nxDS0WithCas, the payload size divide by the number of timeslots must be an integer factor of the number of frames per trunk multi-frame (for example, 16 for E1 trunk and 24 for T1 trunk).
For 1xDS0, the payload size must be a multiple of 2.
For NxDS0, where N > 1, the payload size must be a multiple of the number of timeslots.
For unstructuredE1, unstructuredT1, unstructuredE3 and unstructuredT3, the payload size must be a multiple of 32 bytes.

Configuring the payload size and jitter buffer to values that result in less than 2 packet buffers or greater than 32 packet buffer is not allowed.

Setting the payload size to 0 sets it back to the default value.

Values 0, 16 — 2048

report-alarm

Syntax [**no**] **report-alarm** [**stray**] [**malformed**] [**pktloss**] [**overrun**] [**underrun**] [**rpktloss**] [**rfault**] [**rrdi**]

Context config>service>cpipe>sap>cem

Description This command indicates the type of CEM SAP alarm.

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

Parameters **stray —** Reports the reception of packets not destined for this CES circuit.

malformed — Reports the reception of packet not properly formatted as CES packets.

pktloss — Reports the lack of reception of CES packets.

overrun — Reports reports the reception of too many CES packets resulting in a overrun of the receive jitter buffer.

underrun — Reportsreports the reception of too few CES packets resulting in a overrun of the receive jitter buffer.

rpktloss — Reports hat the remote peer is currently in packet loss status.

rfault — Reports that the remote TDM interface is currently not in service.

rrdi — Reports that the remote TDM interface is currently in RDI status.

rtp-header

Syntax [**no**] **rtp-header**

Context config>service>cpipe>sap>cem

Description This command specifies whether an RTP header is used when packets are transmitted to the packet service network (PSN) by the CEM SAP.

Default no rtp-header

Service Filter and QoS Policy Commands

egress

Syntax **egress**

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

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

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

ingress

Syntax **ingress**

Context config>service>cpipe>sap

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

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

agg-rate-limit

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

Context config>service>cpipe>sap>egress

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

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

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

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

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

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

Values 1 — 40000000, max

qinq-mark-top-only

Syntax [**no**] **qinq-mark-top-only**

Context config>service>cpipe>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 the enabled, only the P-bits/DEI bit in the top Q-tag are marked.

Default no qinq-mark-top-only

qos

Syntax **qos** *policy-i*
no qos

Context config>service>cpipe>sap>egress
config>service>cpipe>sap>ingress

Description This command associates a Quality of Service (QoS) policy with an ingress or egress Service Access Point (SAP) or IP interface.

QoS ingress and egress policies are important for the enforcement of SLA agreements. The policy ID must be defined prior to associating the policy with a SAP or IP interface. If the *policy-id* does not exist, an error will be returned.

The **qos** command is used to associate both ingress and egress QoS policies. The **qos** command only allows ingress policies to be associated on SAP or IP interface ingress and egress policies on SAP or IP interface egress. Attempts to associate a QoS policy of the wrong type returns an error.

Only one ingress and one egress QoS policy can be associated with a SAP or IP interface at one time. Attempts to associate a second QoS policy of a given type will return an error.

By default, no specific QoS policy is associated with the SAP or IP interface for ingress or egress, so the default QoS policy is used.

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

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

Values 1 — 65535

shared-queuing — This keyword can **only** be specified on SAP ingress. The **shared-queueing** keyword specifies the shared queue policy will be used by this SAP. When the value of this object is null, the SAP will use individual ingress QoS queues, instead of the shared ones.

queue-override

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

Context config>service>cpipe>sap>egress
config>service>cpipe>sap>ingress

Description This command enables the context to configure override values for the specified SAP egress or ingress QoS queue. These values override the corresponding ones specified in the associated SAP egress or ingress QoS policy. If the policy was created as a template policy, this command overrides the parameter and its description and queue parameters in the policy.

queue

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

Context config>service>cpipe>sap>egress>queue-override
config>service>cpipe>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.

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

adaptation-rule

Syntax **adaptation-rule** [**pir** *adaptation-rule*}] [**cir** *adaptation-rule*}]

Context config>service>cpipe>sap>egress>queue-override>queue
config>service>cpipe>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.

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>cpipe>sap>egress>queue-override>queue
config>service>cpipe>sap>ingress>queue-override>queue

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

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

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

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

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

  As a special case, when a queue or associated intermediate scheduler is configured with a CIR-weight equal to 0, the system automatically sets the queue's frame based within-cir offered-load to 0, preventing it from receiving bandwidth during the port scheduler's within-cir pass.
- Frame based PIR — The frame based PIR is calculated by multiplying the packet to frame factor with the queue's configured PIR and then adding the result to that PIR. If the queue PIR is set to 7500 octets and the packet to frame factor equals 0.1, the frame based PIR would be 7500 x 1.1 or 8250 octets.
- Frame based within-pir offered-load — The frame based within-pir offered-load is the portion of the frame based offered-load considered to be within the frame based PIR. The frame based within-pir offered-load is the lesser of the frame based offered-load and the frame based PIR. If the frame based offered-load equaled 11000 octets and the frame based PIR equaled 8250 octets, the frame based within-pir offered-load would be limited to 8250 octets. If the frame based offered-load equaled 7000 octets and the frame based PIR equaled 8250 octets, the frame based within-pir offered load would equal 7000 octets.

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

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

The **no** form of this command restores the average frame overhead parameter for the queue to the default value of 0 percent. When set to 0, the system uses the packet based queue statistics for calculating port scheduler priority bandwidth allocation. If the no avg-frame-overhead command is

executed in a queue-override queue id context, the avg-frame-overhead setting for the queue within the sap-egress QoS policy takes effect.

Default 0

Parameters *percent* — This parameter sets the average amount of packet-to-frame encapsulation overhead expected for the queue. This value is not used by the system for egress Ethernet queues.

Values 0.00 — 100.00

cbs

Syntax **cbs** *size-in-kbytes*
no cbs

Context config>service>cpipe>sap>egress>queue-override>queue
config>service>cpipe>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>cpipe>sap>egress>queue-override>queue
config>service>cpipe>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*
no mbs

Context config>service>cpipe>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>cpipe>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 packet's RED slope will not force the discard of the packet. Setting proper CBS parameters and controlling CBS oversubscription is one major safeguard to queue starvation (when a queue does not receive its fair share of buffers). Another is properly setting the RED slope parameters for the needs of services on this port or channel. If the CBS value is larger than the MBS value, an error will occur, preventing the MBS change.

The defined high-prio-only value cannot be greater than the MBS size of the queue. Attempting to change the MBS to a value smaller than the high priority reserve will generate an error and fail execution. Attempting to set the high-prio-only value larger than the current MBS size will also result in an error and fail execution.

The **no** form of this command returns the MBS size 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>cpipe>sap>egress>queue-override>queue
config>service>cpipe>sap>ingress>queue-override>queue

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

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

The CIR defines the rate at which the system prioritizes the queue over other queues competing for the same bandwidth. In-profile packets are preferentially queued by the system at egress and at

subsequent next hop nodes where the packet can traverse. To be properly handled as in- or out-of-profile throughout the network, the packets must be marked accordingly for profiling at each hop.

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

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

Values 0 — 100000000, **max, sum**

Default 0

scheduler-override

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

Context config>service>cpipe>sap>egress
config>service>cpipe>sap>ingress

Description This command specifies the set of attributes whose values have been overridden by 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* [**create**]
no scheduler *scheduler-name*

Context config>service>cpipe>sap>egress>sched-override
config>service>cpipe>sap>ingress>sched-override

Description This command can be used to override specific attributes of the specified scheduler name. A scheduler defines bandwidth controls that limit each child (other schedulers and queues) associated with the scheduler. Scheduler objects are created within the hierarchical tiers of the policy. It is assumed that each scheduler created will have queues or other schedulers defined as child associations. The scheduler can be a child (take bandwidth from a scheduler in a higher tier, except for schedulers created in tier 1). 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 keyword createsa 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>cpipe>sap>egress>sched-override>scheduler
config>service>cpipe>sap>ingress>sched-override>scheduler

Description This command can be used to override specific attributes of the specified scheduler rate. The **rate** command defines the maximum bandwidth that the scheduler can offer its child queues or schedulers. The maximum rate is limited to the amount of bandwidth the scheduler can receive from its parent scheduler. If the scheduler has no parent, the maximum rate is assumed to be the amount available to the scheduler. When a parent is associated with the scheduler, the CIR parameter provides the amount of bandwidth to be considered during the parent scheduler's 'within CIR' distribution phase.

The actual operating rate of the scheduler is limited by bandwidth constraints other than its maximum rate. The scheduler's parent scheduler may not have the available bandwidth to meet the scheduler's needs or the bandwidth available to the parent scheduler could be allocated to other child schedulers or child queues on the parent based on higher priority. The children of the scheduler may not need the maximum rate available to the scheduler due to insufficient offered load or limits to their own maximum rates.

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

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

Parameters *pir-rate* — The **pir** parameter accepts a step multiplier value that specifies the multiplier used to determine the PIR rate at which the queue will operate. A value of 0 to 100000000 or the keyword **max** or **sum** 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 to 250 or the keyword max is accepted. Any other value will result in an error without modifying the current CIR rate.

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

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

Values 0 — 10000000, max, sum

Default sum

scheduler-policy

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

Context config>service>cpipe>sap>ingress
config>service>cpipe>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.

multi-service-site

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

Context config>service>cpipe>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. See multi-service-site on page 116.

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.

tod-suite

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

Context config>service>cpipe>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.

service-mtu

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

Context config>service>cpipe

Description This command configures the service payload (Maximum Transmission Unit – MTU), in bytes, for the service. This MTU value overrides the service-type default MTU. The **service-mtu** defines the payload capabilities of the service. It is used by the system to validate the SAP and SDP binding's operational state within the service.

The service MTU and a SAP's service delineation encapsulation overhead (i.e., 4 bytes for a dot1q tag) is used to derive the required MTU of the physical port or channel on which the SAP was created. If the required payload is larger than the port or channel MTU, then the SAP will be placed in an inoperative state. If the required MTU is equal to or less than the port or channel MTU, the SAP will be able to transition to the operative state.
When binding an SDP to a service, the service MTU is compared to the path MTU associated with the SDP. The path MTU can be administratively defined in the context of the SDP. The default or administrative path MTU can be dynamically reduced due to the MTU capabilities discovered by the tunneling mechanism of the SDP or the egress interface MTU capabilities based on the next hop in the tunnel path.

If the service MTU is larger than the path MTU minus control word length (if applicable), the SDP binding for the service will be placed in an inoperative state with sdp-bind oper flag PathMTUTooSmall.

If the CEM SAP's packet size is larger than the service MTU then the service will be placed in an inoperative state with service oper flag ServiceMTUTooSmall . The CEM SAP packet size is defined as CEM SAP payload-size plus rtp-header size (if applicable).

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

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

Default cpipe: 1514

octets — The size of the MTU in octets, expressed as a decimal integer, between 1 — 9194.

service-name

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

Context config>service>cpipe
config>service>epipe

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 router platforms.

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

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

site

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

Context config>service>epipe

Description This command configures a Epipe 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 service.

CPipe SDP Commands

spoke-sdp

Syntax **spoke-sdp** *sdp-id*[**:***vc-id*] [**no-endpoint**] [**create**]
spoke-sdp *sdp-id:vc-id* [**create**] **endpoint** *endpoint-name* [**icb**]
no spoke-sdp *sdp-id*[**:***vc-id*]

Context config>service>cpipe

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. 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.

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

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.
- The VC type value for a VPLS service is defined as 0x000B.

Values ethernet

no endpoint — removes the association of a spoke SDP with an explicit endpoint name.

endpoint *endpoint-name* **—** Specifies the name of the service endpoint.

icb — Configures the spoke SDP as an inter-chassis backup SDP binding.

egress

Syntax **egress**

Context config>service>cpipe>spoke-sdp

Description This command enables the context to configure egress spoke-SDP context.

ingress

Syntax **ingress**

Context config>service>cpipe>spoke-sdp

Description This command enables the context to configure ingress spoke-SDP context.

vc-label

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

Context config>service>cpipe>spoke-sdp>egress

Description This command configures the spoke-SDP egress VC label.

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

Values 16 — 1048575

vc-label

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

Context config>service>cpipe>spoke-sdp>ingress

Description This command configures the spoke-SDP ingress VC label.

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

Values 2048 — 18431

precedence

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

Context config>service>cpipe>spoke-sdp

Description This command specifies the precedence of the SDP binding when there are multiple SDP bindings attached to one service endpoint. The value of zero can only be assigned to one SDP bind making it the primary SDP bind. When an SDP binding goes down, the next highest precedence SDP binding will begin to forward traffic.

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

Default 4

Parameters *precedence-value* — Specifies the spoke SDP precedence.

Values 1 — 4

primary — Specifies to make this the primary spoke SDP.