Global Service Configuration Commands

Generic Commands

shutdown

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

Context

```
config>eth-cf>mep
config>service>sdp
config>service>sdp>class-forwarding
config>service>sdp>keep-alive
config>service>sdp>forwarding-class
config>service>pw-routing>hop
config>service>sdp>binding>pw-port
config>eth-tunnel>path
config>eth-tunnel>path>eth-cfm>mep
config>eth-tunnel
```

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.

SDP (global) — When an SDP is shutdown at the global service level, all bindings to that SDP are put into the out-of-service state and the SDP itself is put into the administratively and operationally down states. Packets that would normally be transmitted using this SDP binding will be discarded and counted as dropped packets.

SDP (service level) — Shutting down an SDP within a service only affects traffic on that service from entering or being received from the SDP. The SDP itself may still be operationally up for other services.

SDP Keepalives — Enables SDP connectivity monitoring keepalive messages for the SDP ID. Default state is disabled (shutdown) in which case the operational state of the SDP-ID is not affected by the keepalive message state.

description

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

Context config>service>customer
config>service>customer>multi-service-site
config>service>pw-template
config>service>pw-template>split-horizon-group
config>service>sdp
config>eth-tunnel
config>eth-tunnel>path
config>eth-tunnel>path>eth-cfm>mep

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.

new-qinq-untagged-sap

Syntax [**no**] **new-qinq-untagged-sap**

Context config>system>ethernet

Description This command controls the behavior of QinQ SAP y.0 (for example, 1/1/1:3000.0). If the flag is not enabled (no new-qinq-untagged-sap), the y.0 SAP works the same as the y.* SAP (for example, 1/1/1:3000.*); all frames tagged with outer VLAN y and no inner VLANs or inner VLAN x where inner VLAN x is not specified in a SAP y.x configured on the same port (for example, 1/1/1:3000.10).

If the flag is enabled, then the following new behavior immediately applies to all existing and future y.0 SAPs: the y.0 SAP maps all the ingress frames tagged with outer tag VLAN-id of y (qinq-etype) and no inner tag or with inner tag of VLAN-id of zero (0).

Default no new-qinq-untagged-sap. This setting ensures that there will be no disruption for existing usage of this SAP type.

Customer Management Commands

customer

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

Context config>service

Description This command creates a customer ID and customer context used to associate information with a particular customer. Services can later be associated with this customer at the service level.

Each *customer-id* must be unique. The *create* keyword must follow each new **customer** *customer-id* entry.

Enter an existing **customer** *customer-id* (without the *create* keyword) to edit the customer's parameters.

Default **customer 1** always exists on the system and cannot be deleted.

The **no** form of this command removes a *customer-id* and all associated information. Before removing a *customer-id*, all references to that customer in all services must be deleted or changed to a different customer ID.

Parameters *customer-id* — Specifies the ID number to be associated with the customer, expressed as an integer.

Values 1 — 2147483647

create — This keyword is required when first creating the configuration context. Once the context is created, it is possible to navigate into the context without the **create** keyword.

contact

Syntax **contact** *contact-information*
no contact *contact-information*

Context config>service>customer

Description This command allows you to configure contact information for a customer.

Include any customer-related contact information such as a technician's name or account contract name.

Default No contact information is associated with the *customer-id*.

The **no** form of this command removes the contact information from the customer ID.

Parameters *contact-information* — The customer contact information entered as an ASCII character string up to 80 characters in length. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes. Any printable, seven bit ASCII characters may be used within the string.

multi-service-site

Syntax **multi-service-site** *customer-site-name* [**create**]
no multi-service-site *customer-site-name*

Context config>service>customer

Description This command creates a new customer site or edits an existing customer site with the *customer-site-name* parameter. A customer site is an anchor point to create an ingress and egress virtual scheduler hierarchy. When a site is created, it must be assigned to a chassis slot or port with the exception of the 7750 SR-1 in which the slot is set to 1.When scheduler policies are defined for ingress and egress, the scheduler names contained in each policy are created according to the parameters defined in the policy. Multi-service customer sites exist for the sole purpose of creating a virtual scheduler hierarchy and making it available to queues on multiple Service Access Points (SAPs).

The scheduler policy association with the customer site normally prevents the scheduler policy from being deleted until after the scheduler policy is removed from the customer site. The multi-service-site object will generate a log message indicating that the association was deleted due to scheduler policy removal.

When the multi-service customer site is created, an ingress and egress scheduler policy association does not exist. This does not prevent the site from being assigned to a chassis slot or prevent service SAP assignment. After the site has been created, the ingress and egress scheduler policy associations can be assigned or removed at any time.

Default None — Each customer site must be explicitly created.

Parameters *customer-site-name* — Each customer site must have a unique name within the context of the customer. If *customer-site-name* already exists for the customer ID, the CLI context changes to that site name for the purpose of editing the site scheduler policies or assignment. Any modifications made to an existing site will affect all SAPs associated with the site. Changing a scheduler policy association may cause new schedulers to be created and existing queues on the SAPs to no longer be orphaned. Existing schedulers on the site may cease to exist, causing queues relying on that scheduler to be orphaned.

If the *customer-site-name* does not exist, it is assumed that an attempt is being made to create a site of that name in the customer ID context. The success of the command execution depends on the following:

- The maximum number of customer sites defined for the chassis has not been met.
- The *customer-site-name* is valid.
- The **create** keyword is included in the command line syntax (if the system requires it).

When the maximum number of customer sites has been exceeded a configuration error occurs; the command will not execute and the CLI context will not change.

If the *customer-site-name* is invalid, a syntax error occurs; the command will not execute and the CLI context will not change.

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

phone

Syntax [**no**] **phone** *string*

Context config>service>customer *customer-id*

Description This command adds telephone number information for a customer ID.

Default none

The **no** form of this command removes the phone number value from the customer ID.

Parameters *string* — The customer phone number entered as an ASCII string string up to 80 characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes. Any printable, seven bit ASCII characters may be used within the string.

assignment

Syntax **assignment** {**port** *port-id* | **card** *slot-number*}
no assignment

Context config>service>customer>multi-service-site

Description This command assigns a multi-service customer site to a specific chassis slot, port, or channel. This allows the system to allocate the resources necessary to create the virtual schedulers defined in the ingress and egress scheduler policies as they are specified. This also verifies that each SAP assigned to the site exists within the context of the proper customer ID and that the SAP was configured on the proper slot, port, or channel. The assignment must be given prior to any SAP associations with the site.

The **no** form of the command removes the port, channel, or slot assignment. If the customer site has not yet been assigned, the command has no effect and returns without any warnings or messages.

Default None

Parameters **port** *port-id* — The **port** keyword is used to assign the multi-service customer site to the port-id or port-id.channel-id given. When the multi-service customer site has been assigned to a specific port or channel, all SAPs associated with this customer site must be on a service owned by the customer and created on the defined port or channel. The defined port or channelmust already have been pre-provisioned on the system but need not be installed when the customer site assignment is made.

Syntax: *port-id*[:encap-val]

Values	port-id	*slot*/*mda*/*port*[.*channel*]		
		aps-id	aps-*group-id*[.*channel*]	
			aps	keyword
			group-id	1 — 64
			group-id	1 — 16
		bundle-*type-slot/mda.bundle-num*		
			bundle	keyword
			type	ima, ppp
			bundle-num	1 — 256
		bpgrp-id:	**bpgrp**-*type-bpgrp-num*	
			bpgrp	keyword
			type	ima

bpgrp-num 1 — 1280
ccag-id - ccag-<id>.<path-id>[cc-type]
ccag keyword
id 1 — 8
path-ida, b
cc-type[.sap-net | .net-sap]
lag-id lag-*id*
lag keyword
id 1 — 200
lag-id lag-*id*
lag keyword
id 1 — 64

card *slot-number* — The **card** keyword is used to assign the multi-service customer site to the slot-number given. When the multi-service customer site has been assigned to a specific slot in the chassis, all SAPs associated with this customer site must be on a service owned by the customer and created on the defined chassis slot. The defined slot must already have been pre-provisioned on the system but need not be installed when the customer site assignment is made.

Values Any pre-provisioned slot number for the chassis type that allows SAP creation
slot-number 1 — 10

ingress

Syntax ingress

Context config>service>customer>multi-service-site

Description This command enables the context to configure the ingress node associate an existing scheduler policy name with the customer site. The ingress node is an entity to associate commands that complement the association.

egress

Syntax egress

Context config>service>customer>multi-service-site

Description This command enables the context to configure the egress node associate an existing scheduler policy name with the customer site. The egress node is an entity to associate commands that complement the association.

agg-rate-limit

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

Context config>service>customer>multi-service-site>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 optional queue-frame-based-accounting keyword allows the service queues within the SAPs to operate in the frame based accounting mode.

Once egress frame based accounting is enabled on a SAP or Multi-Service Site, all queues associated with the SAP or SAPs will have their rate and CIR values interpreted as frame based values. When shaping, the queues will include the 12 byte Inter-Frame Gap (IFG) and 8 byte preamble for each packet scheduled out the queue. The profiling CIR threshold will also include the 20 byte frame encapsulation overhead. Statistics associated with the queue will also include the frame encapsulation overhead within the octet counters.

The queue-frame-based-accounting keyword does not change the behavior of the agg-rate-limit rate value. Since agg-rate-limit is always associated with egress port based scheduling and egress port based scheduling is dependant on frame based operation, the agg-rate-limit rate is always interpreted as a frame based value.

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

Parameters {**max** | *kilobits-per-second*} **—** The max keyword and kilobits-per-second parameter are mutually exclusive. Either max or a value for kilobits-per-second must follow the agg-rate-limit command.

max — The max keyword specifies that the egress aggregate rate limit for the SAP or the Multi-Service Site is unlimited. Scheduling for the service queues will only be governed by the individual queue parameters and any congestion on the port relative to each queues scheduling priority.

kilobits-per-second — The kilobits-per-second parameter defines an actual egress aggregate rate to which all queues associated with the SAP or Multi-Service Site will be limited. The value must be defined as an integer and is representative of increments of 1000 bits per second.

Values 1 to 40000000

Default max

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

Default Frame based accounting is disabled by default

scheduler-override

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

Context config>service>customer>multi-service-site>ingress
config>service>customer>multi-service-site>egress

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>customer>multi-service-site>ingress>sched-override
config>service>customer>multi-service-site>egress>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>customer>multi-service-site>ingress>sched-override>scheduler
config>service>customer>multi-service-site>egress>sched-override>scheduler

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

The **rate** command defines the maximum bandwidth that the scheduler can offer its child queues or schedulers. The maximum rate is limited to the amount of bandwidth the scheduler can receive from its parent scheduler. If the scheduler has no parent, the maximum rate is assumed to be the amount available to the scheduler. When a parent is associated with the scheduler, the CIR parameter provides the scheduler's amount of bandwidth to be considered during the parent schedulers '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 at which the queue will be allowed to forward packets in a given second, thus shaping the queue's output.

The PIR parameter for SAP ingress queues do not have a negate (**no**) function. To return the 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 100000000 or the keyword **max** or **sum** are 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>customer>multi-service-site>ingress
config>service>customer>multi-service-site>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.

tod-suite

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

Context config>service>cust>multi-service-site

Description This command applies a time-based policy (filter or QoS policy) to the multiservice site. 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. Only the scheduler-policy part of the tod-suite is taken into account. The suite can be applied to more than one multi-service-site.

MRP Commands

mrp

Syntax **mrp**

Context config>service

Description This command configures a Multi-service Route Processor (MRP).

mrp-policy

Syntax [**no**] **mrp-policy** *policy-name*

Context config>service>mrp

Description This command enables the context for a MRP policy. The mrp-policy specifies either a forward or a drop action for the Group BMAC attributes associated with the ISIDs specified in the match criteria. The mrp-policy can be applied to multiple BVPLS services as long as the scope of the policy is template.

Any changes made to the existing policy, using any of the sub-commands, will be applied immediately to all services where this policy is applied. For this reason, when many changes are required on a mrp-policy, it is recommended that the policy be copied to a work area. That work-in-progress policy can be modified until complete and then written over the original mrp-policy. Use the config mrp-policy copy command to maintain policies in this manner.

The **no** form of the command deletes the mrp-policy. An MRP policy cannot be deleted until it is removed from all the SAPs or SDPs where it is applied.

Default no mrp-policy is defined

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

create — This keyword is required when first creating the configuration context. Once the context is created, it is possible to navigate into the context without the **create** keyword.

scope

Syntax **scope** {**exclusive** | **template**}
no scope

Context config>service>mrp>mrp-policy

Description This command configures the filter policy scope as exclusive or template. If the scope of the policy is template and is applied to one or more services, the scope cannot be changed.

The **no** form of the command sets the scope of the policy to the default of template.

Default template

Parameters **exclusive —** When the scope of a policy is defined as exclusive, the policy can only be applied to a single entity (SAP or SDP). Attempting to assign the policy to a second entity will result in an error message. If the policy is removed from the entity, it will become available for assignment to another entity.

template — When the scope of a policy is defined as template, the policy can be applied to multiple SAPs or network ports.

default-action

Syntax **default-action** {**block** | **allow**}

Context config>service>mrp>mrp-policy

Description This command specifies the action to be applied to the MMRP attributes (Group BMACs) whose ISIDs do not match the specified criteria in all of the entries of the mrp-policy.

When multiple default-action commands are entered, the last command will overwrite the previous command.

Default default-action-allow

Parameters **block —** Specifies that all MMRP attributes will not be declared or registered unless there is a specific mrp-policy entry which causes them to be allowed on this SAP/SDP.

allow — Specifies that all MMRP attributes will be declared and registered unless there is a specific mrp-policy entry which causes them to be blocked on this SAP/SDP.

entry

Syntax [**no**] **entry** *entry-id*

Context config>service>mrp>mrp-policy

Description This command creates or edits an mrp-policy entry. Multiple entries can be created using unique entry-id numbers within the policy. The implementation exits the policy on the first match found and executes the actions in accordance with the accompanying action command. For this reason, entries must be sequenced correctly from most to least explicit. An entry may not have any match criteria defined (in which case, everything matches) but must have at least the keyword action for it to be considered complete. Entries without the action keyword will be considered incomplete and hence will be rendered inactive.

The no form of the command removes the specified entry from the mrp-policy. Entries removed from the mrp-policy are immediately removed from all services where the policy is applied.

The no form of the command removes the specified entry-id.

Default none

Parameters *entry-id* — An entry-id uniquely identifies a match criteria and the corresponding action. It is recommended that multiple entries be given entry-ids in staggered increments. This allows users to insert a new entry in an existing policy without requiring renumbering of all the existing entries.

Values 1-65535

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

match

Syntax [**no**] **match**

Context config>service>mrp>mrp-policy>entry

Description This command creates the context for entering/editing match criteria for the mrp-policy entry. When the match criteria have been satisfied the action associated with the match criteria is executed. In the current implementation just one match criteria (ISID based) is possible in the entry associated with the mrp-policy. Only one match statement can be entered per entry.

The **no** form of the command removes the match criteria for the entry-id.

isid

Syntax [**no**] i**sid** *value* **| from** *value* **to** *higher-value*

Context config>service>mrp>mrp-policy>entry>match

Description This command configures an ISID value or a range of ISID values to be matched by the mrp-policy parent when looking at the related MMRP attributes (Group BMACs). The pbb-etype value for the related SAP (inherited from the ethernet port configuration) or for the related SDP binding (inherited from SDP configuration) will be used to identify the ISID tag.

Multiple isid statements are allowed under a match node. The following rules govern the usage of multiple isid statements:

- overlapping values are allowed:
  - isid from 1 to 10
  - isid from 5 to 15
  - isid 16
- the minimum and maximum values from overlapping ranges are considered and displayed. The above entries will be equivalent with “isid from 1 to 16” statement.
- there is no consistency check with the content of isid statements from other entries. The entries will be evaluated in the order of their IDs and the first match will cause the implementation t o execute the associated action for that entry and then to exit the mrp-policy.
- If there are no isid statements under a match criteria but the mac-filter type is isid the following behaviors apply for different actions:
  - For end-station – it treats any ISID value as no match and goes to next entry or default action which must be “block” in this case
  - For allow – it treats any ISID value as a match and allows it
  - For block – it treats any ISID value as a match and blocks it

The **no** form of the command can be used in two ways:

no isid - removes all the previous statements under one match node

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

Default no isid

Parameters *value or higher-value* — Specifies the ISID value in 24 bits. When just one present identifies a particular ISID to be used for matching.

Values 0..16777215

from *value* **to** *higher-value* — Identifies a range of ISIDs to be used as matching criteria.

action

Syntax **action** {**block** | **allow** | **end-station**}
no action

Context config>service>mrp>mrp-policy>entry

Description This command specifies the action to be applied to the MMRP attributes (Group BMACs) whose ISIDs match the specified ISID criteria in the related entry.

The action keyword must be entered for the entry to be active. Any filter entry without the action keyword will be considered incomplete and will be inactive. If neither keyword is specified (no action is used), this is considered a No-Op policy entry used to explicitly set an entry inactive without modifying match criteria or removing the entry itself. Multiple action statements entered will overwrite previous actions parameters when defined. To remove a parameter, use the no form of the action command with the specified parameter.

The **no** form of the command removes the specified action statement. The entry is considered incomplete and hence rendered inactive without the action keyword.

Default no action

Parameters **block —** Specifies that the matching MMRP attributes will not be declared or registered on this SAP/SDP.

allow — Specifies that the matching MMRP attributes will be declared and registered on this SAP/SDP.

end-station — Specifies that an end-station emulation is present on this SAP/SDP for the MMRP attributes related with matching ISIDs. Equivalent action with the block keyword on that SAP/SDP– the attributes associated with the matching ISIDs do not get declared or registered on the SAP/SDP. The matching attributes on the other hand are mapped as static MMRP entries on the SAP/SDP which implicitly instantiates in the data plane as a MFIB entry associated with that SAP/SDP for the related Group BMAC. For the other SAPs/SDPs in the BVPLS with MRP enabled (no shutdown) this means permanent declaration of the matching attributes, same as in the case when the IVPLS instances associated with these ISIDs were locally configured.

If an mrp-policy has end-station action in one entry, the only default action allowed in the policy is block. Also no other actions are allowed to be configured in other entry configured under the policy.

This policy will apply even if the MRP is shutdown on the local SAP/SDP or for the whole BVPLS to allow for manual creation of MMRP entries in the data plane. Specifically the following rules apply:

– If service vpls mrp shutdown then MMRP on all SAP/SDPs is shutdown - MRP PDUs pass-through transparently

– If service vpls mrp no shutdown and endstation statement (even with no ISID values in the related match statement) is used in a mrp-policy applied to SAP/SDP - no declaration is sent on SAP/SDP. The provisioned ISIDs in the match statement are registered on that SAP/SDP and are propagated on all the other MRP enabled endpoints.

copy

Syntax **copy mrp-policy** *source-name* **to** *dest-name*

Context config>service>mrp

Description This command copies existing mrp-policy list entries for a specific policy name to another policy name. The copy command is a configuration level maintenance tool used to create new mrp-policy using existing mrp-policy.

An error will occur if the destination policy name exists.

Parameters **mrp-policy —** Indicates that source-name and dest-name are MRP policy names.

source-name — Identifies the source mrp-policy from which the copy command will attempt to copy. The mrp-policy with this name must exist for the command to be successful.

dest-name — Identifies the destination mrp-policy to which the copy command will attempt to copy. If the mrp-policy with dest-name exist within the system an error message is generated.

renum

Syntax **renum** *old-entry-id* **to** *new-entry-id*

Context config>service>mrp>mrp-policy

Description This command renumbers existing MRP policy entries to properly sequence policy entries. This may be required in some cases since the implementation exits when the first match is found and executes the actions according to the accompanying action command. This requires that entries be sequenced correctly from most to least explicit.

Parameters *old-entry-id* — Specifies the entry number of an existing entry.

Values 1-65535

new-entry-id — Specifies the new entry number to be assigned to the old entry. If the new entry exists, an error message is generated.

Oper Group Commands

oper-group

Syntax **oper-group** *group-name* [**create**]
no oper-group *group-name*

Context config>service

Description This command creates a system-wide group name which can be used to associate a number of service objects (for example, SAPs or pseudowires). The status of the group is derived from the status of its members. The status of the group can then be used to influence the status of non-member objects. FOr example, when a group status i marked as down, the object(s) that monitor the group change their status accordingly.

The **no** form of the command removes thegroup. All the object associations need to be removed before the no command can be executed.

no oper-group

Parameters *group-name* — specifies the operational group identifier up to 32 characters in length.

create — This keyword is required when first creating the configuration context. Once the context is created, it is possible to navigate into the context without the **create** keyword.

hold-time

Syntax **hold-time**

Context config>service>oper-group

Description This command enables the context to configure hold time information.

group up

Syntax **group up** *time* **| no group up**

Context config>service>oper-group>hold-time

Description This command configures the number of seconds to wait before notifying clients monitoring this group when its operational status transitions from down to up. A value of zero indicates that transitions are reported immediately to monitoring clients. The up time option is a must to achieve fast convergence: when the group comes up, the monitoring MH site which tracks the group status may wait without impacting the overall convergence; there is usually a pair MH site that is already handling the traffic.

The **no** form sets the values back to the defaults.

Default 4

Parameters *time* — Specifies the group up time value.

Values 0 — 3600

group down

Syntax **group down** *time* **| no group down**

Context config>service>oper-group>hold-time

Description This command configures the number of seconds to wait before notifying clients monitoring this group when its operational status transitions from up to down.

The **no** form sets the values back to the default.

Pseudowire Commands

pw-routing

Syntax **pw-routing**

Context config>service

Description This command enables the context to configure dynamic multi-segment pseudowire (MS-PW) routing. Pseudowire routing must be configured on each node that will be a T-PE or an S-PE.

Default disabled

block-on-peer-fault

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

Context config>service>pw-template

Description When enabled, this command blocks the transmit direction of a pseudowire when any of the following pseudowire 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

boot-timer

Syntax **boot-timer** *secs*
no boot-timer

Context config>service>pw-routing

Description This command configures a hold-off timer for MS-PW routing advertisements and signaling and is used at boot time.

The **no** form of this command removes a previously configured timer and restores it to its default.

Default 10

Parameters *timer-value* — The value of the boot timer in seconds.

Values 0 — 600

local-prefix

Syntax **local-prefix** *local-prefix* [**create**]
no local-prefix*local-prefix*

Context config>service>pw-routing

Description This command configures one or more node prefix values to be used for MS-PW routing. At least one prefix must be configured on each node that is an S-PE or a T-PE.

The **no** form of this command removes a previously configured prefix, and will cause the corresponding route to be withdrawn if it has been advertised in BGP.

Default no local-prefix.

Parameters *local-prefix* — Specifies a 32 bit prefix for the AII. One or more prefix values, up to a maximum of 16 may be assigned to the 7x50 node. The global ID can contain the 2-octet or 4-octet value of the provider's Autonomous System Number (ASN). The presence of a global ID based on the provider's ASN ensures that the AII for spoke-SDPs configured on the node will be globally unique.

Values <global-id>:<ip-addr>|<raw-prefix>
ip-addr a.b.c.d
raw-prefix1 — 4294967295
global-id1 — 4294967295

advertise-bgp

Syntax **advertise-bgp route-distinguisher** *rd* [**community** *community*]
no advertise-bgp route-distinguisher *rd*

Context config>service>pw-routing

Description This command enables a given prefix to be advertised in MP-BGP for dynamic MS-PW routing.

The no form of this command will explicitly woithdraw a route if it has been previously advertised.

Default no advertise-bgp.

Parameters *rd* — Specifies an 8-octet route distinguisher associated with the prefix. Up to 4 unique route distinguishers can be configured and advertised for a given prefix though multiple instances of the advertise-bgp command. This parameter is mandatory.

Values (6 bytes, other 2 Bytes of type will be automatically generated)
asn:number1 (RD Type 0): 2bytes ASN and 4 bytes locally administered number
ip-address:number2 (RD Type 1): 4bytes IPv4 and 2 bytes locally administered number;

community community — An optional BGP communities attribute associated with the advertisement. To delete a previously advertised community, advertise-bgp route-distinguisher must be run again with the same value for the RD but excluding the community attribute.

Values	*community*	{2-byte-as-number:comm-va1}
	2-byte-asnumber	0— 65535
	comm.-val	0 — 65535

path

Syntax **path** *name* [**create**]
no path name

Context config>service>pw-routing

Description This command configures an explicit path between this 7x50 T-PE and a remote 7x50 T-PE. For each path, one or more intermediate S-PE hops must be configured. A path can be used by multiple multi-segment pseudowires. Paths are used by a 7x50 T-PE to populate the list of Explicit Route TLVs included in the signaling of a dynamic MS-PW.

A path may specify all or only some of the hops along the route to reach a T-PE.

The **no** form of the command removes a specified explicit path from the configuration.

Default no path

Parameters *path-name* — Specifies a locally-unique case-sensitive alphanumeric name label for the MS-PW path of up to 32 characters in length.

hop

Syntax **hop** *hop-index ip-address*
no hop h*op-index*

Context config>service>pw-routing>hop

Description This command configures each hop on an explicit path that can be used by one or more dynamic MS-PWs. It specifies the IP addresses of the hops that the MS-PE should traverse. These IP addresses can correspond to the system IP address of each S-PE, or the IP address on which the T-LDP session to a given S-PE terminates.

The **no** form of this command deletes hop list entries for the path. All the MS-PWs currently using this path are unaffected. Additionally, all services actively using these MS-PWs are unaffected. The path must be shutdown first in order to delete the hop from the hop list. The ‘no hop hop-index’ command will not result in any action, except for a warning message on the console indicating that the path is administratively up.

Default no hop

Parameters *hop-index* — Specifies a locally significant numeric identifier for the hop. The hop index is used to order the hops specified. The LSP always traverses from the lowest hop index to the highest. The hop index does not need to be sequential.

Values 1 — 1024

ip-address — Specifies the system IP address or terminating IP address for the T-LDP session to the S-PE corresponding to this hop. For a given IP address on a hop, the system will choose the appropriate SDP to use.

retry-count

Syntax **retry-count** [10..10000]
no retry-count

Context config>service>pw-routing

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* — Specifies the maximum number of retries before putting the spoke-sdp into the shutdown state.

Values 10 — 10000

retry-timer

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

Context config>service>pw-routing

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

Values 10 – 480

spe-address

Syntax **spe-address** *global-id:prefix*
no spe-address

Context config>service>pw-routing

Description This command configures a single S-PE Address for the node to be used for dynamic MS-PWs. This value is used for the pseudowire switching point TLV used in LDP signaling, and is the value used by pseudowire status signaling to indicate the PE that originates a pseudowire status message. . Configuration of this parameter is mandatory to enable dynamic MS-PW support on a node.

If the S-PE Address is not configured, spoke-sdps that use dynamic MS-PWs and pw-routing local-prefixes cannot be configured on a T-PE. Furthermore, and 7x50 node will send a label release for any label mappings received for FEC129 AII type 2.

The S-PE Address cannot be changed unless the dynamic ms-pw configuration is removed. Furthermore, changing the S-PE Address will also result in all dynamic MS-PWs for which this node is an S-PE being released. It is recommended that the S-PE Address should be configured for the life of an MS-PW configuration after reboot of the 7x50.

The **no** form of this command removes the configured S-PE Address.

Default no spe-address

Parameters *global-id* — Specifies a 4-octet value that is unique to the service provider. For example, the global ID can contain the 2-octet or 4-octet value of the provider's Autonomous System Number (ASN).

Syntax: <global-id:prefix>:<global-id>:{<prefix>|<ipaddress>}

```
global-id 1 — 4294967295
 prefix    1 — 4294967295
ipaddress a.b.c.d
```

static-route

Syntax [**no**] **static-route** *route-name*

Context config>service>pw-routing

Description This command configures a static route to a next hop S-PE or T-PE. Static routes may be configured on either S-PEs or T-PEs.

A default static route is entered as follows:

```
static-route  0:0:next_hop_ip_addresss
```

or

```
static-route 0:0.0.0.0:next_hop_ip_address
```

The **no** form of this command removes a previously configured static route.

Default no static-route

Parameters *route-name* — Specifies the static pseudowire route.

Values	route-name	<global-id>:<prefix>:<next-hop-ip_addr>
	global-id	0 — 4294967295
	prefix	a.b.c.d | 0— 4294967295
	ip_addr	a.b.c.d

pw-template

Syntax [**no**] **pw-template** *sdp-template-id* [**use-provisioned-sdp**] [**create**]

Context config>service

Description This command configures an SDP template.

Parameters *sdp-template-id* — Specifies a number used to uniquely identify a template for the creation of a Service Distribution Point (SDP. The value 0 is used as the null ID.

Values 0, 1 — 2147483647

use-provisioned-sdp — Specifies whether to use an already provisioned SDP. When specified, the tunnel manager will be consulted for an existing active SDP. Otherwise, the default SDP template will be used to use for instantiation of the SDP.

create — This keyword is required when first creating the configuration context. Once the context is created, it is possible to navigate into the context without the **create** keyword.

SDP Commands

sdp

Syntax **sdp** *sdp-id* [**gre** | **mpls**] [**create**]
no sdp *sdp-id*

Context config>service

Description This command creates or edits a Service Distribution Point (SDP). SDPs must be explicitly configured.

An SDP is a logical mechanism that ties a far-end 7750 SR to a particular service without having to specifically define far end SAPs. Each SDP represents a method to reach a 7750 SR router.

One method is IP Generic Router Encapsulation (GRE) which has no state in the core of the network. GRE does not specify a specific path to the 7750 SR. A GRE-based SDP uses the underlying IGP routing table to find the best next hop to the far end router.

The other method is Multi-Protocol Label Switching (MPLS) encapsulation. A router supports both signaled and non-signaled Label Switched Paths (LSPs) through the network. Non-signaled paths are defined at each hop through the network. Signaled paths are communicated by protocol from end to end using Resource ReserVation Protocol (RSVP). Paths may be manually defined or a constraint-based routing protocol (such as OSPF-TE or CSPF) can be used to determine the best path with specific constraints. An LDP LSP can also be used for an SDP when the encapsulation is MPLS. The use of an LDP LSP type or an RSVP/Static LSP type are mutually exclusive except when the mixed-lsp option is enabled on the SDP.

SDPs are created and then bound to services. Many services may be bound to a single SDP. The operational and administrative state of the SDP controls the state of the SDP binding to the service.

If *sdp-id* does not exist, a new SDP is created. When creating an SDP, either the **gre** or the **mpls** keyword must be specified. SDPs are created in the admin down state (**shutdown**) and the **no shutdown** command must be executed once all relevant parameters are defined and before the SDP can be used.

If *sdp-id* exists, the current CLI context is changed to that SDP for editing and modification. For editing an existing SDP, neither the **gre** nor the **mpls** keyword is specified. If a keyword is specified for an existing *sdp-id*, an error is generated and the context of the CLI will not be changed to the specified *sdp-id*.

The **no** form of this command deletes the specified SDP. Before an SDP can be deleted, it must be administratively down (shutdown) and not bound to any services. If the specified SDP is bound to a service, the **no sdp** command will fail generating an error message specifying the first bound service found during the deletion process. If the specified *sdp-id* does not exist an error will be generated.

Default none

Parameters *sdp-id* — The SDP identifier.

Values 1 — 17407

gre — Specifies the SDP will use GRE to reach the far-end router. Only one GRE SDP can be created to a given destination device. Multiple GRE SDPs to a single destination serve no purpose as the path taken to reach the far end is determined by the IGP which will be the same for all SDPs to a given destination and there is no bandwidth reservation in GRE tunnels.

mpls — Specifies the SDP will use MPLS encapsulation and one or more LSP tunnels to reach the far-end device. Multiple MPLS SDPs may be created to a given destination device . Multiple MPLS SDPs to a single destination device are helpful when they use divergent paths.

auto-learn-mac-protect

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

Context config>service>pw-template
config>service>pw-template>split-horizon-group

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

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

Default auto-learn-mac-protect

accounting-policy

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

Context config>service>pw-template
config>service>sdp

Description This command creates the accounting policy context that can be applied to an SDP. An accounting policy must be defined before it can be associated with a SDP. If the *policy-id* does not exist, an error message is generated.

A maximum of one accounting policy can be associated with a SDP at one time. Accounting policies are configured in the **config>log** context.

The **no** form of this command removes the accounting policy association from the SDP, and the acccounting policy reverts to the default.

Default Default accounting policy.

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

Values 1 — 99

bgp-tunnel

Syntax [no] **bgp-tunnel**

Context config>service>sdp

Description This command allows the use of BGP route tunnels available in the tunnel table to reach SDP far-end nodes.

Use of BGP route tunnels are only available with MPLS-SDP. Only one of the transport methods is allowed per SDP - LDP, RSVP-LSP or BGP-Tunnel (BGP-Tunnel is not supported on multi-mode LSP)

The **no** form of the command disables resolving BGP route tunnel LSP for SDP far-end.

Default no bgp-tunnel (BGP tunnel route to SDP far-end is disabled)

booking-factor

Syntax **booking-factor** *percentage*
no booking-factor

Context config>service>sdp

Description This command specifies the booking factor applied against the maximum SDP available bandwidth 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. A value of 0 means no VLL can be admitted into the SDP.

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

Parameters *percentage* — Specifies the percentage of the SDP maximum available bandwidth for VLL call admission. When the value of this parameter is set to zero (0), no new VLL spoke SDP bindings with non-zero bandwidth are permitted with this SDP. Overbooking, >100% is allowed.

Values 0 — 1000 %

Default 100%

collect-stats

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

Context config>service>pw-template
config>service>sdp

Description This command enables accounting and statistical data collection for either the SDP. 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

control-word

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

Description config>service>pw-template

Description This command enables the use of the control word on pseudowire packets in VPLS and enables the use of the control word individually on each mesh-sdp or spoke-sdp. By default, the control word is disabled. When the control word is enabled, all VPLS packets, including the BPDU frames, are encapsulated with the control word when sent over the pseudowire. The T-LDP control plane behavior is the same as in the implementation of control word for VLL services. The configuration for the two directions of the Ethernet pseudowire should match.

The **no** form of the command reverts the mesh SDP or spoke-sdp to the default behavior of not using the control word.

Default no control-word

disable-aging

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

Context config>service>pw-template

Description This command disables MAC address aging across a service.

The **no** form of this command enables aging.

Default no disable-aging

disable-learning

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

Context config>service>pw-template

Description This command enables learning of new MAC addresses.

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

The **no** form of this command enables learning of MAC addresses.

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

discard-unknown-source

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

Context config>service>pw-template

Description When this command is enabled, packets received with an unknown source MAC address will be dropped only if the maximum number of MAC addresses have been reached.

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

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

Default **no discard-unknown**

egress

Syntax **egress**

Context config>service>pw-template

Description This command enables the context to configure spoke SDP binding egress filter parameters.

ingress

Syntax **ingress**

Context config>service>pw-template

Description This command enables the context to configure spoke SDP binding ingress filter parameters.

filter

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

Context config>service>pw-template>egress
config>service>pw-template>ingress

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

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

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

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

qos

Syntax

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

Context

configure>service>apipe>spoke-sdp>egress
configure>service>cpipe>spoke-sdp>egress
configure>service>epipe>spoke-sdp>egress
configure>service>fpipe>spoke-sdp>egress
configure>service>ipipe>spoke-sdp>egress
config>service>vpls>spoke-sdp>egress
config>service>vpls>mesh-sdp>egress
config>service>pw-template>egress
config>service>vprn>interface>spoke-sdp>egress
config>service>ies>interface>spoke-sdp>egress

Description

This command is used to redirect pseudowire packets to an egress port queue-group for the purpose of shaping.

The egress pseudowire shaping provisioning model allows the mapping of one ore more pseudowires to the same instance of queues, or policers and queues, which are defined in the queue-group template.

Operationally, the provisioning model consists of the following steps:

1. Create an egress queue-group template and configure queues only or policers and queues for each FC that needs to be redirected.
2. Apply the queue-group template to the network egress context of all ports where there exists a network IP interface on which the pseudowire packets can be forwarded. This creates one instance of the template on the egress of the port. One or more instances of the same template can be created.
3. Configure FC-to-policer or FC-to-queue mappings together with the redirect to a queue-group in the egress context of a network QoS policy. No queue-group name is specified in this step, which means the same network QoS policy can redirect different pseudowires to different queue-group templates.
4. Apply this network QoS policy to the egress context of a spoke-SPD inside a service or to the egress context of a pseudowire template and specify the redirect queue-group name.

One or more spoke-SPDs can have their FCs redirected to use queues only or queues and policers in the same queue-group instance.

The following are the constraints and rules of this provisioning model:

1. When a pseudowire FC is redirected to use a queue or a policer and a queue in a queue-group and the queue-group name does not exist, the association is failed at the time the user associates the

egress context of a spoke-SPD to the named queue-group. In such a case, the pseudowire packet will be fed directly to the corresponding egress queue for that FC used by the IP network interface on which the pseudowire packet is forwarded. This queue can be a queue-group queue, or the egress shared queue for that FC defined in the network-queue policy applied to the egress of this port. This is the existing implementation and default behavior for a pseudowire packet.

2. When a pseudowire FC is redirected to use a queue or a policer, and a queue in a queue-group and the queue-group name exists, but the policer-id and/or the queue-id is not defined in the queue-group template, the association is failed at the time the user associates the egress context of a spoke-SPD to the named queue-group. In such a case, the pseudowire packet will be fed directly to the corresponding egress queue for that FC used by the IP network interface the pseudowire packet is forwarded on.
3. When a pseudowire FC is redirected to use a queue, or a policer and a queue in a queue-group, and the queue-group name exists and the policer-id or policer-id plus queue-id exist, it is not required to check that an instance of that queue-group exists in all egress network ports which have network IP interfaces. The handling of this is dealt with in the data path as follows:
   a When a pseudowire packet for that FC is forwarded and an instance of the referenced queue-group name exists on that egress port, the packet is processed by the queue-group policer and will then be fed to the queue-group queue.

   b When a pseudowire packet for that FC is forwarded and an instance of the referenced queue-group name does not exist on that egress port, the pseudowire packet will be fed directly to the corresponding egress shared queue for that FC defined in the network-queue policy applied to the egress of this port.
4. If a network QoS policy is applied to the egress context of a pseudowire, any pseudowire FC, which is not explicitly redirected in the network QoS policy, will have the corresponding packets feed directly the corresponding the egress shared queue for that FC defined in the network-queue policy applied to the egress of this port.

When the queue-group name the pseudowire is redirected to exists and the redirection succeeds, the marking of the packet DEI/dot1.p/DSCP and the tunnel DEI/dot1.p/DSCP/EXP is performed; according to the relevant mappings of the (FC, profile) in the egress context of the network QoS policy applied to the pseudowire. This is true regardless, wether an instance of the queue-group exists or not on the egress port to which the pseudowire packet is forwarded. If the packet profile value changed due to egress child policer CIR profiling, the new profile value is used to mark the packet DEI/dot1.p and the tunnel DEI/dot1.p/EXP, but the DSCP is not modified by the policer operation.

When the queue-group name the pseudowire is redirected does not exist, the redirection command is failed. In this case, the marking of the packet DEI/dot1.p/DSCP and the tunnel DEI/dot1.p/DSCP/EXP fields is performed according to the relevant commands in the egress context of the network QoS policy applied to the network IP interface to which the pseudowire packet is forwarded.

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

Parameters *network-policy-id* — Specifies the network policy identification. The value uniquely identifies the policy on the system.

Values 1 — 65535

queue-redirect-group *queue-group-name* **—** This optional parameter specifies that the *queue-group-name* will be used for all egress forwarding class redirections within the network QoS policy ID. The specified *queue-group-name* must exist as a port egress queue group on the port associated with the IP interface.

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

Values 1 — 16384

hash-label

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

Context config>service>pw-template

Description This command enables the use of the hash label on a VLL, VPRN or VPLS service bound to LDP or RSVP SDP as well as to a VPRN service using the autobind mode with the **ldp**, **rsvp-te**, or **mpls** options. This feature is not supported on a service bound to a GRE SDP or for a VPRN service using the autobind mode with the gre option. This feature is also not supported on multicast packets forwarded using RSVP P2MP LPS or mLDP LSP in both the base router instance and in the multicast VPN (mVPN) instance. It is, however, supported when forwarding multicast packets using an IES/VPRN spoke-interface.

When this feature is enabled, the ingress data path is modified such that the result of the hash on the packet header is communicated to the egress data path for use as the value of the label field of the hash label. The egress data path appends the hash label at the bottom of the stack (BoS) and sets the S-bit to one (1).

In order to allow applications where the egress LER infers the presence of the hash label implicitly from the value of the label, the Most Significant Bit (MSB) of the result of the hash is set before copying into the Hash Label. This means that the value of the hash label will always be in the range [524,288 - 1,048,575] and will not overlap with the signaled/static LSP and signaled/static service label ranges. This also guarantees that the hash label will not match a value in the reserved label range.

The (unmodified) result of the hash continues to be used for the purpose of ECMP and LAG spraying of packets locally on the ingress LER. Note, however, that for VLL services, the result of the hash is overwritten and the ECMP and LAG spraying will be based on service-id when ingress SAP shared queuing is not enabled. However, the hash label will still reflect the result of the hash such that an LSR can use it to perform fine grained load balancing of VLL pseudowire packets.

Packets generated in CPM and that are forwarded labeled within the context of a service (for example, OAM packets) must also include a Hash Label at the BoS and set the S-bit accordingly.

The TTL of the hash label is set to a value of 0.

The user enables the signaling of the hash-label capability under a VLL spoke-sdp, a VPLS spoke-sdp or mesh-sdp, or an IES/VPRN spoke interface by adding the **signal-capability** option. In this case, the decision whether to insert the hash label on the user and control plane packets by the local PE is solely determined by the outcome of the signaling process and can override the local PE configuration. The following are the procedures:

- The 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 pseudowire but must not insert the hash label in the user and control packets over that spoke-sdp or

mesh-sdp. If the remote PE does not support the **signal-capability** option, then there are a couple of possible outcomes:

  - If the **hash-label** option was enabled on the local configuration of the spoke-sdp or mesh-sdp at the remote PE, the pseudowire packets received by the local PE will have the hash label included. These packets must be dropped. The only way to solve this is to disable the signaling capability option on the local node which will result in the insertion of the hash label by both PE nodes.
  - If the **hash-label** option is not supported or was not enabled on the local configuration of the spoke-sdp or mesh-sdp at the remote PE, the pseudowire received by the local PE will not have the hash label included.

- The user can enable or disable the signal-capability option in CLI as needed. When doing so, the router 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.

force-vlan-vc-forwarding

Syntax [**no**] **force-vlan-vc-forwarding**

Context config>service>pw-template

Description This command forces vc-vlan-type forwarding in the data path for spoke and mesh SDPs that 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

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 pseudowire packets to an ingress forwarding plane queue-group for the purpose of rate-limiting.

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

Operationally, the provisioning model in the case of the ingress pseudowire shaping feature consists of the following steps:

1. Create an ingress queue-group template and configure policers for each FC which 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 which the pseudowire 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 pseudowires to different queue-group templates.
4. Apply this network QoS policy to the ingress context of a spoke-sdp inside a service or to the ingress context of a pseudowire template and specify the redirect queue-group name.

One or more spoke-sdps can have their FCs redirected to use policers in the same policer queue-group instance.

The following are the constraints and rules of this provisioning model when used in the ingress pseudowire rate-limiting feature:

1. When a pseudowire FC is redirected to use a policer in a named policer queue-group and the queue-group name does not exist, the association is failed at the time the user associates the ingress context of a spoke-sdp to the named queue-group. In such a case, the pseudowire packet will feed directly the ingress network shared queue for that FC defined in the network-queue policy applied to the ingress of the MDA/FP.
2. When a pseudowire FC is redirected to use a policer in a named policer queue-group and the queue-group name exists but the policer-id is not defined in the queue-group template, the association is failed at the time the user associates the ingress context of a spoke-sdp to the named queue-group. In such a case, the pseudowire packet will feed directly the ingress network shared queue for that FC defined in the network-queue policy applied to the ingress of the MDA/FP.
3. When a pseudowire FC is redirected to use a policer in a named policer queue-group and the queue-group name exists and the policer-id is defined in the queue-group template, it is not required to check that an instance of that queue-group exists in all ingress FPs which have network IP interfaces. The handling of this is dealt with in the data path as follows:
   - When a pseudowire packet for that FC is received and an instance of the referenced queue-group name exists on that FP, the packet is processed by the policer and will then feed the per-FP ingress shared queues referred to as “policer-output-queues”.
   - When a pseudowire packet for that FC is received and an instance of the referenced queue-group name does not exist on that FP, the pseudowire packets will be fed directly into the corresponding ingress network shared queue for that FC defined in the network-queue policy applied to the ingress of the MDA/FP.
4. If a network QoS policy is applied to the ingress context of a pseudowire, any pseudowire FC which is not explicitly redirected in the network QoS policy will have the corresponding packets

feed directly the ingress network shared queue for that FC defined in the network-queue policy applied to the ingress of the MDA/FP.

5. If no network QoS policy is applied to the ingress context of the pseudowire, then all packets of the pseudowire will feed:
   - 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 a IES/VPRN spoke interface and from a R-VPLS spoke-sdp which 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 pseudowire is redirected to use a policer queue-group, the classification of the packet for the purpose of FC and profile determination is performed according to default classification rule or the QoS filters defined in the ingress context of the network QoS policy applied to the pseudowire. This is true regardless if an instance of the named policer queue-group exists on the ingress FP the pseudowire 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 pseudowire terminates in IES or VPRN service (spoke-interface).

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

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

Parameters *network-policy-id* — Specifies the network policy identification. The value uniquely identifies the policy on the system.

Values 1 — 65535

fp-redirect-group *queue-group-name* — Specifies the network policy identification. The value uniquely identifies the policy on the system.

Values 1 — 16384

vc-label

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

Context config>service>pw-template>ingress

Description This command configures the ingress VC label.

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

Values 2048 — 18431

limit-mac-move

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

Context config>service>pw-template

Description This command indicates whether or not the mac-move agent will limit the MAC re-learn (move) rate.

Default **blockable**

Parameters **blockable —** The agent will monitor the MAC re-learn rate, and it will block it when the re-learn rate is exceeded.

non-blockable — When specified, a SAP will not be blocked, and another blockable SAP will be blocked instead.

mac-pinning

Syntax [**no**] **mac-pinning**

Context config>service>pw-template

Description Enabling this command will disable re-learning of MAC addresses on other SAPs within the service. The MAC address will remain attached to a given SAP for duration of its age-timer.

The age of the MAC address entry in the FIB is set by the age timer. If **mac-aging** is disabled on a given VPLS service, any MAC address learned on a SAP/SDP with **mac-pinning** enabled will remain in the FIB on this SAP/SDP forever. Every event that would otherwise result in re-learning will be logged (MAC address; original-SAP; new-SAP).

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

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

max-nbr-mac-addr

Syntax **max-nbr-mac-addr** *table-size*
no max-nbr-mac-addr

Context config>service>pw-template

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

When the configured limit has been reached, and discard-unknown-source has been enabled for this SAP or spoke SDP (see discard-unknown-source on page 140), packets with unknown source MAC addresses will be discarded.

The **no** form of the command restores the global MAC learning limitations for the SAP or spoke SDP.

Default no max-nbr-mac-addr

Parameters *table-size* — Specifies the maximum number of learned and static entries allowed in the FDB of this service.

Values 1 — 196607
The chassis-mode C limit: 511999

restrict-protected-src

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

Context config>service>pw-template
config>service>pw-template>split-horizon-group

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

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

Default restrict-protected-src

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

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

mfib-allowed-mda-destinations

Syntax **mfib-allowed-mda-destinations**

Context config>service>pw-template>egress

Description This command enables the context to configure MFIB-allowed MDA destinations.

The allowed-mda-destinations node and the corresponding **mda** command are used on spoke and mesh SDP bindings to provide a list of MDA destinations in the chassis that are allowed as destinations for multicast streams represented by [*,g] and [s,g] multicast flooding records on the VPLS service. The MDA list only applies to IP multicast forwarding when IGMP snooping is enabled on the VPLS service. The MDA list has no effect on normal VPLS flooding such as broadcast, Layer 2 multicast, unknown destinations or non-snooped IP multicast.

At the IGMP snooping level, a spoke or mesh SDP binding is included in the flooding domain for an IP multicast stream when it has either been defined as a multicast router port, received a IGMP query through the binding or has been associated with the multicast stream through an IGMP request by a host over the binding. Due to the dynamic nature of the way that a spoke or mesh SDP binding is associated with one or more egress network IP interfaces, the system treats the binding as appearing on all network ports. This causes all possible network destinations in the switch fabric to be included in the multicast streams flooding domain. The MDA destination list provides a simple mechanism that narrows the IP multicast switch fabric destinations for the spoke or mesh SDP binding.

If no MDAs are defined within the allowed-mda-destinations node, the system operates normally and will forward IP multicast flooded packets associated with the spoke or mesh SDP binding to all switch fabric taps containing network IP interfaces.

The MDA inclusion list should include all MDAs that the SDP binding may attempt to forward through. A simple way to ensure that an MDA that is not included in the list is not being used by the binding is to define the SDP the binding is associated with as MPLS and use an RSVP-TE LSP with a strict egress hop. The MDA associated with the IP interface defined as the strict egress hop should be present in the inclusion list.

If the inclusion list does not currently contain the MDA that the binding is forwarding through, the multicast packets will not reach the destination represented by the binding. By default, the MDA inclusion list is empty.

If an MDA is removed from the list, the MDA is automatically removed from the flooding domain of any snooped IP multicast streams associated with a destination on the MDA unless the MDA was the last MDA on the inclusion list. Once the inclusion list is empty, all MDAs are eligible for snooped IP multicast flooding for streams associated with the SDP binding.

mda

Syntax [**no**] **mda** *mda-id*

Context config>service>pw-template>egress>mfib-mda

Description This command specifies an MFIB-allowed MDA destination for an SDP binding configured in the system.

Parameters *mda-id* — Specifies an MFIB-allowed MDA destination.

Values 1, 2

igmp-snooping

Syntax **igmp-snooping**

Context config>service>pw-template

Description This command enables the Internet Group Management Protocol (IGMP) snooping context.

Default none

fast-leave

Syntax [**no**] **fast-leave**

Context config>service>pw-template>igmp-snooping

Description This command enables fast leave.

When IGMP fast leave processing is enabled, the SR-Series will immediately remove a SAP or SDP from the IP multicast group when it detects an IGMP 'leave' on that SAP or SDP. Fast leave processing allows the switch to remove a SAP or SDP that sends a 'leave' from the forwarding table without first sending out group-specific queries to the SAP or SDP, and thus speeds up the process of changing channels ('zapping').

Fast leave should only be enabled when there is a single receiver present on the SAP or SDP.

When fast leave is enabled, the configured last-member-query-interval value is ignored.

Default no fast-leave

import

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

Context config>service>pw-template>igmp-snooping

Description This command specifies the import routing policy to be used for IGMP packets. Only a single policy can be imported at a time.

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

Default **no import** — No import policy is specified.

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

last-member-query-interval

Syntax **last-member-query-interval** *tenths-of-seconds*
no last-member-query-interval

Context config>service>pw-template>igmp-snooping

Description This command configures the maximum response time used in group-specific queries sent in response to ‘leave’messages, and is also the amount of time between 2 consecutive group-specific queries. This value may be tuned to modify the leave latency of the network. A reduced value results in reduced time to detect the loss of the last member of a group.

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

Default 10

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

Values 1 — 50

max-num-groups

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

Context config>service>pw-template>igmp-snooping

Description This command defines the maximum number of multicast groups that can be joined. If the SR-Series receives an IGMP join message that would exceed the configured number of groups, the request is ignored.

Default no max-num-groups

Parameters *count* — Specifies the maximum number of groups that can be joined.

Values 1 — 1000

query-interval

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

Context config>service>pw-template>igmp-snooping

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

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

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

Default 125

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

Values 2 — 1024

query-response-interval

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

Context config>service>pw-template>igmp-snooping

Description This command configures the IGMP query response interval. If the **send-queries** command is enabled, this parameter specifies the maximum response time advertised in IGMPv2/v3 queries.

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

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

Default 10

Parameters *seconds* — Specifies the length of time to wait to receive a response to the host-query message from the host.

Values 1 — 1023

robust-count

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

Context config>service>pw-template>igmp-snooping

Description If the **send-queries** command is enabled, this parameter allows tuning for the expected packet loss. The robust-count variable allows tuning for the expected packet loss on a subnet and is comparable to a retry count.

If send-queries is not enabled, this parameter will be ignored.

Default 2

Parameters *robust-count* — Specifies the robust count for the SAP or SDP.

Values 2 — 7

send-queries

Syntax [**no**] **send-queries**

Context config>service>pw-template>igmp-snooping

Description This command specifies whether to send IGMP general query messages.

When send-queries is configured, all type of queries generate ourselves are of the configured version. If a report of a version higher than the configured version is received, the report will get dropped and a new wrong version counter will get incremented.

If send-queries is not configured, the version command has no effect. The version used on that SAP/SDP will be the version of the querier. This implies that, for example, when we have a v2 querier, we will never send out a v3 group or group-source specific query when a host wants to leave a certain group.

Default no send-queries

version

Syntax **version** *version*
no version

Context config>service>pw-template>igmp-snooping

Description This command specifies the version of IGMP. This object can be used to configure a router capable of running either value. For IGMP to function correctly, all routers on a LAN must be configured to run the same version of IGMP on that LAN.

When the **send-query** command is configured, all type of queries generate ourselves are of the configured **version**. If a report of a version higher than the configured version is received, the report gets dropped and a new “wrong version” counter is incremented.

If the **send-query** command is not configured, the **version** command has no effect. The version used on that

SAP or SDP will be the version of the querier. This implies that, for example, when there is a v2 querier, a v3 group or group-source specific query when a host wants to leave a certain group will never be sent.

Parameters *version* — Specify the IGMP version.

Values 1, 2, 3

sdp-include

Syntax [**no**] **sdp-include** *group-name*

Context config>service>pw-template

Description This command configures SDP admin group constraints for a PW template.

The admin group name must have been configured or the command is failed. The user can execute the command multiple times to include or exclude more than one admin group. The sdp-include and sdp-exclude commands can only be used with the **use-provisioned-sdp** option. If the same group name is included and excluded within the same PW template, only the exclude option will be enforced.

Any changes made to the admin group sdp-include and sdp-exclude constraints will only be reflected in existing spoke-sdps after the following command has been executed:

tools>perform>service>eval-pw-template>allow-service-impact

When the service is bound to the PW template, the SDP selection rules will enforce the admin group constraints specified in the sdp-include and sdp-exclude commands.

In the SDP selection process, all provisioned SDPs with the correct far-end IP address, the correct tunnel-far-end IP address, and the correct service label signaling are considered. The SDP with the lowest admin metric is selected. If more then one SDP with the same lowest metric are found then the SDP with the highest sdp-id is selected. The type of SDP, GRE or MPLS (BGP/RSVP/LDP) is not a criterion in this selection.

The selection rule with SDP admin groups is modified such that the following admin-group constraints are applied upfront to prune SDPs that do not comply:

- if one or more **sdp-include** statement is part of the pw-template, then an SDP that is a member of one or more of the included groups will be considered. With the **sdp-include** statement, there is no preference for an SDP that belongs to all included groups versus one that belongs to one or fewer of the included groups. All SDPs satisfying the admin-group constraint will be considered and the selection above based on the lowest metric and highest sdp-id is applied.
- if one or more **sdp-exclude** statement is part of the pw-template, then an sdp that is a member of any of the excluded groups will not be considered.

SDP admin group constraints can be configured on all 7x50 services that makes use of the PW template (i.e., BGP-AD VPLS service, BGP-VPLS service, and FEC129 VLL service). In the latter case, only support at a T-PE node is provided.

The **no** form of this command removes the SDP admin group constraints from the PW template.

Default none

Parameters *group-name* — Specifies the name of the SDP admin group. A maximum of 32 characters can be entered.

sdp-exclude

Syntax [**no**] **sdp-exclude** *group-name*

Context config>service>pw-template

Description This command configures SDP admin group constraints for a PW template.

The admin group name must have been configured or the command is failed. The user can execute the command multiple times to include or exclude more than one admin group. The sdp-include and sdp-exclude commands can only be used with the use-provisioned-sdp option. If the same group name is included and excluded within the same PW template, only the exclude option will be enforced.

Any changes made to the admin group sdp-include and sdp-exclude constraints will only be reflected in existing spoke-sdps after the following command has been executed:

tools>perform>service>eval-pw-template>allow-service-impact

When the service is bound to the PW template, the SDP selection rules will enforce the admin group constraints specified in the sdp-include and sdp-exclude commands.

In the SDP selection process, all provisioned SDPs with the correct far-end IP address, the correct tunnel-far-end IP address, and the correct service label signaling are considered. The SDP with the lowest admin metric is selected. If more then one SDP with the same lowest metric are found then the SDP with the highest sdp-id is selected. The type of SDP, GRE or MPLS (BGP/RSVP/LDP) is not a criterion in this selection.

The selection rule with SDP admin groups is modified such that the following admin-group constraints are applied upfront to prune SDPs that do not comply:

- if one or more **sdp-include** statement is part of the pw-template, then an SDP that is a member of one or more of the included groups will be considered. With the **sdp-include** statement, there is no preference for an SDP that belongs to all included groups versus one that belongs to one or fewer of the included groups. All SDPs satisfying the admin-group constraint will be considered and the selection above based on the lowest metric and highest sdp-id is applied.
- if one or more **sdp-exclude** statement is part of the pw-template, then an sdp that is a member of any of the excluded groups will not be considered.

SDP admin group constraints can be configured on all 7x50 services that makes use of the PW template (i.e., BGP-AD VPLS service, BGP-VPLS service, and FEC129 VLL service). In the latter case, only support at a T-PE node is provided.

The **no** form of this command removes the SDP admin group constraints from the PW template.

Default none

Parameters *group-name* — Specifies the name of the SDP admin group. A maximum of 32 characters can be entered.

split-horizon-group

Syntax [**no**] **split-horizon-group** [*group-name*] [*residential-group*]

Context config>service>pw-template

Description This command creates a new split horizon group (SGH).

Comparing a “residential” SGH and a “regular” SHG is that a residential SHG:

- Has different defaults for the SAP/SDP that belong to this group (ARP reply agent enabled (SAP only), MAC pinning enabled). These can be disabled in the configuration.
- Does not allow enabling spanning tree (STP) on a SAP. It is allowed on an SDP.
- Does not allow for downstream broadcast (broadcast / unknown unicast) on a SAP. It is allowed on an SDP.
- On a SAP, downstream multicast is only allowed when IGMP is enabled (for which an MFIB state exists; only IP multicast); on a SDP, downstream mcast is allowed.

When the feature was initially introduced, residential SHGs were also using ingress shared queing by default to increase SAP scaling.

A residential SAP (SAP that belongs to a RSHG) is used to scale the number of SAPs in a single VPLS instance. The limit depends on the hardware used and is higher for residential SAPs (where there is no need for egress multicast replication on residential SAPs) than for regular SAPs. Therefore, residential SAPs are usefull in residential aggregation environments (for example, triple play networks) with a VLAN/subscriber model.

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

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

residential-group — Defines a split horizon group as a residential split horizon group (RSHG). Doing so entails that:

- SAPs which are members of this Residential Split Horizon Group will have:
  - → Double-pass queuing at ingress as default setting (can be disabled)
  - → STP disabled (cannot be enabled)
  - → ARP reply agent enabled per default (can be disabled)
  - → MAC pinning enabled per default (can be disabled)
  - → Downstream Broadcast packets are discarded thus also blocking the unknown, flooded traffic
  - → Downstream Multicast packets are allowed when IGMP snooping is enabled
- Spoke SDPs which are members of this Residential Split Horizon Group will have:
  - → Downstream multicast traffic supported
  - → Double-pass queuing is not applicable
  - → STP is disabled (can be enabled)
  - → ARP reply agent is not applicable (dhcp-lease-states are not supported on spoke SDPs)
  - → MAC pinning enabled per default (can be disabled)

Default A split horizon group is by default not created as a residential-group.

auto-learn-mac-protect

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

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

config>service>pw-template
config>service>pw-template>split-horizon-group

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

When the auto-learn-mac-protect is enabled on an SHG the action only applies to the associated SAPs (no action is taken by default for spoke SDPs in the SHG). In order to enable this function for spoke SDPs within a SHG, the auto-learn-mac-protect must be enabled explicitly under the spoke-SDP. If required, auto-learn-mac-protect can also be enabled explicitly under specific SAPs within the SHG. For more information about auto-learn MAC protect, refer to Auto-Learn MAC Protect on page 616.

Default no auto-learn-mac-protect

restrict-protected-src

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

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

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

- No parameter

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

- alarm-only

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

- discard-frame

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

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

restrict-protected-src can also be enabled explicitly under specific SAPs within the SHG.

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

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

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

Default no alarm-only

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

Default no discard-frame

Default no restrict-protected-src

restrict-unprotected-dst

Syntax **restrict-unprotected-dst** *alarm-only*
no restrict-unprotected-dst

Context config>service>pw-template>split-horizon-group
config>service>vpls>split-horizon-group
config>service>vpls>sap

Description This command indicates how the system will forward packets destined to an unprotected MAC address, either manually added using the mac-protect command or automatically added using the auto-learn-mac-protect command. While enabled all packets entering the configured SAP or SAPs within a split-horizon-group (but not spoke or mesh-SDPs) will be verified to contain a protected destination MAC address. If the packet is found to contain a non-protected destination MAC, it will be discarded. Detecting a non-protected destination MAC on the SAP will not cause the SAP to be placed in the operationally down state. No alarms are generated.

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

Default no restrict-unprotected-dst

vc-type

Syntax **vc-type** {**ether** | **vlan**}

Context config>service>pw-template

Description This command overrides the default VC type signaled for the binding to the far end 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.

Parameters

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

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

vlan-vc-tag

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

Context config>service>pw-template

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.

adv-mtu-override

Syntax [**no**] **adv-mtu-override**

Context config>service>sdp

Description This command overrides the advertised VC-type MTU of all spoke-sdp's of L2 services using this SDP-ID. When enabled, the router signals a VC MTU equal to the service MTU, which includes the Layer 2 header. It also allows this router to accept an MTU advertized by the far-end PE which value matches either its advertised MTU or its advertised MTU minus the L2 headers.

By default, the router advertizes a VC-MTU equal to the L2 service MTU minus the Layer 2 header and always matches its advertized MTU to that signaled by the far-end PE router, otherwise the spoke-sdp goes operationally down.

When this command is enabled on the SDP, it has no effect on a spoke-sdp of an IES/VPRN spoke interface using this SDP-ID. The router continues to signal a VC MTU equal to the net IP interface MTU, which is min{ip-mtu, sdp operational path mtu - L2 headers}. The router also continues to make sure that the advertized MTU values of both PE routers match or the spoke-sdp goes operationally down.

The **no** form of the command disables the VC-type MTU override and returns to the default behavior.

Default no adv-mtu-override

binding

Syntax **binding**

Context config>service>sdp

Description The command enables the context to configure SDP bindings.

port

Syntax **port** [*port-id* | *lag-id*]
no ort

Context config>service>sdp>binding

Description This command specifies the port or lag identifier, to which the pseudowire ports associated with the underlying SDP are bound. If the underlying SDP is re-routed to a port or lag other than the specified one, the pseudowire ports on the SDP are operationally brought down.

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

Default none

Parameters *port-id* — The identifier of the port in the slot/mda/port format.

lag-id — Specifies the LAG identifier.

pw-port

Syntax **pw-port** *pw-port-id* [**vc-id** *vc-id*] [**create**]
no pw-port

Context config>service>sdp>binding

Description This command creates a pseudowire port.

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

Default none

Parameters *pw-port-id* — Specifies a unique identifier of the pseudowire port.

Values 1 — 10239

vc-id *vc-id* **—** Specifies a virtual circuit identifier signaled to the peer.

Values 1 — 4294967295

encap-type

Syntax **encap-type** {**dot1q**|**qinq**}
no encap-type

Context config>service>sdp>binding>pw-port

Description This command sets the encapsulation type for the pseudowire port as dot1q or qinq.

Default dot1q

Parameters **dot1q —** Specifies **dot1q** encapsulation type.

qinq — Specifies **qinq** encapsulation type.

vc-type

Syntax **vc-type** {**ether**|**vlan**}
no vc-type

Context config>service>sdp>binding>pw-port

Description This command sets the forwarding mode for the pseudowirepseudowireport. The vc-type is signaled to the peer, and must be configured consistently on both ends of the pseudowire. vc-type VLAN is only configurable with dot1q encapsulation on the pseudowire port. The tag with vc-type vlan only has significance for transport, and is not used for service delineation or ESM. The top (provider tag) is stripped while forwarding out of the pseudowire, and a configured vlan-tag (for vc-type vlan) is inserted when forwarding into the pseudowire. With vc-type ether, the tags if present (max 2), are transparently preserved when forwarding in our out of the pseudowire.

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

Default ether

Parameters **ether —** Specifies **ether** as the virtual circuit (VC) associated with the SDP binding.

vlan — Specifies **vlan** as the virtual circuit (VC) associated with the SDP binding.

vlan-vc-tag

Syntax **vlan-vc-tag** *vlan-id*
no vc-type

Context config>service>sdp>binding>pw-port

Description This command sets tag relevant for vc-type vlan mode. This tag is inserted in traffic forwarded into the pseudowire.

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

Default 0

Parameters *vlan-id* — Specifies the VLAN ID value.

Values 0 — 4094

egress

Syntax **egress**

Context config>service>sdp>binding>pw-port

Description This command enters egress configuration context for the vport.

Default none

shaper

Syntax [**no**] **shaper**

Context config>service>sdp>binding>pw-port>egress

Description This command configures an egress shaping option for use by a pseudowire port.

Default no shaper.

class-forwarding

Syntax **class-forwarding** [**default-lsp** *lsp-name*]
no class-forwarding

Context config>service>sdp

Description This command enables the forwarding of a service packet over the SDP based on the class of service of the packet. Specifically, the packet is forwarded on the RSVP LSP or static LSP whose forwarding class matches that of the packet. The user maps the system forwarding classes to LSPs using the **config>service>sdp>class-forwarding>fc** command. If there is no LSP that matches the packet's forwarding class, the default LSP is used. If the packet is a VPLS multicast/broadcast packet and the user did not explicitly specify the LSP to use under the **config>service>sdp>class-forwarding>multicast-lsp** context, then the default LSP is used.

VLL service packets are forwarded based on their forwarding class only if shared queuing is enabled on the ingress SAP. Shared queuing must be enabled on the VLL ingress SAP if class-forwarding is enabled on the SDP the service is bound to. Otherwise, the VLL packets will be forwarded to the LSP which is the result of hashing the VLL service ID. Since there are eight entries in the ECMP table for an SDP, one LSP ID for each forwarding class, the resulting load balancing of VLL service ID is weighted by the number of times an LSP appears on that table. For instance, if there are eight LSPs, the result of the hashing will be similar to when class based forwarding is disabled on the SDP. If there are fewer LSPs, then the LSPs which were mapped to more than one forwarding class, including the default LSP, will have proportionally more VLL services forwarding to them.

Class-based forwarding is not supported on a spoke SDP used for termination on an IES or VPRN service. All packets are forwarded over the default LSP.

The **no** form of the command deletes the configuration and the SDP reverts back to forwarding service packets based on the hash algorithm used for LAG and ECMP.

Default **no class-forwarding** — Packets of a service bound to this SDP will be forwarded based on the hash algorithm used for LAG and ECMP.

Parameters **default-lsp** *lsp-name* **—** Specifies the default LSP for the SDP. This LSP name must exist and must have been associated with this SDP using the *lsp-name* configured in the **config>service>sdp>lsp** context. The default LSP is used to forward packets when there is no available LSP which matches the packet's forwarding class. This could be because the LSP associated with the packet's forwarding class is down, or that the user did not configure a mapping of the packet's forwarding class to an LSP using the **config>service>sdp>class-forwarding>fc** command. The default LSP is also used to forward VPLS service multicast/broadcast packets in the absence of a user configuration indicating an explicit association to one of the SDP LSPs.

Note that when the default LSP is down, the SDP is also brought down. The user will not be able to enter the class-forwarding node if the default LSP was not previously specified. In other words the class-forwarding for this SDP will remain shutdown.

enforce-diffserv-lsp-fc

Syntax [**no**] **enforce-diffserv-lsp-fc**

Context config>service>sdp>class-forwarding

Description This command enables checking by RSVP that a Forwarding Class (FC) mapping to an LSP under the SDP configuration is compatible with the Diff-Serv Class Type (CT) configuration for this LSP.

When the user enables this option, the service manager enquires with RSVP if the FC is supported by the LSP. RSVP checks if the FC maps to the CT of the LSP, for example, the default class-type value or the class-type value entered at the LSP configuration level.

If RSVP did not validate the FC, then the service manager will return an error and the check has failed. In this case, packets matching this FC will be forwarded over the default LSP. Any addition of an LSP to an SDP that will not satisfy the FC check will also be rejected.

The service manager does no validate the default-lsp FC-to-CT mapping. Whether or not the FC is validated, the default-lsp will always end up being used in this case.

RSVP will not allow the user to change the CT of the LSP until no SDP with class-based forwarding enabled and the **enforce-diffserv-lsp-fc** option enabled is using this LSP. All other SDPs using this LSP are not concerned by this rule.

The SDP will continue to enforce the mapping of a single LSP per FC. However, when **enforce-diffserv-lsp-fc** enabled, RSVP will also enforce the use of a single CT per FC as per the user configured mapping in RSVP.

If class-forwarding is enabled but **enforce-diffserv-lsp-fc** is disabled, forwarding of the service packets will continue to be based on the user entered mapping of FC to LSP name without further validation as per the existing implementation. The CT of the LSP does not matter in this case.

If class-forwarding is not enabled on the SDP, forwarding of the service packets will continue to be based on the ECMP/LAG hash routine. The CT of the LSP does not matter in this case.

The **no** form of this command reverts to the default value which is to use the user entered mapping of FC to LSP name.

Default no enforce-diffserv-lsp-fc

far-end

Syntax **far-end** *ip-address* | {**node-id** *node-id* [**global-id** *global-id*]}
no far-end

Context config>service>sdp

Description This command configures the system IP address of the far-end destination router for the Service Distribution Point (SDP) that is the termination point for a service.

The far-end IP address must be explicitly configured. The destination IP address must be a 7750 SR system IP address.

If the SDP uses GRE for the destination encapsulation, the *ip-address* is checked against other GRE SDPs to verify uniqueness. If the *ip-address* is not unique within the configured GRE SDPs, an error is generated and the *ip-address* is not associated with the SDP. The local device may not know whether the *ip-address* is actually a system IP interface address on the far end device.

If the SDP uses MPLS encapsulation, the **far-end** *ip-address* is used to check LSP names when added to the SDP. If the "**to** IP address" defined within the LSP configuration does not exactly match the SDP **far-end** *ip-address*, the LSP will not be added to the SDP and an error will be generated. Alternatively, and SDP that uses MPLS can have an MPLS-TP node with an MPLS-TP node-id and (optioanlly) global-id. In this case, the SDP must use an MPLS-TP LSP and the SDP **signaling** parameter must be set to **off**.

An SDP cannot be administratively enabled until a **far-end** *ip-address* or MPLS-TP node-id is defined. The SDP is operational when it is administratively enabled (**no shutdown**) and the **far-end** *ip-address* is contained in the IGP routing table as a host route. OSPF ABRs should not summarize host routes between areas. This can cause SDPs to become operationally down. Static host routes (direct and indirect) can be defined in the local dev ice to alleviate this issue.

The **no** form of this command removes the currently configured destination IP address for the SDP. The *ip-address* parameter is not specified and will generate an error if used in the **no far-end** command. The SDP must be administratively disabled using the **config service sdp shutdown** command before the **no far-end** command can be executed. Removing the far end IP address will cause all *lsp-name* associations with the SDP to be removed.

Default none

Parameters *ip-address* — The system address of the far-end 7750 SR for the SDP in dotted decimal notation.

node-id *node-id* **—** The MPLS-TP Node ID of the far-end system for the SDP, either in dotted decimal notaion (a.b.c.d) or an unsigned 32-bit integer (1 – 4294967295). This parameter is mandatory for an SDP using an MPLS-TP LSP.

global-id *global-id* **—** The MPLS-TP Global ID of the far-end system for the SDP, in an unsigned 32-bit integer (0 – 4294967295). This parameter is optonal for an SDP using an MPLS-TP LSP. If note entered, a default value for the Global ID of '0' is used. A global ID of '0' indicates that the far end node is in the same domain as the local node. The user must explicitly configure a Global ID if its value is non-zero.

fc

Syntax **fc** {**be** | **l2** | **af** | **l1** | **h2** | **ef** | **h1** | **nc**} **lsp** *lsp-name*
no fc {**be** | **l2** | **af** | **l1** | **h2** | **ef** | **h1** | **nc**}

Context config>service>sdp>forwarding-class

Description This command makes an explicit association between a forwarding class and an LSP. The LSP name must exist and must have been associated with this SDP using the command config>service>sdp>lsp. Multiple forwarding classes can be associated with the same LSP. However, a forwarding class can only be associated with a single LSP in a given SDP. All subclasses will be assigned to the same LSP as the parent forwarding class.

Default none

Parameters **lsp** *lsp-name* — Specifies the RSVP or static LSP to use to forward service packets which are classified into the specified forwarding class.

multicast-lsp

Syntax **multicast-lsp** *lsp-name*
no multicast-lsp

Context config>service>sdp>forwarding-class

Description This command specifies the RSVP or static LSP in this SDP to use to forward VPLS multicast and broadcast packets. The LSP name must exist and must have been associated with this SDP using the command config>service>sdp>lsp. In the absence of an explicit configuration by the user, the default LSP is used.

Default default-lsp-name

ldp

Syntax [**no**] **ldp**

Context config>service>sdp

Description This command enables LDP-signaled LSP's on MPLS-encapsulated SDPs.

In MPLS SDP configurations *either* one LSP can be specified *or* LDP can be enabled. The SDP **ldp** and **lsp** commands are mutually exclusive. If an LSP is specified on an MPLS SDP, then LDP cannot be enabled on the SDP. To enable LDP on the SDP when an LSP is already specified, the LSP must be removed from the configuration using the **no lsp *lsp-name*** command.

Alternatively, if LDP is already enabled on an MPLS SDP, then an LSP cannot be specified on the SDP. To specify an LSP on the SDP, the LDP must be disabled. The LSP must have already been created in the **config>router>mpls** context with a valid far-end IP address. The above rules are relaxed when the mixed-lsp option is enabled on the SDP.

Default no ldp (disabled)

lsp

Syntax **lsp** *lsp-name*
no lsp *lsp-name*

Context config>service>sdp

Description This command creates associations between one or more label switched paths (LSPs) and an Multi-Protocol Label Switching (MPLS) Service Distribution Point (SDP). This command is implemented *only* on MPLS-type encapsulated SDPs.

In MPLS SDP configurations *either* one LSP can be specified *or* LDP can be enabled. The SDP **ldp** and **lsp** commands are mutually exclusive. If an LSP is specified on an MPLS SDP, then LDP cannot be enabled on the SDP. To enable LDP on the SDP when an LSP is already specified, the LSP must be removed from the configuration using the **no lsp *lsp-name*** command.

Alternatively, if LDP is already enabled on an MPLS SDP, then an LSP cannot be specified on the SDP. To specify an LSP on the SDP, the LDP must be disabled. The LSP must have already been created in the **config>router>mpls** context. with a valid far-end IP address. RSVP must be enabled.

If no LSP is associated with an MPLS SDP, the SDP cannot enter the operationally up state. The SDP can be administratively enabled (**no shutdown)** with no LSP associations. The *lsp-name* may be shutdown, causing the association with the SDP to be operationally down (the LSP will not be used by the SDP).

Up to 16 LSP names can be entered on a single command line.

The **no** form of this command deletes one or more LSP associations from an SDP. If the *lsp-name* does not exist as an association or as a configured LSP, no error is returned. An *lsp-name* must be removed from all SDP associations before the *lsp-name* can be deleted from the system. The SDP must be administratively disabled (**shutdown)** before the last *lsp-name* association with the SDP is deleted.

Default none

Parameters *lsp-name* — The name of the LSP to associate with the SDP. An LSP name is case sensitive and is limited to 32 ASCII 7-bit printable characters with no spaces. If an exact match of *lsp-name* does not already exist as a defined LSP, an error message is generated. If the *lsp-name* does exist and the LSP **to** IP address matches the SDP **far-end** IP address, the association is created.

metric

Syntax **metric** *metric*
no metric

Context config>service>sdp

Description This command specifies the metric to be used within the tunnel table manager for decision making purposes. When multiple SDPs going to the same destination exist, this value is used as a tie-breaker by tunnel table manager users such as MP-BGP to select the route with the lower value.

Parameters *metric* — Specifies the SDP metric.

Values 0 — 65535

mixed-lsp-mode

Syntax [**no**] **mixed-lsp-mode**

Context config>service>sdp

Description This command enables the use by an SDP of the mixed-LSP mode of operation. This command indicates to the service manager that it must allow a primary LSP type and a backup LSP type in the same SDP configuration. For example, the **lsp** and **ldp** commands are allowed concurrently in the SDP configuration. The user can configure one or two types of LSPs under the same SDP. Without this command, these commands are mutually exclusive.

The user can configure an RSVP LSP as a primary LSP type with an LDP LSP as a backup type. The user can also configure a BGP RFC 3107 BGP LSP as a backup LSP type.

If the user configures an LDP LSP as a primary LSP type, then the backup LSP type must be an RFC 3107 BGP labeled route.

At any given time, the service manager programs only one type of LSP in the linecard that will

activate it to forward service packets according to the following priority order:

6. RSVP LSP type. Up to 16 RSVP LSPs can be entered by the user and programmed by the service manager in ingress linecard to load balance service packets. This is the highest priority LSP type.
7. LDP LSP type. One LDP FEC programmed by service manager but ingress IOM can use up to 16 LDP ECMP paths for the FEC to load balance service packets when ECMP is enabled on the node.
8. BGP LSP type. One RFC 3107-labeled BGP prefix programmed by the service manager. The ingress IOM can use more than one next-hop for the prefix.

In the case of the RSVP/LDP SDP, the service manager will program the NHLFE(s) for the active LSP type preferring the RSVP LSP type over the LDP LSP type. If no RSVP LSP is configured or all configured RSVP LSPs go down, the service manager will re-program the IOM with the LDP LSP if available. If not, the SDP goes operationally down.

When a higher priority type LSP becomes available, the service manager reverts back to this LSP at the expiry of the sdp-revert-time timer or the failure of the currently active LSP, whichever comes first. The service manager then re-programs the IOM accordingly. If the infinite value is configured, then the SDP reverts to the highest priority type LSP only if the currently active LSP failed.

Note however, that LDP uses a tunnel down damp timer which is set to three seconds by default. When the LDP LSP fails, the SDP will revert to the RSVP LSP type after the expiry of this timer. For an immediate switchover this timer must be set to zero. Use the **configure>router>ldp>tunnel-down-damp-time** command.

If the user changes the value of the sdp-revert-time timer, it will take effect only at the next use of the timer. Any timer which is outstanding at the time of the change will be restarted with the new value.

If class based forwarding is enabled for this SDP, the forwarding of the packets over the RSVP LSPs will be based on the FC of the packet as in current implementation. When the SDP activates the LDP LSP type, then packets are forwarded over the LDP ECMP paths using the regular hash routine.

In the case of the LDP/BGP SDP, the service manager will prefer the LDP LSP type over the BGP LSP type.

The service manager will re-program the IOM with the BGP LSP if available otherwise it brings down the SDP operationally.

Also Note the following difference in behavior of the LDP/BGP SDP compared to that of an RSVP/LDP SDP. For a given /32 prefix, only a single route will exist in the routing table: the IGP route or the BGP route. Thus, either the LDP FEC or the BGP label route is active at any given time. The impact of this is that the tunnel table needs to be re-programmed each time a route is deactivated and the other is activated. Furthermore, the SDP revert-time cannot be used since there is no situation where both LSP types are active for the same /32 prefix.

The **no** form of this command disables the mixed-LSP mode of operation. The user first has to remove one of the LSP types from the SDP configuration or the command will fail.

Default no mixed-lsp-mode

sdp-revert-time

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

Context config>service>sdp>mixed-lsp-mode

Description This command configures the delay period the SDP must wait before it reverts to a higher priority LSP type when one becomes available.

The **no** form of the command resets the timer to the default value of 0. This means the SDP reverts immediately to a higher priority LSP type when one becomes available.

Default 0

Parameters *seconds* — Specifies the delay period, in seconds, that the SDP must wait before it reverts to a higher priority LSP type when one becomes available. A value of zero means the SDP reverts immediately to a higher priority LSP type when one becomes available.

Values 0 — 600

infinite — This keyword forces the SDP to never revert to another higher priority LSP type unless the currently active LSP type is down.

sdp-group

Syntax [**no**] **sdp-group** *group-name*

Context config>service>sdp

Description This command configures the SDP membership in admin groups.

The user can enter a maximum of one (1) admin group name at once. The user can execute the command multiple times to add membership to more than one admin group. The admin group name must have been configured or the command is failed. Admin groups are supported on an SDP of type GRE and of type MPLS (BGP/RSVP/LDP). They are also supported on an SDP with the mixed-lsp-mode option enabled.

The **no** form of this command removes this SDP membership to the specified admin group.

Default none

Parameters *group-name* — Specifies the name of the SDP admin group. A maximum of 32 charactrs can be entered.

group-name

Syntax **group-name** *group-name* **value** *group-value*
no group-name *group-name*

Context config>service>sdp-group

Description This command defines SDP administrative groups, referred to as SDP admin groups.

SDP admin groups provide a way for services using a PW template to automatically include or exclude specific provisioned SDPs. SDPs sharing a specific characteristic or attribute can be made members of the same admin group. When users configure a PW template, they can include and/or exclude one or more admin groups. When the service is bound to the PW template, the SDP selection rules will enforce the admin group constraints specified in the **sdp-include** and **sdp-exclude** commands.

A maximum of 32 admin groups can be created. The group value ranges from zero (0) to 31. It is uniquely associated with the group name at creation time. If the user attempts to configure another group name for a group value that is already assigned to an existing group name, the SDP admin group creation is failed. The same happens if the user attempts to configure an SDP admin group with a new name but associates it to a group value already assigned to an existing group name.

The **no** option of this command deletes the SDP admin group but is only allowed if the group-name is not referenced in a pw-template or SDP.

Default none

Parameters *group-name* — Specifies the name of the SDP admin group. A maximum of 32 characters can be entered.

value *group-value* **—** Specifies the group value associated with this SDP admin group. This value is unique within the system.

Values 0—31

signaling

Syntax **signaling** {**off** | **tldp** | **bgp**}

Context config>service>sdp

Description This command specifies the signaling protocol used to obtain the ingress and egress pseudowire labels in frames transmitted and received on the SDP. When signaling is *off* then labels are manually configured when the SDP is bound to a service. The signalling value can only be changed while the administrative status of the SDP is down. Additionally, the signaling can only be changed on an SDP if that SDP is not in use by BGP-AD or BGP-VPLS. BGP signaling can only be enabled if that SDP does not already have pseudowires signaled over it. Also, BGP signaling is not supported with mixed mode LSP SDPs.

The **no** form of this command is not applicable. To modify the signaling configuration, the SDP must be administratively shut down and then the signaling parameter can be modified and re-enabled.

Default tldp

Parameters **off —** Ingress and egress signal auto-labeling is not enabled. If this parameter is selected, then each service using the specified SDP must manually configure VPN labels. This configuration is independent of the SDP's transport type, GRE, MPLS (RSVP or LDP).

tldp — Ingress and egress pseudowire signaling using T-LDP is enabled. Default value used when BGP AD automatically instantiates the SDP.

bgp — Ingress and egress pseudowire signaling using BGP is enabled. Default value used when BGP VPLS automatically instantiates the SDP.

tunnel-far-end

Syntax **tunnel-far-end** *ip-address*
no tunnel-far-end [*ip-address*]

Context config>service>sdp

Description This command enables the user to specify an SDP tunnel destination address that is different from the configuration in the SDP far-end option.

The SDP must be shutdown first to add or change the configuration of the **tunnel-far-end** option.

When this option is enabled, service packets are encapsulated using an LDP LSP with a FEC prefix matching the value entered in ip-address. By default, service packets are encapsulated using an LDP LSP with a FEC prefix matching the address entered in the SDP far-end option.

The T-LDP session to the remote PE is still targeted to the address configured under the **far-end option**. This means that targeted "hello" messages are sent to the far-end address, which is also the LSR-ID of the remote node. TCP based LDP messages, such as initialization and label mapping messages, are sent to the address specified in the transport-address field of the "hello" message received from the remote PE. This address can be the same as the remote PE LSR-ID, or a different address. This feature works, however, if the signaling option in the SDP is set to off instead of tldp, in which case, the service labels are statically configured.

This feature operates on an SDP of type LDP only. It can be used with VLL, VPLS, and VPRN services when an explicit binding to an SDP with the **tunnel-far-end** is specified. It also operates with a spoke interface on an IES or VPRN service. Finally, this feature operates with a BGP AD based VPLS service when the **use-provisioned-sdp** option is enabled in the pseudowire template.

This feature is not supported in an SDP of type MPLS when an RSVP LSP name is configured under the SDP. It also does not work with a mixed-lsp SDP.

The **no** form of this command disables the use of the **tunnel-far-end** option and returns to using the address specified in the far-end.

Default no tunnel-far-end

Parameters *ip-address* — The system address of the far-end router for the SDP in dotted decimal notation.

path-mtu

Syntax **path-mtu** *bytes*

no path-mtu

Context config>service>sdp

Description This command configures the Maximum Transmission Unit (MTU) in bytes that the Service Distribution Point (SDP) can transmit to the far-end device router without packet dropping or IP fragmentation overriding the SDP-type default path-mtu.

The default SDP-type **path-mtu** can be overridden on a per SDP basis. Dynamic maintenance protocols on the SDP like RSVP may override this setting.

If the physical **mtu** on an egress interface or PoS channel indicates the next hop on an SDP path cannot support the current **path-mtu**, the operational **path-mtu** on that SDP will be modified to a value that can be transmitted without fragmentation.

The **no** form of this command removes any **path-mtu** defined on the SDP and the SDP will use the system default for the SDP type.

Default The default **path-mtu** defined on the system for the type of SDP is used.

network-domain

Syntax **network-domain** *network-domain-name*
no network-domain

Context config>service>sdp

Description This command assigns a given SDP to a given network-domain. The network-domain is then taken into account during sap-ingress queue allocation for VPLS SAP.

The network-domain association can only be done in a base-routing context. Associating a network domain with an loop-back or system interface will be rejected. Associating a network-domain with an interface that has no physical port specified will be accepted, but will have no effect as long as a corresponding port, or LAG, is undefined.

A single SDP can only be associated with a single network-domain.

Default per default “default” network domain is assigned

pbb-etype

Syntax **pbb-etype** [**0x0600**..**0xffff**]
no pbb-etype

Context configure>service>sdp

Default 0x88E7

Description This command configures the Ethertype used for PBB.

Values **0x0600**..**0xffff:** 1536 — 65535 (accepted in decimal or hex)

vlan-vc-etype

Syntax **vlan-vc-etype** *0x0600..0xffff*
no vlan-vc-etype [*0x0600..0xffff*]

Context config>service>sdp

Description This command configures the VLAN VC EtherType.

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

Default no vlan-vc-etype

Parameters *0x0600..0xffff* — Specifies a valid VLAN etype identifier.

SDP Keepalive Commands

keep-alive

Syntax **keepalive**

Context config>service>sdp

Description Context for configuring SDP connectivity monitoring keepalive messages for the SDP ID.

SDP-ID keepalive messages use SDP Echo Request and Reply messages to monitor SDP connectivity. The operating state of the SDP is affected by the keepalive state on the SDP-ID. SDP Echo Request messages are only sent when the SDP-ID is completely configured and administratively up. If the SDP-ID is administratively down, keepalives for that SDP-ID are disabled. SDP Echo Requests (when sent for keepalive messages) are always sent with the *originator-sdp-id*. All SDP-ID keepalive SDP Echo Replies are sent using generic IP/GRE OAM encapsulation.

When a keepalive response is received that indicates an error condition, the SDP ID will immediately be brought operationally down. Once a response is received that indicates the error has cleared and the **hold-down-time** interval has expired, the SDP ID will be eligible to be put into the operationally up state. If no other condition prevents the operational change, the SDP ID will enter the operational state.

A set of event counters track the number of keepalive requests sent, the size of the message sent, non-error replies received and error replies received. A keepalive state value is kept indicating the last response event. A keepalive state timestamp value is kept indicating the time of the last event. With each keepalive event change, a log message is generated indicating the event type and the timestamp value.

The table below describes keepalive interpretation of SDP echo reply response conditions and the effect on the SDP ID operational status.

Result of Request	Stored Response State	Operational State
keepalive request timeout without reply	Request Timeout	Down
keepalive request not sent due to non-existent *orig-sdp-id*[a]	Orig-SDP Non-Existent	Down
keepalive request not sent due to administratively down *orig-sdp-id*	Orig-SDP Admin-Down	Down
keepalive reply received, invalid origination-id	Far End: Originator-ID Invalid	Down
keepalive reply received, invalid responder-id	Far End: Responder-ID Error	Down
keepalive reply received, No Error	Success	Up (If no other condition prevents)

[a] This condition should not occur.

hello-time

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

Context config>service>sdp>keep-alive

Description Configures the time period between SDP keepalive messages on the SDP-ID for the SDP connectivity monitoring messages.

The **no** form of this command reverts the **hello-time** *seconds* value to the default setting.

Default **hello-time 10** — 10 seconds between keepalive messages

seconds — The time period in seconds between SDP keepalive messages, expressed as a decimal integer.

Values 1 — 3600

hold-down-time

Syntax **hold-down-time** *seconds*
no hold-down-time

Context config>service>sdp>keep-alive

Description Configures the minimum time period the SDP will remain in the operationally down state in response to SDP keepalive monitoring.

This parameter can be used to prevent the SDP operational state from “flapping” by rapidly transitioning between the operationally up and operationally down states based on keepalive messages.

When an SDP keepalive response is received that indicates an error condition or the **max-drop-count** keepalive messages receive no reply, the *sdp-id* will immediately be brought operationally down. If a keepalive response is received that indicates the error has cleared, the *sdp-id* will be eligible to be put into the operationally up state only after the **hold-down-time** interval has expired.

The **no** form of this command reverts the **hold-down-time seconds** *value* to the default setting.

Default **hold-down-time 10** — The SDP is operationally down for 10 seconds after an SDP keepalive error.

Parameters *seconds* — The time in seconds, expressed as a decimal integer, the *sdp-id* will remain in the operationally down state before it is eligible to enter the operationally up state. A value of 0 indicates that no **hold-down-time** will be enforced for *sdp-id*.

Values 0 — 3600

max-drop-count

Syntax **max-drop-count** *count*
no max-drop-count

Context config>service>sdp>keep-alive

Description This command configures the number of consecutive SDP keepalive failed request attempts or remote replies that can be missed after which the SDP is operationally downed. If the **max-drop-count** consecutive keepalive request messages cannot be sent or no replies are received, the SDP-ID will be brought operationally down by the keepalive SDP monitoring.

The **no** form of this command reverts the **max-drop-count** *count* value to the default settings.

Default **max-drop-count 3**

Parameters *count* — The number of consecutive SDP keepalive requests that are failed to be sent or replies missed, expressed as a decimal integer.

Values 1 — 5

message-length

Syntax **message-length** *octets*
no message-length

Context config>service>sdp>keep-alive

Description This command configures the SDP monitoring keepalive request message length transmitted.

The **no** form of this command reverts the **message-length** *octets* value to the default setting.

Default 0 — The message length should be equal to the SDP's operating path MTU as configured in the **path-mtu** command. If the default size is overridden, the actual size used will be the smaller of the operational SDP-ID Path MTU and the size specified.

octets — The size of the keepalive request messages in octets, expressed as a decimal integer. The **size** keyword overrides the default keepalive message size.

Values 40 — 9198

timeout

Syntax **timeout** *timeout*
no timeout

Context config>service>sdp>keep-alive

Description This command configures the time interval that the SDP waits before tearing down the session.

Default 5

Parameters *timeout* — The timeout time, in seconds.

Values 1 — 10

ETH-CFM Configuration Commands

eth-cfm

Syntax **eth-cfm**

Context config

Description This command enables the context to configure 802.1ag CFM parameters.

mep

Syntax **mep** *mep-id* **domain** *md-index* **association** *ma-index* [**vlan** *vlan-id*]
no mep *mep-id* **domain** *md-index* **association** *ma-index* [**vlan** *vlan-id*]

Context config>port>ethernet>eth-cfm
config>lag>eth-cfm
config>router>if>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

vlan-id — Specific to tunnel facility MEPs which means this option is only applicable to the `lag>eth-cfm>` context. Used to specify the outer vlan id of the tunnel.

Values 1 — 4094

ais-enable

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

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

Description This command enables the reception of AIS messages.

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

client-meg-level

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

Context config>port>ethernet>eth-cfm>mep>ais-enable
config>lag>eth-cfm> mep>ais-enable

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

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

Parameters *level* — Specifies the client MEG level.

Values 1 — 7

Default 1

interval

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

Context config>port>ethernet>eth-cfm>mep>ais-enable
config>lag>eth-cfm> mep>ais-enable

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

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

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

Default 1

priority

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

Context config>port>ethernet>eth-cfm>mep>ais-enable
config>lag>eth-cfm> mep>ais-enable

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

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

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

Values 0 — 7

Default 7

ccm-enable

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

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

Description This command specifies the priority of the CCM and LTM messages transmitted by the MEP. Since CCM does not apply to the Router Facility MEP only the LTM priority is of value under that context.

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

Default *priority* — Specifies the priority value

Values 0 — 7

Default 7

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>port>ethernet>eth-cfm>mep
config>lag>eth-cfm>mep
config>router>if>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*]

The **no** form of the command disables eth-test capabilities.

test-pattern

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

Context config>port>ethernet>eth-cfm>mep>eth-test
config>lag>eth-cfm>mep>eth-test
config>router>if>eth-cfm>mep>eth-test

Description This command specifies the test pattern of the ETH-TEST frames. This does not have to be configured the same on the sender and the receiver.

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

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

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

crc-enable — Generates a CRC checksum.

Default all-zeros

low-priority-defect

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

Context config>port>ethernet>eth-cfm>mep>eth-test
config>lag>eth-cfm>mep>eth-test

Description This command specifies the lowest priority defect that is allowed to generate a fault alarm. This setting is also used to determine the fault state of the MEP which, well enabled to do so, causes a network reaction.

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>port>ethernet>eth-cfm>mep
config>lag>eth-cfm>mep
config>router>if>eth-cfm>mep

Description This command specifies the MAC address of the MEP.

The **no** form of the command reverts to the MAC address of the MEP back to the default, that of the port, since this is SAP based.

Default no mac-address

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.

facility-fault

Syntax [**no**] **facility-fault**

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

Description Allows the facility MEP to move from alarming only to network actionable function. This means a facility MEP will not merely report the defect conditions but will be able to action based on the transition of the MEP state. Without this command the facility MEP will only monitor and report and conditions of the MEP do not affect related services.

Default no facility-fault

tunnel-fault

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

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

```
config>service>epipe>sap>eth-cfm
config>service>ipipe>eth-cfm
config>service>ipipe>sap>eth-cfm
config>service>ies>eth-cfm
config>service>ies>if>sap>eth-cfm
config>service>ies>sub-if>grp-if>sap>eth-cfm
config>service>vprn>eth-cfm
config>service>vprn>if>sap>eth-cfm
config>service>vprn>sub-if>grp-if>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 the **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)

domain

Syntax **domain** *md-index* [**format** {**dns** | **mac** | **none** | **string**}] **name** *md-name* **level** *level*

domain *md-index*

no domain *md-index*

Context config>eth-cfm

Description This command configures Connectivity Fault Management domain parameters.

The **no** form of the command removes the MD index parameters from the configuration.

Parameters *md-index* — Specifies the Maintenance Domain (MD) index value.

Values 1 — 4294967295

format {**dns** | **mac** | **none** | **string**} **—** Specifies a value that represents the type (format).

Values	**dns**:	Specifies the DNS name format.
	mac:	X:X:X:X:X:X-u X: [0..FF]h u: [0..65535]d
	none:	Specifies a Y.1731 domain format and the only format allowed to

execute Y.1731 specific functions.

string	Specifies an ASCII string.

Default string

name *md-name* **—** Specifies a generic Maintenance Domain (MD) name.

Values 1 — 43 characters

level *level* **—** Specifies the integer identifying the maintenance domain level (MD Level). Higher numbers correspond to higher maintenance domains, those with the greatest physical reach, with the highest values for customers' CFM packets. Lower numbers correspond to lower maintenance domains, those with more limited physical reach, with the lowest values for single bridges or physical links.

Values 0 — 7

association

Syntax **association** *ma-index* [**format** {**icc-based** | **integer** | **string** | **vid** | **vpn-id**}] **name** *ma-name*
association *ma-index*
no association *ma-index*

Context config>eth-cfg>domain

Description This command configures the Maintenance Association (MA) for the domain.

ma-index — Specifies the MA index value.

Values 1 — 4294967295

format {**icc-based** | **integer** | **string** | **vid** | **vpn-id**} **—** Specifies a value that represents the type (format).

Values

icc-based:	Only applicable to a Y.1731 context where the domain format is configured as none. Allows for exactly a 13 character name.
integer:	0 — 65535 (integer value 0 means the MA is not attached to a VID.)
string:	raw ascii
vid:	0 — 4095
vpn-id:	RFC-2685, *Virtual Private Networks Identifier* xxx:xxxx, where x is a value between 00 and FF. for example 00164D:AABBCCDD

Default integer

name ***ma-name*** **—** Specifies the part of the maintenance association identifier which is unique within the maintenance domain name.

Values 1 — 45 characters

bridge-identifier

Syntax [**no**] **bridge-identifier** *bridge-id*

Context config>eth-cfm>domain>association

Description This command configures the service ID for the domain association. The value must be configured to match the *service-id* of the service where MEPs for this association will be created. Note that there is no verification that the service with a matching *service-id* exists. This is not used for facility MEPs as they are not tied to services.

Parameters *bridge-id* — Specifies the bridge ID for the domain association.

Values 1 — 2147483647

mhf-creation

Syntax **mhf-creation** {**default** | **none** | **explicit** | **static**}
no mhf-creation

Context config>eth-cfm>domain>association>bridge-identifier

Description This command determines whether to allow MIP creation for the MA. Use of the none, default and explicit parameters are only allowed for MHFs (MIPs) that are not associated with a configured Primary VLAN. The static parameter is only applicable to MHFs (MIPs) that are associated with a Primary VLAN.

Default none

Parameters **default —** Specifies MHFs (MIPs) can be created for this SAP or Spoke-Sdp without the requirement for a MEP at some lower MA level.

none — Specifies that no MHFs (MIPs) can be created for this SAP or Spoke-SDP.

explicit — Specifies that MHFs (MIPs) can be created for this SAP or Spoke-Sdp only if a MEP is created at some lower MD Level. There must be at least one lower MD Level MEP provisioned on the same SAP or Spoke-SDP.

static — Specifies the exact level of the MHF (MIP) that will be created for this SAP. Multiple MHFs (MIPs) are allowed as long as the MD Level hierarchy is properly configured for the particular Primary VLAN. Ingress MHFs (MIPs) with primary VLAN are not supported on SDP Bindings.

mip-ltr-priority

Syntax **mip-ltr-priority** *priority*
no mip-ltr-priority

Context config>eth-cfm>domain>association>bridge-identifier

Description This command allows the operator to set the priority of the Linktrace Response Message (ETH-LTR) from a MIP for this association. If this command is not specified a LTR priority of 7 will be used.

Default no mip-ltr-priority

Parameters *priority* — Specifies the priority of the Linktrace Response Message (ETH-LTR) from a MIP for this association.

Values 0 — 7

vlan

Syntax **vlan** *vlan-id*
no vlan

Context config>eth-cfm>domain>association>bridge-identifier

Description This command configures the bridge-identifier primary VLAN ID. Note that it is informational only, and no verification is done to ensure MEPs on this association are on the configured VLAN.

Parameters *vlan-id* — Specifies a VLAN ID monitored by MA.

Values 0 — 4094

ccm-interval

Syntax **ccm-interval** *interval*
no ccm-interval

Context config>eth-cfm>domain>association

Description This command configures the CCM transmission interval for all MEPs in the association.

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

Default 10 seconds

Parameters *interval* — Specifies the interval between CCM transmissions to be used by all MEPs in the MA.

Values 10 milliseconds, 100 milliseconds, 1 second, 10 seconds, 60 seconds, 600 seconds, 100 milliseconds

remote-mep

Syntax [**no**] **remote-mepid** *mep-id* **remote-mac** {*unicast-da* | default}

Context config>eth-cfm>domain>association

Description This command identifies remote maintenance association endpoint (MEP) the systems is expecting to receive packets form. Optionally, the operator may configure a unciast MAC address associated with the remote-mep. This unicast value will replace the default layer two class 1 multicast address that is typically associated with ETH-CC packets.

Note: This command is not supported with sub second CCM intervals. **unicast-da** may only be configured when a single remote MEP exists in the association.

Default multicast class 1 address

Parameters **remote-mep** *mep-id* — Specifies the remote MEP identifier.

Values *mep-id* 1 — 8191

remote-mac {*unicast-da* | default}

Values unicast-da —The unicast layer two destination address in the form xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx.

default — Removes the unicast address and reverts back to class 1 multicast.

remote-mepid

Syntax [**no**] **remote-mepid** *mep-id*

Context config>eth-cfm>domain>association

Description This command configures the remote maintenance association end point (MEP) identifier.

Parameters *mep-id* — Maintenance association end point identifier of a remote MEP whose information from the MEP database is to be returned.

Values 1 — 8191

ccm-hold-time

Syntax **ccm-hold-time down** *delay-down*
no ccm-hold-time

Context config>eth-cfm>domain>association

Description This command allows a sub second CCM enabled MEP to delay a transition to a failed state if a configured remote CCM peer has timed out. The MEP will remain in the UP state for 3.5 times CCM interval + down-delay.

The no form of this command removes the additional delay

Default 0 second

Parameters **down —** Specifies the amount of time to delay in 100ths of a second

Values 0-1000

slm

Syntax **slm**

Context config>eth-cfm

Description This is the container that provides the global configuration parameters for ITU-T Synthetic Loss Measurement (ETH-SL).

inactivity-timer

Syntax **inactivity-timer** *timeout*
[**no**] **inactivity-timer**

Context config>eth-cfm>slm

Description The time the responder keeps a test active. Should the time between packets exceed this values within a test the responder will mark the previous test as complete. It will treat any new packets from a peer with the same test-id, source-mac and MEP-ID as a new test responding with the sequence number one.

Default 100 seconds

Parameters **timeout —** Specifies the amount of time in seconds

Values 10 100

ccm-hold-time

Syntax **ccm-hold-time down** *delay-down*
no ccm-hold-time

Context config>eth-cfm>domain>association

Description This command allows a sub second CCM enabled MEP to delay a transition to a failed state if a configured remote CCM peer has timed out. The MEP will remain in the UP state for 3.5 times CCM interval + down-delay.

The no form of this command removes the additional delay

Default 0 second

Parameters **down —** Specifies the amount of time to delay in 100ths of a second

Values 0-1000

system

Syntax **system**

Context config>eth-cfm

Description This command configures Connectivity Fault Management General System parameters.

grace-tx-enable

Syntax **grace-tx-enable**
[no] grace-tx-enable

Context config>eth-cfm>system

Description This command enables and disables the transmission of ETH-VSM messages to delay CCM timeout and AIS churn during ISSU and soft reset functions.

Default grace-tx-enable

redundancy

Syntax **redundancy**

Context config>eth-cfm

Description This command provides the context under which the ETH-CFM redundancy parameters are to be configured

Default none

mc-lag

Syntax **mc-lag**

Context config>eth-cfm>redundancy

Description This command provides the context under which the MC-LAG specific ETH-CFM redundancy parameters are to be configured

Default none

propagate-hold-time

Syntax **propagate-hold-time *second*>**
no propagate-hold-time

Context config>eth-cfm>redundancy>mc-lag

Description Configure the delay, in seconds, that fault propagation is delayed because of port or MC-LAG state changes. This provides the amount of time for system stabilization during a port state changes that may be protected by MC-LAG. This command requires the standby-mep-shutdown command in order to take effect.

Default 1 second

Parameters **seconds** — The amount of time in seconds, zero means no delay.

Values 0-60

standby-mep-shutdown

Syntax **standby-mep-shutdown**
no standby-mep-shutdown

Context config>eth-cfm>redundancy>mc-lag

Description System wide command that enables MEPs to track the state of MC-LAG. This allows MEPs on the standby MC-LAG to act administratively down.

Default no standby-mep-shutdown

ETH-Tunnel Commands

eth-tunnel

Syntax **eth-tunnel** *tunnel-index*
no eth-tunnel tunnel-index

Context config

Description This command configures a unique Ethernet Tunnel Identifier for an Ethernet Tunnel Group.

The no form of the command removes the index ID from the configuration.

Default none

Parameters *tunnel-index* — Specifies a tunnel index identifier.

Values 1 — 1024

ccm-hold-time

Syntax **ccm-hold-time** { **down** *down-timeout* **|** **up** *up-timeout* }
no ccm-hold-time

Context config>eth-tunnel

Description This command allows a sub second CCM enabled MEP to delay a transition to a failed state if a configured remote CCM peer has timed out. The MEP will remain in the UP state for 3.5 times CCM interval + down-delay.

The **no** form of this command removes the additional delay

Default **down** *down-timeout* **—** Specifies the time, in centiseconds, used for the hold-timer for associated Continuity Check (CC) Session down event dampening. This guards against reporting excessive member operational state transitions.

This is implemented by not advertising subsequent transitions of the CC state to the Ethernet Tunnel Group until the configured timer has expired.

Values 0 — 1000

Default 0

up *up-timeout* **—** Specifies the time, in deciseconds, used for the hold-timer for associated Continuity Check (CC) Session up event dampening. This guards against reporting excessive member operational state transitions.

This is implemented by not advertising subsequent transitions of the CC state to the Ethernet Tunnel Group until the configured timer has expired.

Values 0 — 5000

Default 20

ethernet

Syntax **ethernet**

Context config>eth-tunnel

Description This command enables the context to configure Ethernet parameters for the Ethernet tunnel.

encap-type

Syntax **encap-type** {**dot1q**|**qinq**}
no encap-type

Context config>eth-tunnel>ethernet

Description This command configures the encapsulation method used to distinguish customer traffic on a LAG. The encapsulation type is configurable on a LAG port. The LAG port and the port member encapsulation types must match when adding a port member.

If the encapsulation type of the LAG port is changed, the encapsulation type on all the port members will also change. The encapsulation type can be changed on the LAG port only if there is no interface associated with it. If the MTU is set to a non default value, it will be reset to the default value when the encap type is changed.

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

Default dot1q

Parameters **dot1q —** Specifies that frames carry 802.1Q tags where each tag signifies a different service.

qinq — Specifies the qinq encapsulation method.

mac

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

Context config>eth-tunnel>ethernet

Description This command assigns a specific MAC address to an Ethernet port, Link Aggregation Group (LAG), Ethernet tunnel, or BCP-enabled port or sub-port.

Only one MAC address can be assigned to a port. When multiple **mac** commands are entered, the last command overwrites the previous command. When the command is issued while the port is operational, IP will issue an ARP, if appropriate, and BPDU's are sent with the new MAC address.

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

Default A default MAC address is assigned by the system from the chassis MAC address pool.

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 addresses6-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.

lag-emulation

Syntax **lag-emulation**

Context config>eth-tunnel

Description This command enables the context to configure eth-tunnel loadsharing parameters/

access

Syntax **access**

Context config>eth-tunnel>lag-emulation

Description This command enables the context to configure eth-tunnel loadsharing access parameters

adapt-qos

Syntax **adapt-qos** {**distribute**|**link**}
no adapt-qos

Context config>eth-tunnel>lag-emulation>access

Description This command specifies how the LAG queue and virtual scheduler buffering and rate parameters are adapted over multiple active MDAs. This command applies only to access LAGs.

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

Parameters **distribute** — Creates an additional internal virtual scheduler per IOM as parent of the configured SAP queues and vitual schedulers per LAG member port on that IOM. This internal virtual scheduler limits the total amount of egress bandwidth for all member ports on the IOM to the bandwidth specified in the egress qos policy.

link — Specifies that the LAG will create the SAP queues and virtual schedulers with the actual parameters on each LAG member port.

per-fp-ing-queuing

Syntax [**no**] **per-fp-ing-queuing**

Context config>eth-tunnel>lag-emulation>access

Description This command specifies whether a more efficient method of queue allocation for the LAG should be utilized.

The **no** form of the command disables the method of queue allocation.

path-threshold

Syntax **path-threshold** *num-paths*
no path-threshold

Context config>eth-tunnel>lag-emulation

Description This command configures whether a more efficient method of queue allocation for Ethernet Tunnel Group SAPs should be utilized.

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

Default **no per-fp-ing-queuing**

Parameters *num-paths* — Specifies the behavior for the eth-tunnel if the number of operational members is equal to or below a threshold level.

Values 0 — 15

path

Syntax **path**

Context config>eth-tunnel

Description This command configures one of the two paths supported under the Ethernet tunnel.

The **no** form of this command removes the path from under the Ethernet tunnel. If this is the last path, the associated SAP need to be un-configured before the path can be deleted.

Default no path

Parameters *path-index* — Specifies the identifier for the path.

Values 1 — 16

control-tag

Syntax **control-tag** *qtag*[.*qtag*]
no control-tag

Context config>eth-tunnel>path

Description This command specifies the VLAN-ID to be used for Ethernet CFM and G.8031 control plane exchanges. If the operator wants to replace an existing control-tag, the parent path needs to be in shutdown state, then deleted and recreated before a new control-tag can be specified.

The **no** form of this command is used just to indicate that a control-tag is not configured. The procedure described above, based on ‘no path’ command must be used to un-configure/change the control-tag assigned to the path.

Default no control tag specified

Parameters *vlan-id* — specifies the value of the VLAN ID to be used for the control tag.

Values 0 — 4094

eth-cfm

Syntax **eth-cfm**

Context config>eth-tunnel>path

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

mep

Syntax [**no**] **mep** *mep-id* **domain** *md-index* **association** *ma-index*

Context config>eth-tunnel>path>eth-cfm

Description This command provisions an 802.1ag 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

ccm-enable

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

Context config>eth-tunnel>path>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>eth-tunnel>path>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

Context config>eth-tunnel>path>eth-cfm>mep

Description This command inserts additional padding in the CCM packets.

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

Parameters *ccm-padding* — Specifies the additional padding in the CCM packets.

Values 3 — 1500 octets

control-mep

Syntax [**no**] **control-mep**

Context config>eth-tunnel>path>eth-cfm>mep

Description This command enables the Ethernet ring control on the MEP. The use of control-mep command is mandatory for a ring. MEP detection of failure using CCM may be enabled or disabled independently of the control mep.

The **no** form of this command disables Ethernet ring control.

Default no control-mep

eth-test-enable

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

Context config>eth-tunnel>path>eth-cfm>mep

Description This command enables eth-test functionality on MEP. 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 done for both the provisioning and test to ensure the MEP is an Y.1731 MEP (MEP provisioned with domain format none, association format icc-based). If not, the operation fails. An error message in the

CLI and SNMP will indicate the problem.

bit-error-threshold

Syntax **bit-error-threshold** *bit-errors*

Context config>eth-ring>path>eth-cfm>mep

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

Default 1

Parameters *bit-errors* — Specifies the lowest priority defect.

Values 0 — 11840

test-pattern

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

Context config>eth-ring>path>eth-cfm>mep>eth-test-enable

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

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

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

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

crc-enable — Generates a CRC checksum.

Default all-zeros

low-priority-defect

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

Context config>eth-tunnel>path>eth-cfm>mep

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

Default remErrXcon

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>eth-tunnel>path>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>eth-tunnel>path>eth-cfm>mep

Description This command enables one way delay threshold time limit.

Default 3 seconds

Parameters *priority* — Specifies the value for the threshold.

Values 0 — 600

member

Syntax **member** *port-id*
no member

Context config>eth-tunnel>path

Description This command configures the path member.

The **no** form of the command removes the port-id from the configuration.

Default none

Parameters *port-id* — Specifies the path member

Values slot/mda/port

precedence

Syntax **precedence** {**primary|secondary**}

Context config>eth-tunnel>path

Description This command specifies the precedence to be used for the path. Only two precedence options are supported: **primary** and **secondary**.

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

Default secondary

Parameters **primary | secondary —** specifies the path precedence as either primary or secondary.

protection-type

Syntax **protection-type** {**g8031-1to1|loadsharing**}

Context config>eth-tunnel

Description This command configures the model used for determining which members are actively receiving and transmitting data.

When the value is set to 'g8031-1to1 (1)', as per G.8031 spec, only two members are allowed, and only one of them can be active at one point in time.

When the value is set to 'loadsharing (2)', multiple members can be active at one point in time.

Default g8031-1to1

revert-time

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

Context config>eth-tunnel

Description This command configures the revert time for an Eth tunnel. It ranges from 60 seconds to 720 second by 1 second intervals.

The **no** form of this command this command means non-revertive mode and revert time essentially is 0 meaning the revert timers are not set.

Default 300 seconds

Parameters *value* — Specifies the guard-time.

Values 60 — 720 seconds

Tools Perform Commands

tools

Syntax **tools**

Context root

Description This command enables the context to enable useful tools for debugging purposes.

Default none

Parameters **dump —** Enables dump tools for the various protocols.

perform — Enables tools to perform specific tasks.

perform

Syntax perform

Context tools

Description This command enables the context to enable tools to perform specific tasks.

Default none

service

Syntax **services**

Context tools>perform

Description This command enables the context to configure tools for services.

id

Syntax **id** *service-id*

Context tools>perform>service

Description This command enables the context to configure tools for a specific service.

Parameters *service-id* — Specify an existing service ID.

Values 1 — 2147483647

loopback

Syntax **loopback**

Context tools>perform>service>id

Description Tools for placing and removing saps and SDP bindings in data loopback. Overwrite will occur for any SAP or SDP Binding when issuing a subsequent loopback command on the same SAP or SDP Binding.

Interactions: Loopback functions are only applicable to epipe, PBB ePipe, VPLS, I-VPLS and PBB core service contexts.

sap

Syntax **sap** *sap-id* **start** *mode* [**mac-swap** [**mac** *ieee-address* [**all**]]]
sap *sap-id* **stop**

Context tools>perform>service>loopback

Description This command places and removes the specific SAP in loopback mode for reflecting traffic back in the direction of the received stream. This is only applicable to Ethernet based SAPs.

Parameters *sap-id* —

null *port-id* | *lag-id*
dot1q *port-id* | l*ag-id* :qtag1
qinq *port-id* | l*ag-id* :qtag1.qtag2
port-id *slot*/*mda*/*port*
lag-id lag-*id*

lag	keyword
id	[1..200]
qtag1	[0..4094]
qtag2	[*|0..4094]

start — keyword that places the sap in loopback mode.

mode **ingress** | **egress** : keywords that specifies the location on the loopback in relation to the SAP.

ingress — Traffic arriving at the sap-ingress will be reflected back out the same sap.

egress — Traffic arriving at the sap-egress will be reflected back into the service in the direction of the original source.

stop — removes the SAP from loopback mode.

mac-swap — enable source address and destination address swapping for the reflected packets when the arriving packet is unicast. Any broadcast and multicast packets arriving on a looped point will be dropped.

mac — *ieee-address* optionally configure the source MAC address used in the reflected packet when the arriving packet is a broadcast or multicast. This does not apply to arriving unicast packets.

Value: 6-byte unicast mac-address in the form
xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx

all — configured ieee-address is used as the source address for all reflected packets regardless of the arriving destination.

Default [**no**] **mac-swap** – no swapping of mac addresses are performed without specifying this option and any non-unicast destined packets will not be reflected back to the source.

sdp

Syntax **sdp** *sdp-id:vc-id* **start** *mode* [**mac-swap** [**mac** *ieee-address* [**all**]]]
sdp *sdp-id:vc-id* **stop**

Context tools>perform>service>loopback

Description This command places the specific MPLS SDP binding in loopback mode for reflecting traffic back in the direction of the received stream. This is only applicable to MPLS SDP Bindings.

Parameters *sdp-id:vc-id —*

sdp-id	[1..17407]
vc-id	[1.. 4294967295]

start — keyword that places the sap in loopback mode.

mode **ingress** | **egress** : keywords that specifies the location on the loopback in relation to the MPLS SDP Binding.

ingress — Traffic arriving at the sap-ingress will be reflected back out the same sap.

egress — Traffic arriving at the sap-egress will be reflected back into the service in the direction of the original source.

stop — rkeyword that removes the MPLS SD- binding from loopback mode.

mac-swap — enable source address and destination address swapping for the reflected packets when the arriving packet is unicast. Any broadcast and multicast packets arriving on a looped point will be dropped.

mac — *ieee-address* optionally configure the source MAC address used in the reflected packet when the arriving packet is a broadcast or multicast. This does not apply to arriving unicast packets.

Value: 6-byte unicast mac-address in the form
xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx

all — configured ieee-address is used as the source address for all reflected packets regardless of the arriving destination.

Default [**no**] **mac-swap** – no swapping of mac addresses are performed without specifying this option and any non-unicast destined packets will not be reflected back to the source.