Triple Play Subscriber Management Configuration Commands

Generic Commands

description

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

Context config>subscr-mgmt>authentication-policy
config>subscr-mgmt>host-tracking
config>subscr-mgmt>sla-profile
config>subscr-mgmt>sla-profile>egress>ip-filter>entry
config>subscr-mgmt>sla-profile>ingress>ip-filter>entry
config>subscr-mgmt>sub-ident-policy
config>subscr-mgmt>sub-profile
config>service>vpls>gsmp>group
config>log>accounting-policy
config>service>vprn>redundant-interface
config>service>ies>redundant-interface
config>service>ies>subscriber-interface
config>service>ies>subscriber-interface>group-interface
config>service>ies>subscriber-interface>grp-if>dhcp
config>service>ies>sub-if>grp-if>srrp
config>service>vprn>subscriber-interface
config>service>vprn>sub-if>dhcp
config>service>vprn>subscriber-interface>group-interface
config>service>vprn>subscriber-interface>grp-if>sap
config>service>vprn>sub-if>grp-if>srrp
config>service>vprn>subscriber-interface>grp-if>dhcp
config>service>vprn>gsmp>group
config>service>vprn>gsmp>group>neighbor
config>redundancy>multi-chassis>peer
config>subscr-mgmt>cat-map
config>sub-mgmt>diameter-policy
config>sub-mgmt>credit-control-policy
config>sub-mgmt>host-lockout>policy
config>subscr-mgmt>sub-mcac-policy
config>aaa>route-downloader
config>port>ethernet>access>egress

Description This command creates a text description stored in the configuration file for a configuration context.

The **description** command associates a text string with a configuration context to help identify the context in the configuration file.

The **no** form of this command removes any description string from the context.

Default No description is associated with the configuration context.

Parameters *description-string* — A text string describing the entity. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters excluding double quotes. If the

string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

shutdown

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

Context config>subscr-mgmt>sub-ident-policy>primary
config>subscr-mgmt>sub-ident-policy>secondary
config>subscr-mgmt>sub-ident-policy>tertiary
config>service>vpls>sap>sub-sla-mgmt
config>service>vpls>gsmp
config>service>vpls>gsmp>group
config>service>vpls>gsmp>group>neighbor
config>service>vprn>redundant-interface
config>service>vprn>redundant-interface>spoke-sdp
config>service>vprn>subscriber-interface
config>service>vprn>subscriber-interface>group-interface
config>service>vprn>subscriber-interface>grp-if>dhcp
config>service>vprn>sub-if>grp-if>srrp
config>service>ies>subscriber-interface
config>service>ies>subscriber-interface>grp-if>dhcp
config>service>ies>sub-if>grp-if>srrp
config>service>ies>redundant-interface
config>service>ies>sub-if>grp-if>arp-host
config>service>vprn>gsmp>group>neighbor
config>redundancy>multi-chassis>peer
config>redundancy>multi-chassis>peer>mc-lag
config>redundancy>multi-chassis>peer>sync
config>service>ies>sub-if>dhcp
config>subscr-mgmt>diameter-policy>diameter-base>peer
config>subscr-mgmt>sub-mcac-policy
config>aaa>route-downloader

Description The **shutdown** command administratively disables the entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics. Many entities must be explicitly enabled using the **no shutdown** command.

The **shutdown** command administratively disables an entity. The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they can be deleted.

Unlike other commands and parameters where the default state is not indicated in the configuration file, **shutdown** and **no shutdown** are always indicated in system generated configuration files.

Shutting down a subscriber interface will operationally shut down all child group interfaces and SAPs. Shutting down a group interface will operationally shut down all SAPs that are part of that group-interface.

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

Default no shutdown

subscriber-mgmt

Syntax **subscriber-mgmt**

Context config

Description This command enables the context to configure subscriber management entities. A subscriber is uniquely identified by a subscriber identification string. Each subscriber can have several DHCP sessions active at any time. Each session is referred to as a subscriber host and is identified by its IP address and MAC address.

All subscriber hosts belonging to the same subscriber are subject to the same hierarchical QoS (HQoS) processing. The HQoS processing is defined in the **sub-profile** (the subscriber profile). A sub-profile refers to an existing scheduler policy (configured in **the configure>qos>scheduler-policy** context) and offers the possibility to overrule the rate of individual schedulers within this policy.

Because all subscriber hosts use the same scheduler policy instance, they must all reside on the same complex.

ANCP and GSMP Commands

ancp

Syntax **ancp**

Context config>subscr-mgmt
config>subscr-mgmt>sub-prof

Description This command enables the context to configure Access Node Control Protocol (ANCP) parameters.

ancp-policy

Syntax **ancp-policy** *name*

Context config>subscr-mgmt>ancp

Description This command creates an Access Node Control Protocol (ANCP) policy. The policy is associated with either the ANCP string (static case) or subscriber-profile (dynamic case) and defines the behavior of the hosts belonging to these profiles.

ANCP polices control rates and subscribers based on port-up/port-down messages from the access node. When configured, the 7750 SR should stop SHCV to a host that is part of a port defined to be down (by port-down message). When the node receives a port-up message for a port that was in port-down, state the node will initiate the SHCV process immediately to verify connectivity.

When ANCP is used with Enhanced Subscriber Management, the ANCP string last associated with the subscriber will be used. All hosts of a subscriber will be updated with the new ANCP string.

Default No policies are defined.

Parameters *name* — Configures the ANCP policy name.

ancp-policy

Syntax **ancp-policy** *name*

Context config>subscr-mgmt>sub-prof>ancp

Description This command specifies an existing Access Node Control Protocol (ANCP) policy to associate with the subscriber profile. The policy is associated with either the ANCP string (static case) or subscriber-profile (dynamic case) and defines the behavior of the hosts belonging to these profiles.

Default No policies are defined.

Parameters *name* — Specifies an existing ANCP policy name.

ingress

Syntax **ingress**

Context config>subscr-mgmt>ancp>ancp-policy

Description This command configures ingress ANCP policy parameters.

rate-adjustment

Syntax **rate-adjustment** *adjusted-percent*
no rate-adjustment

Context config>subscr-mgmt>ancp>ancp-policy>ingress
config>subscr-mgmt>ancp>ancp-policy>egress

Description This command configures a rate adjustment for the scheduler. The **rate-adjustment** command should be used when the rate returned by the DSLAM is calculated with different encapsulation than the 7750 SR. The node will adjust the rate by the percent specified as:

DSLAM_RATE*adjust-rate/100 — rate-reduction.

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

Default none

Parameters *adjusted-percent* — Specifies a rate adjustment for the scheduler.

Values 1 — 200

Default 100

rate-reduction

Syntax **rate-reduction** *kilobit-per-second*
no rate-reduction

Context config>subscr-mgmt>ancp>ancp-policy>ingress
config>subscr-mgmt>ancp>ancp-policy>egress

Description This command defines a constant rate reduction to the rate specified by the DSLAM. The **rate-reduction** command should be used if the node should adjust the rate to a value that is offset (for example by a fixed multicast dedicated bandwidth) compared to the total available on the DSLAM.

When set, the rate will be:

DSLAM_RATE*adjust-rate/100 — rate-reduction

Default none

rate-monitor

Syntax **rate-monitor** *kilobit-per-second* [**alarm**]
no rate-monitor

Context config>subscr-mgmt>ancp>ancp-policy>ingress
config>subscr-mgmt>ancp>ancp-policy>egress

Description This command configures the rate monitor level.

Default none

Parameters *kilobit-per-second* — Specifies the rate, in kilobits, below which the system will generate an event.

alarm — When the monitored rate is below the configured value the system generates an alarm (trap) to the management system. The trap includes the rate as well as the ANCP policy name and the ANCP string.

rate-modify

Syntax **rate-modify** {**scheduler** *scheduler-name* | **arbiter** *arbiter-name*}
no rate-modify

Context config>subscr-mgmt>ancp>ancp-policy>ingress

Description This command configures ingress rate modify scheduler parameters.

Default none

Parameters **scheduler** *scheduler-name* **—** Specifies a scheduler name.

arbiter *arbiter-name* **—** Specifies an arbiter name

egress

Syntax **egress**

Context config>subscr-mgmt>ancp>ancp-policy

Description This command configures egress ANCP policy parameters.

rate-modify

Syntax **rate-modify** {**scheduler** *scheduler-name* | **arbiter** *arbiter-name*}
rate-modify agg-rate-limit
no rate-modify

Context config>subscr-mgmt>ancp>ancp-policy>egress

Description This command configures egress rate modify scheduler parameters.

Default none

Parameters **agg-rate-limit —** specifies that the maximum total rate for all subscriber egress queues for each subscriber associated with the policy.

scheduler *scheduler-name* — Specify a scheduler name.

arbiter *arbiter-name* **—** Specifies an arbiter name

port-down

Syntax [**no**] **port-down**

Context config>subscr-mgmt>ancp>ancp-policy

Description This command specifies the number of GSMP portdown messages received in this ANCP session.

disable-shcv

Syntax [**no**] **disable-shcv** [**alarm**] [**hold-time** *seconds*]

Context config>subscr-mgmt>ancp>ancp-policy>port-down

Description When this command is configured, the node will suspend SHCV for the hosts defined with this ANCP policy until the access node sends a port-up message. When the hold-time parameter is used, the node will suspend SHCV for the period of time defined. If the hold-time parameter is not defined the node will suspend SHCV until a port-up message is received.

If the optional alarm flag is used the node should send a SHCV alarm before suspending.

Default no disable-shcv

ancp-static-map

Syntax **ancp-static-map**

Context config>subscr-mgmt>ancp

Description This command enables the context to configure a static ANCP name map.

Default ancp-static-map

entry

Syntax **entry key** *ancp-string* **customer** *customer-id* **multi-service-site** *customer-site-name* **ancp-policy** *policy-name*
entry key *ancp-string* **sap** *sap-id* **ancp-policy** *policy-name*
no entry key *ancp-string* **customer** *customer-id* **multi-service-site** *customer-site-name*
no entry key *ancp-string* **sap** *sap-id*

Context config>subscr-mgmt>ancp>static-map

Description This command configures an ANCP name. When ANCP is configured to provide rate adaptation without the use of enhanced subscriber management, this command will define how to map an ANCP key (usually the circuit-id of the DSLAM port) to either a SAP and a scheduler name (when a Multi-Service Site (MSS) is not used) or a customer, site and scheduler name when MSS is used.

Different ANCP names may be used with the same SAPs or customer ID/MSS combinations to allow schedulers within the policy to be mapped to the ANCP names. An ANCP string and SAP combination may reference only one ancp-policy. An ANCP string and customer and site-name combination may reference a single ancp-policy.

Default none

Parameters **key** *ancp-string* — Specify the ASCII representation of the DSLAM circuit-id name.

customer *customer-id* — Specify the associated existing customer name.

multi-service-site *customer-site-name* — Specify the associated customer's configured MSS name.

ancp-policy *policy-name* — Specify an existing ANCP policy name.

sap *sap-id* — Specifies the physical port identifier portion of the SAP definition. See Common Service Commands on page 1740 for *sap-id* command syntax.

gsmp

Syntax **gsmp**

Context config>service>vpls
config>service>vprn

Description This command enables the context to configure GSMP connections maintained in this service.

Default not enabled

group

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

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

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

ancp

Syntax **ancp**

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

Description This command configures ANCP parameters for this GSMP group.

dynamic-topology-discover

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

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

Description This command enables the ANCP dynamic topology discovery capability.

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

oam

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

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

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

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

hold-multiplier

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

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

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

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

Values 1 — 100

idle-filter

Syntax [**no**] **idle-filter**

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

Description This command when applied will filter out new incoming ANCP messages while subscriber "DSL-line-state" is IDLE. The command takes effect at the time that it is applied. Existing subscriber already in IDLE state are not purged from the database.

Default no idle-filter

keepalive

Syntax **keepalive** *seconds*
no keepalive

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

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

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

Values 1 — 25

neighbor

Syntax **neighbor** *ip-address* [**create**]
no neighbor *ip-address*

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

Description This command configures a GSMP ANCP neighbor.

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

local-address

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

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

Description This command configures the source ip-address used in the connection towards the neighbor. The local address is optional. If specified the node will accept connections only for that address in the service running ANCP. The address may be created after the reference but connections will not be

accepted until it is created. If the local address is not used, the system accepts connections on any interface within the routing context.

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

priority-marking

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

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

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

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

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

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

Values 0 — 7

persistency-database

Syntax [**no**] **persistency-database**

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

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

Default no persistency-database

RADIUS Policy Commands

radius-coa-port

Syntax **radius-coa-port** {**1647**|**1700**|**1812**|**3799**}
no radius-coa-port

Context config>aaa

Description This command configures the system-wide UDP port number that RADIUS is listening on for CoA and Disconnect messages

The **no** form of the command resets the default UDP port to 3799.

Default 3799

Parameters {**1647**|**1700**|**1812**|**3799**} — Specifies the udp port number for RADIUS CoA and Disconnect Messages.

authentication-policy

Syntax **authentication-policy** *name* [**create**]
no authentication-policy

Context config>subscr-mgmt

Description This command creates the context to configure RADIUS server parameters for session authentication. The policies can be applied to an IES or VPRN interface, or a VPLS SAP.

The **no** form of the command removes the RADIUS server configuration for session authentication.

RADIUS servers can be configured for three different applications:

1. For authentication of dynamic Triple Play subscriber sessions, under `config>subscr-mgmt>authentication-plcy`
2. For 802.1x port authentication, under `config>system>security>dot1x>radius-plcy`
3. For CLI login users, under `config>system>radius`

Default none

Parameters *name* — The name of the profile. The string is case sensitive and limited to 32 ASCII 7-bit printable characters with no spaces.

radius-accounting-policy

Syntax **radius-accounting-policy** *name*
no radius-accounting-policy

Context config>subscr-mgmt
config>subscr-mgmt>sub-prof

Description This command specifies a subscriber RADIUS based accounting policy.

Parameters *name* — The name of the policy. The string is case sensitive and limited to 32 ASCII 7-bit printable characters with no spaces.

accept-authorization-change

Syntax [**no**] **accept-authorization-change**

Context config>subscr-mgmt>auth-policy

Description This command specifies whether or not the system should handle the CoA messages initiated by the RADIUS server, and provide for mid-session interval changes of policies applicable to subscriber hosts.

Default no accept-authorization-change

accept-script-policy

Syntax **accept-script-policy** *policy-name*
no accept-script-policy

Context config>subscr-mgmt>auth-policy

Description This command configures a RADIUS script policy used to change the RADIUS attributes of the incoming Access-Accept messages.

Parameters *policy-name* — Configures a Python script policy to modify Access-Accept messages.

access-loop-options

Syntax [**no**] **access-loop-options**

Context config>subscr-mgmt>auth-plcy>include-radius-attribute
config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables inclusion of access loop information: Broadband Forum (BBF) access loop characteristics, DSL line state and DSL type. The BBF access loop characteristics are returned as BBF specific RADIUS attributes where DSL line state and DSL type are returned as Alcatel-Lucent specific RADIUS VSA's.

Information obtained via the ANCP protocol has precedence over information received in PPPoE Vendor Specific BBF tags or DHCP Vendor Specific BBF Options.

If ANCP is utilized and interim accounting update is enabled, any "Port Up" event from GSMP will initiate in an interim update. "Port Up" messages can include information such as an update on the current subscriber actual-upstream-speed. The next interim accounting message will be from "port up" triggering point.

Default no access-loop-options

host-accounting

Syntax [**no**] **host-accounting** [**interim-update**]

Context config>subscr-mgmt>acct-plcy

Description This command enables per host accounting mode. In host accounting mode, the acct-session-id is generated per host. This acct-session-id is uniformly included in all accounting messages (START/INTERIM-UPDTATE/STOP) and it can be included in RADIUS Access-Request message.

Accounting counters are based on the queue counters and as such are aggregated for all host sharing the queues within an sla-profile instance (non HSMDA) or a subscriber (HSMDA). CoA and LI is supported based on the acct-session-id of the host.

Default no host-accounting

Parameters **interim-update** — Without this keyword only START and STOP accounting messages are generated when the host is established/terminated. This is equivalent to a time-based accounting where only the duration of the session is required.

include-radius-attribute

Syntax [**no**] **include-radius-attribute**

Context config>subscr-mgmt>auth-plcy
config>subscr-mgmt>acct-plcy

Description This command enables the context to specify the RADIUS parameters that the system should include into RADIUS authentication-request messages.

use-std-acct-attributes

Syntax [**no**] **use-std-acct-attributes**

Context config>subscr-mgmt>radius-acct-policy>include-radius-attribute

Description This command sets use of RADIUS standard define attributes for sending accounting information.

Default Per default ALU VSAs are sent.

acct-authentic

Syntax [**no**] **acct-authentic**

Context config>subscr-mgmt>auth-policy>include-radius-attribute
config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables the generation of the acct-authentic RADIUS attribute.

acct-delay-time

Syntax [**no**] **acct-delay-time**

Context config>subscr-mgmt>auth-policy>include-radius-attribute
config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables the generation of the acct-delay-time RADIUS attribute.

all-authorized-session-addresses

Syntax [**no**] **all-authorized-session-addresses**

Context config>subscr-mgmt>acct-plcy>include-radius-attribute

Description Applicable for session-accounting mode only.

With this flag enabled, all IP address attributes explicitly enabled to be included are the following:

- delegated-ipv6-prefix
- framed-ip-address
- framed-ip-netmask
- framed-ipv6-prefix
- ipv6-address

These are included if the corresponding addresses or prefixes are authorized (via access-accept or ludb) and independent if they are used or not.

Default no all-authorized-session-addresses

called-station-id

Syntax [**no**] **called-station-id**

Context config>subscr-mgmt>auth-policy>include-radius-attribute
config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command includes called station id attributes.

The **no** form of the command excludes called station id attributes.

calling-station-id

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

Context config>service>ies>if>sap
config>service>ies>sub-if>grp-if>sap
config>service>vpls>sap
config>service>vprn>if>sap
config>service>vprn>sub-if>grp-if>sap
config>subscr-mgmt>auth-plcy>include-radius-attribute
config>subscr-mgmt>acct-plcy>include>include-radius-attribute

Description This command enables the inclusion of the calling-station-id attribute in RADIUS authentication requests and RADIUS accounting messages. The value inserted is set at the SAP level. If no **calling-station-id** value is set at the SAP level, the **calling-station-id** attribute will not be sent.

Default no calling-station-id

Parameters **mac —** Specifies that the mac-address will be sent.

remote-id — Specifies that the remote-id will be sent.

sap-id — Specifies that the sap-id will be sent.

sap-string — Specifies that the value is the inserted value set at the SAP level. If no **calling-station-id** value is set at the SAP level, the **calling-station-id** attribute will not be sent.

access-loop-options

Syntax [**no**] **access-loop-options**

Context config>subscr-mgmt>auth-plcy>include-radius-attribute
config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables inclusion of access loop information: Broadband Forum (BBF) access loop characteristics, DSL line state and DSL type. The BBF access loop characteristics are returned as BBF specific RADIUS attributes where DSL line state and DSL type are returned as Alcatel-Lucent specific RADIUS VSA's.

Information obtained via the ANCP protocol has precedence over information received in PPPoE Vendor Specific BBF tags or DHCP Vendor Specific BBF Options.

acct-session-id

Syntax [**no**] **acct-session-id**

Context configure>subscr-mgmt>auth-plcy>include-radius-attribute

Description The **acct-session-id** attribute for each subscriber host will be generated at the very beginning of the session initiation. This command will enable or disable sending this attribute to the RADIUS server in the Access-Request messages regardless of whether the accounting is enabled or not. The **acct-session-id** attribute can be used to address the subscriber hosts from the RADIUS server in the CoA Request.

The `acct-session-id` attribute will be unique per subscriber host network wide. It is a 22bytes long field comprised of the system MAC address along with the creation time and a sequence number in a hex format.

Default Disabled

circuit-id

Syntax [**no**] **circuit-id**

Context config>subscr-mgmt>auth-policy>include-radius-attribute
config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables the generation of the agent-circuit-id for RADIUS.

delegated-ipv6

Syntax [**no**] **delegated-ipv6**

Context config>subscr-mgmt>auth-policy>include-radius-attribute
config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables the generation of the delegated-ipv6 RADIUS attribute.

dhcp-options

Syntax [**no**] **dhcp-options**

Context config>subscr-mgmt>auth-plcy>include-radius-attribute

Description This command enables insertion of RADIUS VSA containing all dhcp-options from dhcp-discover (or dhcp-request) message. The VSA contains all dhcp-options in a form of the string. If required (the total length of all dhcp-options exceeds 255B), multiple VSAs are included.

Default no dhcp-options

dhcp-vendor-class-id

Syntax [**no**] **dhcp-vendor-class-id**

Context config>subscr-mgmt>auth-plcy>include-radius-attribute

Description This command includes the "[26-6527-36] Alc-DHCP-Vendor-Class-Id" attribute in RADIUS accounting messages. The content of the DHCP Vendor-Class-Identifier option (60) is mapped in this attribute.

Default no dhcp-vendor-class-id

framed-interface-id

Syntax [**no**] **framed-interface-id**

Context config>subscr-mgmt>auth-policy>include-radius-attribute
config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables the generation of the framed-interface-id RADIUS attribute.

framed-ip-addr

Syntax [**no**] **framed-ip-addr**

Context config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables the inclusion of the framed-ip-addr attribute.

framed-ip-netmask

Syntax [**no**] **framed-ip-netmask**

Context config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables the inclusion of the framed-ip-netmask attribute.

framed-ipv6-prefix

Syntax [**no**] **framed-ipv6-prefix**

Context config>subscr-mgmt>auth-policy>include-radius-attribute
config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables the generation of the framed-ipv6-prefix RADIUS attribute.

framed-ipv6-route

Syntax [**no**] **framed-ipv6-route**

Context config>subscr-mgmt>acct-plcy>include-radius-attribute

Description When enabled, all valid [99] Framed-IPv6-Route attributes as received in the RADIUS authentication phase and associated with an instantiated IPv6 wan host will be included in the RADIUS accounting request messages. The state of the Framed-IPv6-Route (installed, shadowed, hostInactive, etc.) is not taken into account for reporting in the accounting request messages.

Default no framed-ipv6-route

framed-route

Syntax [**no**] **framed-route**

Context config>subscr-mgmt>acct-plcy>include-radius-attribute

Description When enabled, all valid [22] Framed-Route attributes as received in the RADIUS authentication phase and associated with an instantiated IPv4 host will be included in the RADIUS accounting request messages. The state of the Framed-Route (installed, shadowed, hostInactive, etc.) is not taken into account for reporting in the accounting request messages.

Default no framed-route

ipv6-address

Syntax [**no**] **framed-ipv6-address**

Context config>subscr-mgmt>auth-policy>include-radius-attribute
config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables the generation of the ipv6-address RADIUS attribute.

mac-address

Syntax [**no**] **mac-address**
config>subscr-mgmt>auth-policy>include-radius-attribute
config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables the generation of the client MAC address RADIUS attribute.

nas-identifier

Syntax [**no**] **nas-identifier**

Context config>subscr-mgmt>auth-policy>include-radius-attribute
config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables the generation of the nas-identifier RADIUS attribute.

nas-port

Syntax [**no**] **nas-port** *bit-specification binary-spec*

Context config>subscr-mgmt>auth-policy>include-radius-attribute
config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables the generation of the nas-port RADIUS attribute. You enter decimal representation of a 32-bit string that indicates your port information. This 32-bit string can be compiled based on different information from the port (data types). By using syntax number-of-bits data-type you indicate how many bits from the 32 bits are used for the specific data type. These data types can be combined up to 32 bits in total. In between the different data types 0's and/or 1's as bits can be added.

The **no** form of this command disables your nas-port configuration.

Parameters *bit-specification binary-spec* — Specifies the NAS-Port attribute

Values

binary-spec	<bit-specification> <binary-spec>
bit-specification	0 | 1 | <bit-origin>
bit-origin	*<number-of-bits><origin>
number-of-bits	1 — 32
origin	o | i | s | m | p

o outer VLAN ID
i inner VLAN ID
s slot number
m MDA number
p port number or lag-id

Sample

```
*12o*12i00*2s*2m*2p => oooo oooo oooo iiii iiii iiii 00ss mmpp
If outer vlan = 0 & inner vlan = 1 & slot = 3 & mda = 1 & port = 1
=>  0000 0000 0000 0000 0000 0001 0011 0101 => nas-port = 309
```

nas-port-id

Syntax [**no**] **nas-port-id** [**prefix-string** *string*] [**suffix** *suffix-option*]

Context config>subscr-mgmt>auth-policy>include-radius-attribute
config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables the generation of the nas-port-id RADIUS attribute. Optionally, the value of this attribute (the SAP-id) can be prefixed by a fixed string and suffixed by the circuit-id or the remote-id of the client connection. If a suffix is configured, but no corresponding data is available, the suffix used will be 0/0/0/0/0/0.

Parameters **prefix-string** *string* — Specifies that a user configurable string will be added to the RADIUS NAS port attribute, up to 8 characters in length.

suffix *suffix-option* — Specifies the suffix type to be added to the RADIUS NAS oort attribute.

Values circuit-id, remote-id

nas-port-type

Syntax **nas-port-type**
nas-port-type [0..255]
no nas-port-type

Context config>subscr-mgmt>auth-plcy>include-radius-attribute
config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables the generation of the nas-port-type RADIUS attribute. If set to **nas-port-type**, the following will be sent: values: 32 (null-encap), 33 (dot1q), 34 (qinq), 15 (DHCP hosts). The **nas-port-type** can also be set as a specified value, with an integer from 0 to 255.

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

Default no nas-port-type

Parameters **0 — 255** — Specifies an enumerated integer that specifies the value that will be put in the RADIUS nas-port-type attribute.

nat-port-range

Syntax [**no**] **nat-port-range**

Context config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables the generation of the of nat-port-range attribute.

Default no nat-port-range

pppoe-service-name

Syntax [**no**] **pppoe-service-name**

Context config>subscr-mgmt>auth-policy>include-radius-attribute
config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables the generation of the pppoe-service-name RADIUS attribute.

remote-id

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

Context config>subscr-mgmt>auth-policy>include-radius-attribute
config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables the generation of the agent-remote-id for RADIUS.

tunnel-server-attrs

Syntax [**no**] **tunnel-server-attrs**

Context config>subscr-mgmt>auth-policy>include-radius-attribute

Description This command includes tunnel-server attribute.

sla-profile

Syntax [**no**] **sla-profile**

Context config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command specifies that SLA profile attributes should be included into RADIUS accounting messages.

sub-profile

Syntax [**no**] **sub-profile**

Context config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command specifies that subscriber profile attributes should be included into RADIUS accounting messages.

subscriber-id

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

Context config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command specifies that subscriber ID attributes should be included into RADIUS accounting messages.

tunnel-server

Syntax [**no**] **tunnel-server**

Context config>subscr-mgmt>auth-policy>include-radius-attribute
config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables the generation of the tunnel-server RADIUS attribute.

user-name

Syntax [**no**] user-name

Context config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables the inclusion of the user-name attribute.

The **no** form of the command disables the inclusion of the user-name attribute.

Default no user-name

password

Syntax **password** *password* [**hash** | **hash2**]
no password

Context config>subscr-mgmt>auth-policy

Description This command sets a password that is sent with **user-name** in every RADIUS authentication request sent to the RADIUS server upon receipt of DHCP discover or request messages. If no password is provided, an empty password will be sent.

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

Default none

Parameters *password* — A text string containing the password. Allowed values are any string up to 64 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

hash — Specifies the key is entered in an encrypted form. If the **hash** parameter is not used, the key is assumed to be in a non-encrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** parameter specified.

hash2 — Specifies the key is entered in a more complex encrypted form. If the **hash2** parameter is not used, the less encrypted **hash** form is assumed.

ppp-user-name

Syntax **ppp-user-name append** *domain-name*
ppp-user-name default-domain *domain-name*
ppp-user-name replace *domain-name*
ppp-user-name strip
no ppp-user-name

Context config>subscr-mgmt>auth-plcy

Description The command configures the PPP user name operation. Use this command to specify the operation to perform on the PAP/CHAP user name.

The no form of the command reverts to the default.

Default no ppp-user-name

Parameters **append** *domain-name* — Astring specified by tmnxSubAuthPlcyPppDomain, preceded with a '@', is appended to the PAP/CHAP user name.

default-domain *domain-name* — The same action is performed as with appendDomain, but only if the PAP/CHAP user name does not already contain a domain name.

replace ***domain-name*** — All characters after a '@' delimiter are replaced with the string specified by tmnxSubAuthPlcyPppDomain.

strip — Any '@' character and all subsequent characters are removed from the PAP/CHAP user name.

pppoe-access-method

Syntax **pppoe-access-method** {**none** | **padi** | **pap-chap**}
no pppoe-access-method

Context config>subscr-mgmt>auth-plcy

Description This command indicates the authentication method used towards the RADIUS server in case the policy is used for PPPoE.

Parameters **none** — Indicates that the client will be authenticated by the local user database defined under the group interface and not through RADIUS.

padi — Indicates that the client will be authenticated by RADIUS as soon as the PADI packet comes in (there is no PPP authentication done in the session in this case).

pap-chap — Indicates that the RADIUS authentication of the client will be delayed until the authentication protocol phase in the PPP session (PAP or CHAP) and authentication will be performed with the user name and PAP password / CHAP response supplied by the client.

queue-instance-accounting

Syntax **queue-instance-accounting** [**interim-update**]
no queue-instance-accounting

Context config>subscr-mgmt>acct-plcy

Description This command enables per queue-instance-accounting. A stream of accounting messages (START/INTERIM-UPDATE/STOP) is generated per queuing instance. A queuing instance is equivalent to an sla-profile instance on non HSMDA based hardware and to subscriber on HSMDA based hardware. Accounting session id is generated per queuing instance and this accounting session id CANNOT be included in RADIUS Access-Request message. Queue instance counters represent volume based aggregation for all hosts sharing the queuing instance.

CoA and LI is supported based on the acct-session-id of the queuing instance.

Default interim-update

Parameters **interim-update —** specifies whether accounting messages are sent for the queue-instance. The queue-instance is the subscriber on High Scale MDA (HSMDA), or the SLA profile instance otherwise.

radius-authentication-server

Syntax **radius-authentication-server**

Context config>subscr-mgmt>acct-plcy

Description This command creates the context for defining RADIUS authentication server attributes under a given session authentication policy.

access-algorithm

Syntax **access-algorithm** {**direct** | **round-robin**}
no access-algorithm

Context config>subscr-mgmt>auth-plcy-srvr
config>subscr-mgmt>acct-plcy>server

Description This command configures the algorithm used to access the list of configured RADIUS servers.

Parameters **direct —** Specifies that the first server will be used as primary server for all requests, the second as secondary and so on.

round-robin — Specifies that the first server will be used as primary server for the first request, the second server as primary for the second request, and so on. If the router gets to the end of the list, it starts again with the first server.

fallback-action

Syntax **fallback-action accept**
fallback-action user-db *local-user-db-name*
no fallback-action

Context config>subscr-mgmt>auth-plcy-srvr
config>subscr-mgmt>auth-plcy

Description This command configures the action when no RADIUS server is available.

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

Default no fallback-action

hold-down-time

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

Context config>subscr-mgmt>auth-plcy>radius-auth-server

Description This command determines the interval during which no new communication attempts will be made to a RADIUS server that is marked **down** to prevent immediately overloading the server when it is starting up. The only exception is when all servers in the authentication policy are marked **down**; in that case they will all be used again to prevent failures on new client connections.

Default 30

Parameters *seconds* — Specifies the hold time before re-using a RADIUS server that was down.

Values 30 — 900

router

Syntax **router** *router-instance*
router *service-name*
no router

Context config>subscr-mgmt>auth-plcy-srvr
config>subscr-mgmt>acct-plcy>server

Description This command specifies the virtual router instance applicable for the set of configured RADIUS servers. This value cannot be changed once a RADIUS server is configured for this policy. When the value is zero, both base and management router instances are matched.

Parameters *router-instance* — Specifies the virtual router instance.

Values

router-name:	Base, management
service-id:	1 — 2147483647
service-name:	Specifies the service name up to 64 characters in length.

retry

Syntax **retry** *count*
no retry

Context config>subscr-mgmt>auth-plcy-srvr
config>subscr-mgmt>acct-plcy>server

Description This command configures the number of times the router attempts to contact the RADIUS server for authentication, if not successful the first time.

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

Default 3

Parameters *count* — The retry count.

Values 1 — 10

radius-server-policy

Syntax **radius-server-policy <radius-server-policy-name>**
no radius-server-policy

Context config>subscr-mgmt>auth-plcy
config>subscr-mgmt>acct-plcy

Description This command references an existing radius-server-policy (available under the config>aaa context) for use in subscriber management authentication and accounting.

When configured in an authentication-policy, following CLI commands are ignored in the policy to avoid conflicts:

- all commands in the radius-authentication-server context
- accept-authorization-change
- coa-script-policy
- accept-script-policy
- request-script-policy

When configured in a radius-accounting-policy, following CLI commands are ignored in the policy to avoid conflicts:

- all commands in the radius-accounting-server context
- acct-request-script-policy

The **no** form of the command removes the radius-server-policy reference from the configuration

Default no radius-server-policy

Parameters *radius-server-policy-name* — Specifies the RADIUS server policy.

server

Syntax **server** *server-index* **address** *ip-address* **secret** *key* [**hash** | **hash2**] [**port** *port-num*] [**coa-only**] [**pending-requests-limit** *limit*]
no server *index*

Context config> subscr-mgmt>auth-policy>radius-auth-server
config>subscr-mgmt>acct-plcy>server

Description This command adds a RADIUS server and configures the RADIUS server IP address, index, and key values.

Up to sixteen RADIUS servers can be configured at any one time in a RADIUS authentication policy. Only five can be used for authentication, all other servers should be configured as coa-only servers. In a RADIUS accounting policy, up to five RADIUS servers can be configured. RADIUS servers are accessed in order from lowest to highest index for authentication or accounting requests until a response from a server is received. A higher indexed server is only queried if no response is received from a lower indexed server (which implies that the server is not available). If a response from a server is received, no other RADIUS servers are queried.

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

Default No RADIUS servers are configured.

Parameters *server-index* — The index for the RADIUS server. The index determines the sequence in which the servers are queried for authentication requests. Servers are queried in order from lowest to highest index.

Values 1 — 16 (a maximum of 5 authentication servers)

address *ip-address* **—** The IP address of the RADIUS server. Two RADIUS servers cannot have the same IP address. An error message is generated if the server address is a duplicate.

secret *key* **—** The secret key to access the RADIUS server. This secret key must match the password on the RADIUS server.

Values secret-key: Up to 20 characters in length.
hash-key: Up to 33 characters in length.
hash2-ke: Up to 55 characters in length.

hash — Specifies the key is entered in an encrypted form. If the **hash** parameter is not used, the key is assumed to be in a non-encrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** parameter specified.

hash2 — Specifies the key is entered in a more complex encrypted form. If the **hash2** parameter is not used, the less encrypted **hash** form is assumed.

port *port-num* **—** Specifies the UDP port number on which to contact the RADIUS server for authentication.

Values 1 — 65535

coa-only — Specifies Change-of-Authorization Messages only. Servers that are marked with the coa-only flag will not be used for authentication, but they will be able to accept RADIUS CoA messages, independent of the accept-authorization-change setting in the authentication policy.

For authentication purposes, the maximum number of servers is 5. All other servers may only be used as coa-only servers.

pending-requests-limit *limit* — Specifies the maximum number of outstanding RADIUS authentication requests for this authentication server.

Default The default value when not configured is 4096.

Values 1 — 4096

source-address

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

Context config>subscr-mgmt>auth-plcy-srvr
config>subscr-mgmt>acct-plcy>server

Description This command configures the source address of the RADIUS packet.

The system IP address must be configured in order for the RADIUS client to work. See Configuring a System Interface in the 7750 SR OS Router Configuration Guide. Note that the system IP address must only be configured if the source-address is not specified. When the **no source-address** command is executed, the source address is determined at the moment the request is sent. This address is also used in the nas-ip-address attribute: over there it is set to the system IP address if **no source-address** was given.

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

Default System IP address

Parameters *ip-address* — The IP prefix for the IP match criterion in dotted decimal notation.

Values 0.0.0.0 - 255.255.255.255

timeout

Syntax **timeout** *seconds*
no timeout

Context config>subscr-mgmt>auth-plcy-srvr
config>subscr-mgmt>acct-plcy>server

This command configures the number of seconds the router waits for a response from a RADIUS server.

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

Default 3 seconds

Parameters *seconds* — The number of seconds the router waits for a response from a RADIUS server, expressed as a decimal integer.

Values 1 — 90

session-accounting

Syntax **session-accounting** [**interim-update**] [**host-update**]
no session-accounting

Context config>subscr-mgmt>acct-plcy

Description This command enables per session accounting mode. In per session accounting mode, the acct-session-id is generated per session. This acct-session-id is uniformly included in all accounting messages (START/INTERIM-UPDTATE/STOP) and it can be included in RADIUS Access-Request message.

This accounting mode of operation can be used only in PPPoE environment with dual-stack host in which case both hosts (IPv4 and IPv6) are considered part of the same session. In addition to regular interim-updates, *triggered* interim-updates are sent by a host joining or leaving the session.

When an IPv4/v6 address is allocated, or released from a dual-stack host, a triggered interim-update message is immediately sent. This triggered interim-update message reflects the change in the IP address. The triggered interim-update has no effect on the interval at which the regular interim updates are scheduled.

Accounting counters are based on the queue counters and as such are aggregated for all host sharing the queues within an sla-profile instance (non HSMDA) or a subscriber (HSMDA).

CoA and LI is supported based on the acct-session-id of the session.

Default no session-accounting

Parameters **interim-update —** Without this keyword only START and STOP accounting messages are generated when the session is established/terminated. This is equivalent to a time-based accounting where only the duration of the session is required.

host-update — This keyword indicates that host updates messages are sent. INTERIM-UPDATE messages can be generated (volume based accounting) by selecting this keyword..

session-id-format

Syntax **session-id-format** {**description** | **number**}
no session-id-format

Context config>subscr-mgmt>acct-plcy

Description This command specifies the format for the acct-session-id attribute used in RADIUS accounting requests.

Parameters **description —** Specifies to use a string containing following information <subscriber>@<sap-id>@<SLA-profile>_<creation-time>.

number — Specifies to use a unique number generated by the OS to identify a given session.

update-interval

Syntax **update-interval** *minutes*
no update-interval

Context config>subscr-mgmt>acct-plcy

Description This command specifies the interval at which accounting data of subscriber hosts will be updated in a RADIUS Accounting Interim-Update message. Requires interim-update to be enabled when specifying the accounting mode in the radius accounting policy.

A RADIUS specified interim interval (attribute [85] Acct-Interim-Interval) overrides the CLI configured value.

Parameters *minutes* — Specifies the interval, in minutes, at which accounting data of subscriber hosts will be updated.

Values 5 — 259200

update-interval-jitter

Syntax **update-interval-jitter absolute** *seconds*
no update-interval-jitterl

Context config>subscr-mgmt>acct-plcy

Description This command specifies the absolute maximum random delay introduced on the update interval between two accounting interim update messages. The effective maximum random delay value is the minimum of the configured absolute jitter value and 10% of the configured update-interval.

A value of zero will send the accounting interim update message without introducing an additional random delay.

The **no** form of the command sets the default to 10% of the configured update-interval.

Default no update-interval-jitter

This corresponds with 10% of the configured update-interval

Parameters **absolute** *seconds* — specifies the absolute maximum jitter value in seconds.

Values 0 — 36000

use-std-acct-attributes

Syntax [**no**] **use-std-acct-attributes**

Context config>subscr-mgmt>radius-acct-policy

Description This command enables the use of standard accounting attributes.

The **no** form of the command disables the use of standard accounting attributes.

re-authentication

Syntax [**no**] **re-authentication**

Context config>subscr-mgmt>auth-policy

Description This command enables authentication process at every DHCP address lease renewal s only if RADIUS did not reply any special attributes (for example, authentication only, no authorization).

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

Default disabled

request-script-policy

Syntax **request-script-policy** *policy-name*
no request-script-policy

Context config>subscr-mgmt>auth-policy

Description This command specifies the RADIUS script policy used to change the RADIUS attributes of the outgoing Access-Request messages.

Default none

Parameters *policy-name* — Configures a Python script policy to modify Access-Request messages.

send-acct-stop-on-fail

Syntax **send-acct-stop-on-fail** {[**on-request-failure**] [**on-reject**] [**on-accept-failure**]}
no send-acct-stop-on-fail

Context config>subscr-mgmt>auth-policy

Description This command activates the reporting of RADIUS authentication failures of a PPPoE session to a RADIUS accounting server with an Accounting Stop message.

Three failure categories can be enabled separately:

- **on-request-failure**: All failure conditions between the sending of an Access-Request and the reception of an Access-Accept or Access-Reject.
- **on-reject**: When an Access-Reject is received
- **on-accept-failure**: All failure conditions that appear after receiving an Access-Accept and before successful instantiation of the host or session.

The RADIUS accounting policy to be used for sending the Accounting Stop messages must be obtained prior to RADIUS authentication via local user database pre-authentication.

Default no send-acct-stop-on-fail

user-name-format

Syntax **user-name-format** *format* [**mac-format** *mac-format*]
user-name-format *format* **append** [*domain-name*] [**mac-format** *mac-format*]
user-name-format *format* **append domain-name**
user-name-format *format* **default-domain** *domain-name*[**mac-format** *mac-format*]
user-name-format *format* **replace** *domain-name*[**mac-format** *mac-format*]
user-name-format *format* **strip** [**mac-format** *mac-format*]
no user-name-format

Context config>subscr-mgmt>auth-policy

Description This command defines the format of the "user-name" field in the session authentication request sent to the RADIUS server.

The **no** form of the command switches to the default format, **mac**.

Default By default, the MAC source address of the DHCP DISCOVER message is used in the user-name field.

Parameters *format* — Specifies the user name format in RADIUS message.

Values ascii-converted-circuit-id, ascii-converted-tuple, circuit-id, dhcp-client-vendor-opts, mac, mac-giaddr, tuple

ascii-converted-circuit-id — Identical to circuit-id, but the user name will be sent to the RADIUS server as a string of hex digits, for use if there is binary data in the circuit-id.

ascii-converted-tuple — Identical to tuple, but the circuit-id part of the user name will be sent to the RADIUS server as a string of hex digits, for use if there is binary data in the circuit-id.

circuit-id — If the system serves as a DHCP relay server which inserts option 82 info, the user name will be formatted as defined under DHCP information option. If the system is not a DHCP relay server, the circuit-id will be taken from option 82 in the received DHCP message. If no circuit-id can be found, the DHCP-msg is rejected.

dhcp-client-vendor-opts — Creates a concatenation of the DHCP client-identifier option (option 60), a "@" delimiter and the DHCP vendor-class identifier options. The two option strings are parsed for any characters which are non-printing are considered invalid and must be converted to underscore "_" characters. In addition, any space character (hex 20) and @ character (hex 40) are also converted to underscore. The character set considered valid is ASCII hex 21 through hex 3F, and hex 41 through hex 7E. Any character outside this set will be converted into an underscore (hex 5F) character.

mac — The MAC source address of the DHCP DISCOVER message is used in the user-name field. The format of the MAC address string used as the user name in the RADIUS authentication requests uses lowercase hex digits, and ":" as the inter-digit separator, for example, 00:11:22:aa:bb:cc is valid but 00-11-22-AA-BB-CC will return an error. The RADIUS server must be configured accordingly, otherwise the authentication request will fail.

mac-giaddr — Specifies that MAC giaddr indicates the format used to identify the user towards the RADIUS server.

tuple — The concatenation of MAC source address and circuit-ID are used in the user-name field.

mac-format — Specifies how a MAC address is represented when contacting a RADIUS server. This is only used while the value of is equal to the DHCP client vendor options and if the MAC address is used by default of the DHCP client vendor options.

Examples:	ab:	00:0c:f1:99:85:b8 Alcatel-Lucent 7xxx style
	XY-	00-0C-F1-99-85-B8 IEEE canonical style
	mmmm.	0002.03aa.abff Cisco style

append — Specifies the data type which is is an enumerated integer that indicates what needs to be appended to the user-name sent to the RADIUS server.

Values 1 — nothing
2 — domain name

domain — In some instances it is desired to add a domain only to usernames which have omitted the domain (@domain). In these instances a default-domain can be appended to usernames which lack a @domain

append — Adds a "@" delimiter and the specified string after the PAP/CHAP username. No allowance is made for the presence of an existing domain or @ delimited.

replace — Replaces the character-string after the "@" delimiter with the string specified.

strip — Removes all characters after and including the "@" delimiter.

Example:

```
Command: append
String: domainA-1.com
PAP/CHAP User:someuser
Resulting User:someuser@domainA-1.com

Command: append
String: domainA-1.com
PAP/CHAP User:someuser@existing-domain.net
Resulting User:someuser@existing-domain.net@domainA-1.com

Command: strip
String:
PAP/CHAP User:someuser@existing-domain.net
Resulting User:someuser

Command: replace
String: domainA-1.com
PAP/CHAP User:someuser@existing-domain.net
Resulting User:someuser@domainA-1.com

Command: default-domain
String:domainA-1.com
PAP/CHAP User:someuser@existing-domain.net
Resulting User:someuser@existing-domain.net

Command: default-domain
String: domainA-1.com
PAP/CHAP User:someuser
Resulting User:someuser@domainA-1.com
```

Diameter Policy Commands

diameter-policy

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

Context config>sub-mgmt

Description This command configures aa diameter policy.

Parameters *diameter-policy-name* — Specifies the name of the diameter policy. There is a maximum of 32 displayable characters.

diameter-base

Syntax **diameter-base**

Context config>sub-mgmt>diameter-policy

Description This command configures a diameter base.

peer

Syntax **peer** *name* [**create**]
no peer name

Context config>sub-mgmt>diameter-policy>diameter-base

Description This command creates a new peer.

Default no peer is defined

Parameters *name* — Specifies the peer name, up to a maximum of 32 displayable characters

address

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

Context config>sub-mgmt>diameter-policy>diameter-base>peer

Description This command configures the address.

Parameters *ip-address* — Specifies the IPv4 or IPv6 address.

transport

Syntax **transport** {**tcp | stcp**} [**port** *port-num*]
no transport

Context config>sub-mgmt>diameter-policy>diameter-base>peer

Description This command specifies the transport protocol and transport port.

Default tcp transport and 3868 as port number

Parameters *port-num* — Specifies the port number.

Values 1-6535

tcp | stcp — Specifies the protocol.

destination-host

Syntax **destination-host** *destination-host-string*
no destination-host

Context config>sub-mgmt>diameter-policy>diameter-base>peer

Description This command configurs the destination host.

Default no destination-host

Parameters *destination-host-string* — Specifies the destination host name, maximum 80 displayable characters.

destination-realm

Syntax **destination-realm** *destination-realm-string*
no destination-realm

Context config>sub-mgmt>diameter-policy>diameter-base>peer

Description This command configures the destination realm.

Default no destination-realm

Parameters *destination-realm-string* — Specifies the destination realm name, maximum 80 displayable characters.

watchdog-timer

Syntax [**no**] **watchdog-timer** *seconds*

Context config>sub-mgmt>diameter-policy>diameter-base>peer
config>sub-mgmt>diameter-policy>diameter-base

Description This command sets the maximum amount of time the node waits for a diameter peer to respond to awatchdog packet. When the watchdog timer expires, a DWR is sent to the diameter peer and the watchdog timer is reset. If a DWA is not received before the next expiration of the watchdog timer, transport failure to the diameter peer has occurred. The configuration at peer level overrules the value configured at diamater-base level for the given peer.

Default 30 seconds at diameter-base level

Default value at peer is taken from diameter-base.

Parameters *seconds* — Specifies the amount of time, in seconds.

Values 1-1000

connection-timer

Syntax [**no**] **connection-timer** *seconds*

Context config>sub-mgmt>diameter-policy>diameter-base>peer
config>sub-mgmt>diameter-policy>diameter-base

Description This command sets maximum amount of time the node attempts to reconnect to a diameter peer after a connection to the peer has been brought down due to a transport failure. There are certain exceptions to this rule, such as peer which terminated the transport connection indicating that it does not wish to communicate. A value of 0 means that the connection will not be retried. The configuration at peer level overrules the value configured at diamater-base level for the given peer.

Default 30 seconds at diameter-base level

The default value at peer is taken from diameter-base.

Parameters *seconds* — Specifies the amount of time, in seconds.

Values 1 — 1000

preference

Syntax **preference** *preference*
no preference

Context config>sub-mgmt>diameter-policy>diameter-base>peer
config>sub-mgmt>diameter-policy>diameter-base

Description This command configures the preference given to this DIAMETER policy peer with respect to the other peers associated with this DIAMETER policy.

If multiple peers are available for this policy, only the available peer with the highest preference will be used.

If multiple peers with the same preference are available, one of them will be used.

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

Default 50

Parameters *preference* — Specifies the preference of this DIAMETER policy peer.

Values 1 — 100

transaction-timer

Syntax [**no**] **transaction-timer** *seconds*

Context config>sub-mgmt>diameter-policy>diameter-base>peer
config>sub-mgmt>diameter-policy>diameter-base

Description This command sets maximum amount of time the node waits for a diameter peer to respond before trying another peer. The configuration at peer level overrules the value configured at diamater-base level for the given peer.

Default 30 seconds at diameter-base level

Default value at peer is taken from diameter-base.

Parameters *seconds* — Specifies the amount of time, in seconds.

Values 1-1000

vendor-support

Syntax **vendor-support** [**3gpp**|**vodafone**]
no vendor-support

Context config>sub-mgmt>diameter-policy>diameter-base

Description This command selects which vendor specific attributes are used in DCCA sessions

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

Default 3gpp

Parameters **3gpp —** Specifies the 3gpp diameter policy vendor type.

vodafone — Specifies the Vodaphone diameter policy vendor type.

router

Syntax **router service** s*ervice-name*
router *router-instance*
no router

Context config>sub-mgmt>diameter-policy>diameter-base

Description This command specifies the virtual router in which the diameter connection(s) will be established by this diameter policy.

Parameters *router-instance* — Specifies the router name.

Values router-instance: *router-name*|*service-id*
router-name: Base, management
service-id: 1 — 2147483647

Default Base

service-name — Specifies the VPRN service ID.

source-address

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

Context config>sub-mgmt>diameter-policy>diameter-base

Description This command configures the source address.

Default no source-address; system-ip address is used instead

Parameters *ip-address* — Specifies the UC IPv4 or IPv6 IP address.

dcca

Syntax **dcca**

Context config>sub-mgmt>diameter-policy

Description This command configures the DCCA.

cc-session-failover

Syntax [**no**] **cc-session-failover**

Context config>sub-mgmt>diameter-policy>dcca

Description This command configures the CC session failover.

Default cc-session-failover

cc-session-failover-handler

Syntax **cc-session-failover-handler** [{**terminate | continue | retry-and-terminate**}]
no cc-session-failover-handler

Context config>sub-mgmt>diameter-policy>dcca

Description This command configures the CC session failover handler.

Default terminate

tx-timer

Syntax [**no**] **tx-timer** *seconds*

Context config>sub-mgmt>diameter-policy>dcca

Description This command configures the TX timer.

Default 10 seconds

Parameters *seconds* — Specifies the TX timer in seconds.

Values 1-1000

include-avp

Syntax **include-avp**

Context config>sub-mgmt>diameter-policy>dcca

Description This command configures AVP.

service-context-id

Syntax **service-context-id** *string*
no service-context-id

Context config>sub-mgmt>diameter-policy>dcca>include-avp

Description This command configures the service context ID.

Default no service-context-id

Parameters *string* — Specifies the service context ID, maximum 32 displayable characters.

subscription-id

Syntax **subscription-id** {**circuit-id**|**subscriber-id**|**imsi**|**msisdn**|**imei**} [**type** {**e164**|**private**|**imsi**}]
no subscription-id

Context config>sub-mgmt>diameter-policy>dcca>include-avp

Description This command specifies specifies the origin of the information to send in the Subscription-Id-Data AVP.

The type parameter specifies the type of the identifier stored in the Subscription-Id-Data AVP.

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

Default subscriber-id
type private

3gpp-imsi

Syntax **3gpp-imsi** {**circuit-id**|**subscriber-id**|**imsi**}
no 3gpp-imsi

Context config>sub-mgmt>diameter-policy>dcca>include-avp

Description This command specifies the origin of th information to send in the AVP.
The **no** form of the command reverts to the default value.

Default subscriber-id

Parameters **circuit-id** — Specifies the circuit-id DCCA IMSI AVP parameter.
subscriber-id — Specifies the subscriber-id DCCA IMSI AVP parameter.
imsi — Specifies the imsi DCCA IMSI AVP parameter.

called-station-id

Syntax **called-station-id** *string*
no called-station-id

Context config>sub-mgmt>diameter-policy>dcca>include-avp

Description This command configures the called station ID.

Default no called-station-id

Parameters *string* — Specifies the called station ID, maximum 64 characters

radius-user-name

Syntax [**no**] **radius-user-name**

Context config>sub-mgmt>diameter-policy>dcca>include-avp

Description This command configures the RADIUS user name.

out-of-credit-reporting

Syntax **out-of-credit-reporting** {**final**|**quota-exhausted**}
no out-of-credit-reporting

Context config>sub-mgmt>diameter-policy>dcca

Description This command changes the reporting reason in an intermediate interrogation when the final granted units have been consumed and a corresponding out-of-credit-action different from "disconnect-host" is started.

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

Default final

Filter Commands

filter

Syntax **filter**

Context configure

Description This command manages the configuration of filters.

copy

Syntax **copy** {**ip-filter | mac-filter | ipv6-filter**} *src-filter-id* [**src-entry** *src-entry-id*] **to** *dst-filter-id* [**dst-entry** *dst-entry-id*] [**overwrite**]

Context configure>filter

Description This command copies filters and its entries.

Parameters *src-filter-id* — Specifies the source filter ID.

Values 1..65535

src-entry-id — Specifies the source entry ID.

Values 1..65535

dst-filter-id — Specifies the destination filter ID.

Values 1..65535

dst-entry-id — Specifies the destination entry ID.

Values 1..65535

overwrite — Specifies an overwrite.

ip-filter

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

Context configure>filter

Description This command configures an IP filter.

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

Values 1..65535

ipv6-filter

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

Context configure>filter

Description This command configures an IPv6 filter.

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

Values 1..65535

default-action

Syntax **default-action** *drop*|*forward*

Context configure>filter>ip-filter
configure>filter>ipv6-filter

Description This command configures default-action for the IP or IPv6 filter.

Parameters *drop*|*forward* — This keyword specifies the filter action.

entry

Syntax **entry** *entry-id* [**time-range** *time-range-name*] [**create**]
no entry *entry-id*

Context configure>filter>ip-filter
configure>filter>ipv6-filter

Description This command configures an IP or IPv6 filter entry.

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

Values 1..65535

time-range-name — Specifies the time range name.

Values 32 charas max

action

Syntax **action** *drop*|*forward*
no action

Context config>filter>ip-filter>entry
config>filter>ipv6-filter>entry

Description This command configures actions for the IP or IPv6 filter entry.

Parameters *drop|forward* — Specifies the filter action.

log

Syntax **log** *log-id*
no log

Context config>filter>ip-filter>entry
config>filter>ipv6-filter>entry

Description This command configures the log for the IP or IPv6 filter entry.

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

Values 101..199

match

Syntax **match** [**next-header** *next-header*]
no match

Context config>filter>ip-filter>entry
config>filter>ipv6-filter>entry

Description This command configures the match criteria for the IP or IPv6 filter entry.

Parameters *next-header* — Specifies the protocol numbers accepted in DHB.

Values [1..42|45..49|52..29|61..255]

Values none | crtp | crudp | egp | eigrp | encap | ether-i p | gre | icmp | idrp | igmp | igp | ip | ipv6 | ipv6-icmp | ipv6-no-nxt | isis | iso-ip | l2tp | ospf-igp | pim | pnni | ptp | rdp | rsvp | stp | tcp | udp | vrrp * udp/tcp wildcard

dscp

Syntax **dhcp**
no dhcp

Context config>filter>ip-filter>entry>match
config>filter>ipv6-filter>entry>match

Description This command configures DSCP match condition.

dst-ip

Syntax **dst-ip**
no dst-ip

Context config>filter>ip-filter>entry>match
config>filter>ipv6-filter>entry>match

Description This command configures the destination IP or IPv6 address match condition.

dst-port

Syntax **dst-port**
no dst-port

Context config>filter>ip-filter>entry>match
config>filter>ipv6-filter>entry>match

Description This command configures the destination port match condition.

icmp-code

Syntax **icmp-code**
no icmp-code

Context config>filter>ip-filter>entry>match
config>filter>ipv6-filter>entry>match

Description This command configures the ICMP code match condition.

icmp-type

Syntax **icmp-type**
no icmp-type

Context config>filter>ip-filter>entry>match
config>filter>ipv6-filter>entry>match

Description This command configures the ICMP type match condition.

src-ip

Syntax **src-ip**
no src-ip

Context config>filter>ip-filter>entry>match
config>filter>ipv6-filter>entry>match

Description This command configures the source IP or IPv6 address match condition.

src-port

Syntax **src-port**
no src-port

Context config>filter>ip-filter>entry>match
config>filter>ipv6-filter>entry>match

Description This command configures the source port match condition.

tcp-ack

Syntax **tcp-ack**
no tcp-ack

Context config>filter>ip-filter>entry>match
config>filter>ipv6-filter>entry>match

Description This command configures the TCP ACK match condition.

tcp-syn

Syntax **tcp-syn**
no tcp-syn

Context config>filter>ip-filter>entry>match
config>filter>ipv6-filter>entry>match

Description This command configures the TCP SYN match condition.

group-inserted-entries

Syntax **group-inserted-entries application** *application* **location** *location*

Context config>filter>ip-filter
config>filter>ipv6-filter

Description This command groups auto-inserted entries.

Parameters *application* — Specifies the application.

Values radius | credit-control

location — Specifies the location.

Values top | bottom

renum

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

Context config>filter>ip-filter
config>filter>ipv6-filter

Description This command renumbers an IP or IPv6 filter entry.

Parameters *old-entry-id* — Specifies the old entry ID to be renumbered.

Values 1..65535

new-entry-id — Specifies the new entry ID.

Values 1..65535

scope

Syntax **scope exclusive | template**
no scope

Context config>filter>ip-filter
config>filter>ipv6-filter

Description This command configures the scope for the IP or IPv6 filter.

Parameters **exclusive | template** — Specifies the type of policy.

shared-radius-filter-wmark

Syntax **shared-radius-filter-wmark low** *low-watermark* **high** *high-watermark*
no shared-radius-filter-wmark

Context config>filter>ip-filter
config>filter>ipv6-filter

Description This command defines the thresholds that will be used to raise a respective alarm when the number of shared filter copies increases.

Default no shared-radius-filter-wmark

Parameters *low-watermark* — specifies low threshold for the number of shared filter copies

Values 0-8000

high-watermark — specifies high threshold for the number of shared filter copies

Values 0-8000

sub-insert-radius

Syntax **sub-insert-radius start-entry** *entry-id* **count** *count*
no sub-insert-radius

Context config>filter>ip-filter
config>filter>ipv6-filter

Description This command defines the range of filter entries which will be reserved for entries created based on information (match criteria and action) from RADIUS auth-response messages.

The **no** version of the command disables the insertion, which means that information from auth-response messages cannot be stored in the filter, and the corresponding host will not be created in the system.

Default per default insertion is disabled

Parameters *entry-id* — An integer defining the lowest entry of the range.

count — An integer defining the number of entries in the range.

sub-insert-credit-control

Syntax **sub-insert-credit-control start-entry** *entry-id* **count** *count*
no sub-insert-credit-control

Context config>filter>ip-filter
config>filter>ipv6-filter

Description This command defines the range of filter entries that will be reserved for entries created based on information (match criteria and action) configured under the category-map configuration tree to enforce reduced-service level in case of credit exhaustion.

The **no** version of the command disables the insertion, which means that entries will not be installed even though the credit for the given category and subscriber-host has been exhausted.

Default per default insertion is disabled

Parameters *entry-id* — An integer defining the lowest entry of the range.

count — An integer defining the number of entries in the range.

sub-insert-shared-radius

Syntax **sub-insert-shared-radius start-entry** *entry-id* **count** *count*
no sub-insert-shared-radius

Context config>filter>ip-filter
config>filter>ipv6-filter
config>filter>ip-filter
config>filter>ipv6-filter

Description This command defines the range of filter entries that will be reserved for shared filter entries received in RADIUS messages.

The no version of the command disables the insertion resulting in a host setup failure when shared filter attributes are received in a RADIUS authentication response.

Default no sub-insert-shared-radius

Parameters *entry-id* — specifies the lowest entry of the range.

Values 1-65535

count — specifies the number of entries in the range.

Values 1-65535

sub-insert-wmark

Syntax **sub-insert-wmark** [**low** *percentage*] [**high** *percentage*]
no sub-insert-wmark

Context config>filter>ip-filter
config>filter>ipv6-filter

Description This command defines the thresholds that will be used to raise a respective alarm to provide monitoring of the resources for subscriber-specific filter insertion.

The **no** version of the command sets the default values for the respective thresholds.

Default low - 90%

high - 95%

Parameters *percentage* — Defines in percentage the threshold used to raise an alarm.

Values 1-100, integer

RADIUS Accounting Policy Custom Record Commands

custom-record

Syntax [**no**] **custom-record**

Context config>subscr-mgmt>acct-plcy

Description This command enables the context to configure the layout and setting for a custom accounting record associated with this accounting policy.

The **no** form of the command reverts the configured values to the defaults.

override-counter

Syntax [**no**] **override-counter** *override-counter-id*

Context config>log>acct-policy>cr

Description This command enables the context to configure Application Assurance override counter parameters.

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

Parameters *override-counter-id* — Specifies the override counter ID.

Values 1 — 8

e-counters

Syntax **e-counters** [**all**]
no e-counters

Context config>log>acct-policy>cr>override-cntr
config>log>acct-policy>cr>queue
config>log>acct-policy>cr>ref-override-cntr
config>log>acct-policy>cr>ref-queue

Description This command configures egress counter parameters for this custom record.

The **no** form of the command

Parameters **all** — Includes all counters.

i-counters

Syntax **i-counters** [**all**]
no i-counters

Context config>log>acct-policy>cr>override-cntr
config>log>acct-policy>cr>ref-override-cntr
config>log>acct-policy>cr>ref-queue

Description This command configures ingress counter parameters for this custom record.

The **no** form of the command

Parameters **all —** Includes all counters.

queue

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

Context config>log>acct-policy>cr

Description This command specifies the queue-id for which counters will be collected in this custom record. The counters that will be collected are defined in egress and ingress counters.

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

Parameters *queue-id* — Specifies the queue-id for which counters will be collected in this custom record.

in-profile-octets-discarded-count

Syntax [**no**] **in-profile-octets-discarded-count**

Context config>log>acct-policy>cr>oc>e-count
config>log>acct-policy>cr>roc>e-count
config>log>acct-policy>cr>queue>e-count
config>log>acct-policy>cr>ref-queue>e-count

Description This command includes the in-profile octets discarded count.

The **no** form of the command excludes the in-profile octets discarded count.

in-profile-octets-forwarded-count

Syntax [**no**] **in-profile-octets-forwarded-count**

Context config>log>acct-policy>cr>oc>e-count
config>log>acct-policy>cr>roc>e-count
config>log>acct-policy>cr>queue>e-count
config>log>acct-policy>cr>ref-queue>e-count

Description This command includes the in-profile octets forwarded count.

The **no** form of the command excludes the in-profile octets forwarded count.

in-profile-packets-discarded-count

Syntax [**no**] **in-profile-packets-discarded-count**

Context config>log>acct-policy>cr>oc>e-count
config>log>acct-policy>cr>roc>e-count
config>log>acct-policy>cr>queue>e-count
config>log>acct-policy>cr>ref-queue>e-count

Description This command includes the in-profile packets discarded count.

The **no** form of the command excludes the in-profile packets discarded count.

in-profile-packets-forwarded-count

Syntax [**no**] **in-profile-packets-forwarded-count**

Context config>log>acct-policy>cr>oc>e-count
config>log>acct-policy>cr>roc>e-count
config>log>acct-policy>cr>queue>e-count
config>log>acct-policy>cr>ref-queue>e-count

Description This command includes the in-profile packets forwarded count.

The **no** form of the command excludes the in-profile packets forwarded count.

out-profile-octets-discarded-count

Syntax [**no**] **out-profile-octets-discarded-count**

Context config>log>acct-policy>cr>oc>e-count
config>log>acct-policy>cr>roc>e-count
config>log>acct-policy>cr>queue>e-count
config>log>acct-policy>cr>ref-queue>e-count

Description This command includes the out of profile packets discarded count.

The **no** form of the command excludes the out of profile packets discarded count.

out-profile-octets-forwarded-count

Syntax [**no**] **out-profile-octets-forwarded-count**

Context config>log>acct-policy>cr>oc>e-count
config>log>acct-policy>cr>roc>e-count

config>log>acct-policy>cr>queue>e-count
config>log>acct-policy>cr>ref-queue>e-count

Description This command includes the out of profile octets forwarded count.

The **no** form of the command excludes the out of profile octets forwarded count.

out-profile-packets-discarded-count

Syntax [**no**] **out-profile-packets-discarded-count**

Context config>log>acct-policy>cr>oc>e-count
config>log>acct-policy>cr>roc>e-count
config>log>acct-policy>cr>queue>e-count
config>log>acct-policy>cr>ref-queue>e-count

Description This command includes the out of profile packets discarded count.

The **no** form of the command excludes the out of profile packets discarded count.

out-profile-packets-forwarded-count

Syntax [**no**] **out-profile-packets-forwarded-count**

Context config>log>acct-policy>cr>oc>e-count
config>log>acct-policy>cr>roc>e-count
config>log>acct-policy>cr>queue>e-count
config>log>acct-policy>cr>ref-queue>e-count

Description This command includes the out of profile packets forwarded count.

The **no** form of the command excludes the out of profile packets forwarded count.

The **no** form of the command

all-octets-offered-count

Syntax [**no**] **all-octets-offered-count**

Context config>log>acct-policy>cr>oc>i-count
config>log>acct-policy>cr>roc>i-count
config>log>acct-policy>cr>queue>i-count
config>log>acct-policy>cr>ref-queue>i-count

Description This command includes all octets offered in the count.

The **no** form of the command excludes the octets offered in the count.

Default no all-octets-offered-count

all-packets-offered-count

Syntax [**no**] **all-packets-offered-count**

Context config>log>acct-policy>cr>oc>i-count
config>log>acct-policy>cr>roc>i-count
config>log>acct-policy>cr>queue>i-count
config>log>acct-policy>cr>ref-queue>i-count

Description This command includes all packets offered in the count.

The **no** form of the command excludes the packets offered in the count.

Default no all-packets-offered-count

high-octets-discarded-count

Syntax [**no**] **high-octets-discarded-count**

Context config>log>acct-policy>cr>oc>i-count
config>log>acct-policy>cr>roc>i-count
config>log>acct-policy>cr>queue>i-count
config>log>acct-policy>cr>ref-queue>i-count

Description This command includes the high octets discarded count.

The **no** form of the command excludes the high octets discarded count.

Default no high-octets-discarded-count

high-octets-offered-count

Syntax [**no**] **high-octets-offered-count**

Context config>log>acct-policy>cr>oc>i-count
config>log>acct-policy>cr>roc>i-count
config>log>acct-policy>cr>queue>i-count
config>log>acct-policy>cr>ref-queue>i-count

Description This command includes the high octets offered count.

The **no** form of the command excludes the high octets offered count.

high-packets-discarded-count

Syntax [**no**] **high-packets-discarded-count**

Context config>log>acct-policy>cr>oc>i-count
config>log>acct-policy>cr>roc>i-count
config>log>acct-policy>cr>queue>i-count

config>log>acct-policy>cr>ref-queue>i-count

Description This command includes the high packets discarded count.

The **no** form of the command excludes the high packets discarded count.

Default no high-packets-discarded-count

high-packets-offered-count

Syntax [**no**] **high-packets-offered-count**

Context config>log>acct-policy>cr>oc>i-count
config>log>acct-policy>cr>roc>i-count
config>log>acct-policy>cr>queue>i-count
config>log>acct-policy>cr>ref-queue>i-count

Description This command includes the high packets offered count.

The **no** form of the command excludes the high packets offered count.

Default no high-packets-offered -count

in-profile-octets-forwarded-count

Syntax [**no**] **in-profile-octets-forwarded-count**

Context config>log>acct-policy>cr>oc>i-count
config>log>acct-policy>cr>roc>i-count
config>log>acct-policy>cr>queue>i-count
config>log>acct-policy>cr>ref-queue>i-count

Description This command includes the in profile octets forwarded count.

The **no** form of the command excludes the in profile octets forwarded count.

Default no in-profile-octets-forwarded-count

in-profile-packets-forwarded-count

Syntax [**no**] **in-profile-packets-forwarded-count**

Context config>log>acct-policy>cr>oc>i-count
config>log>acct-policy>cr>roc>i-count
config>log>acct-policy>cr>queue>i-count
config>log>acct-policy>cr>ref-queue>i-count

Description This command includes the in profile packets forwarded count.

The **no** form of the command excludes the in profile packets forwarded count.

Default no in-profile-packets-forwarded-count

low-octets-discarded-count

Syntax [**no**] **low-octets-discarded-count**

Context config>log>acct-policy>cr>oc>i-count
config>log>acct-policy>cr>roc>i-count
config>log>acct-policy>cr>queue>i-count
config>log>acct-policy>cr>ref-queue>i-count

Description This command includes the low octets discarded count.

The **no** form of the command excludes the low octets discarded count.

Default no low-octets-discarded-count

low-packets-discarded-count

Syntax [**no**] **low-packets-discarded-count**

Context config>log>acct-policy>cr>oc>i-count
config>log>acct-policy>cr>roc>i-count
config>log>acct-policy>cr>queue>i-count
config>log>acct-policy>cr>ref-queue>i-count

Description This command includes the low packets discarded count.

The **no** form of the command excludes the low packets discarded count.

Default no low-packets-discarded-count

low-octets-offered-count

Syntax [**no**] **low-octets-offered-count**

Context config>log>acct-policy>cr>oc>i-count
config>log>acct-policy>cr>roc>i-count
config>log>acct-policy>cr>queue>i-count
config>log>acct-policy>cr>ref-queue>i-count

Description This command includes the low octets discarded count.

The **no** form of the command excludes the low octets discarded count.

low-packets-offered-count

Syntax [**no**] **low-packets-offered-count**

Context config>log>acct-policy>cr>oc>i-count
config>log>acct-policy>cr>roc>i-count
config>log>acct-policy>cr>queue>i-count

config>log>acct-policy>cr>ref-queue>i-count

Description This command includes the low packets discarded count.

The **no** form of the command excludes the low packets discarded count.

out-profile-octets-forwarded-count

Syntax [**no**] **out-profile-octets-forwarded-count**

Context config>log>acct-policy>cr>oc>i-count
config>log>acct-policy>cr>roc>i-count
config>log>acct-policy>cr>queue>i-count
config>log>acct-policy>cr>ref-queue>i-count

Description This command includes the out of profile octets forwarded count.

The **no** form of the command excludes the out of profile octets forwarded count.

Default no out-profile-octets-forwarded-count

out-profile-packets-forwarded-count

Syntax [**no**] **out-profile-packets-forwarded-count**

Context config>log>acct-policy>cr>oc>i-count
config>log>acct-policy>cr>roc>i-count
config>log>acct-policy>cr>queue>i-count
config>log>acct-policy>cr>ref-queue>i-count

Description This command includes the out of profile packets forwarded count.

The **no** form of the command excludes the out of profile packets forwarded count.

Default no out-profile-packets-forwarded-count

uncoloured-octets-offered-count

Syntax [**no**] **uncoloured-packets-offered-count**

Context config>log>acct-policy>cr>queue>i-count
config>log>acct-policy>cr>ref-queue>i-count

Description This command includes the uncoloured octets offered in the count.

The **no** form of the command excludes the uncoloured octets offered in the count.

uncoloured-packets-offered-count

Syntax [**no**] **uncoloured-packets-offered-count**

Context config>log>acct-policy>cr>queue>i-count
config>log>acct-policy>cr>ref-queue>i-count

Description This command includes the uncoloured packets offered count.

The **no** form of the command excludes the uncoloured packets offered count.

ref-aa-specific-counter

Syntax **ref-aa-specific-counter any**
no ref-aa-specific-counter

Context config>log>acct-policy>cr

Description This command

The **no** form of the command

ref-override-counter

Syntax **ref-override-counter** *ref-override-counter-id*
ref-override-counter all
no ref-override-counter

Context config>log>acct-policy>cr

Description This command configures a reference override counter.

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

Default no ref-override-counter

ref-queue

Syntax **ref-queue** *queue-id*
ref-queue all
no ref-queue

Context config>log>acct-policy>cr

Description This command configures a reference queue.

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

Default no ref-queue

significant-change

Syntax **significant-change** *delta*
no significant-change

Context config>log>acct-policy>cr

Description This command configures the significant change required to generate the record.

Parameters *delta* — Specifies the delta change (significant change) that is required for the custom record to be written to the xml file.

Values 0 — 4294967295

BGP Peering Policy Commands

bgp-peering-policy

Syntax **bgp-peering-policy** *policy-name* [**create**]
no bgp-peering-policy *policy-name*

Context config>subscr-mgmt

Description This command configures the name of the BGP peering policy.

Parameters *policy-name* — Specifies the BGP peer policy name up to 32 characters in length.

advertise-inactive

Syntax [**no**] **advertise-inactive**

Context config>subscr-mgmt>bgp-prng-plcy

Description This command enables or disables the advertising of inactive BGP routers to other BGP peers.

By default, BGP only advertises BGP routes to other BGP peers if a given BGP route is chosen by the route table manager as the most preferred route within the system and is active in the forwarding plane. This command allows system administrators to advertise a BGP route even though it is not the most preferred route within the system for a given destination.

Default no advertise-inactive

aggregator-id-zero

Syntax [**no**] **aggregator-id-zero**

Context config>subscr-mgmt>bgp-prng-plcy

Description This command is used to set the router ID in the BGP aggregator path attribute to zero when BGP aggregates routes. This prevents different routers within an AS from creating aggregate routes that contain different AS paths.

When BGP is aggregating routes, it adds the aggregator path attribute to the BGP update messages. By default, BGP adds the AS number and router ID to the aggregator path attribute.

When this command is enabled, BGP adds the router ID to the aggregator path attribute. The **no** form of the command used at the global level reverts to default where BGP adds the AS number and router ID to the aggregator path attribute.

Default no aggregator-id-zero — BGP adds the AS number and router ID to the aggregator path attribute.

as-override

Syntax [**no**] **as-override**

Context config>subscr-mgmt>bgp-prng-plcy

Description This command replaces all instances of the peer's AS number with the local AS number in a BGP route's AS_PATH.

This command breaks BGP's loop detection mechanism. It should be used carefully.

Default as-override is not enabled by default.

auth-keychain

Syntax **auth-keychain** *name*
no auth-keychain

Context config>subscr-mgmt>bgp-prng-plcy

Description This command configures the BGP authentication key for all peers.

The keychain allows the rollover of authentication keys during the lifetime of a session.

Default no auth-keychain

Parameters *name* — Specifies the name of an existing keychain, up to 32 characters, to use for the specified TCP session or sessions.

authentication-key

Syntax **authentication-key** [*authentication-key* | *hash-key*] [**hash** | **hash2**]
no authentication-key

Context config>subscr-mgmt>bgp-prng-plcy

Description This command configures the BGP authentication key.

Authentication is performed between neighboring routers before setting up the BGP session by verifying the password. Authentication is performed using the MD-5 message-based digest. The authentication key can be any combination of letters or numbers from 1 to 16.

The no form of the command removes the authentication password from the configuration and effectively disables authentication.

Default Authentication is disabled and the authentication password is empty.

Parameters *authentication-key* — The authentication key. The key can be any combination of ASCII characters up to 255 characters in length (unencrypted). If spaces are used in the string, enclose the entire string in quotation marks (“ ”).

hash-key — The hash key. The key can be any combination of ASCII characters up to 342 characters in length (encrypted). If spaces are used in the string, enclose the entire string in quotation marks (“ ”).

This is useful when a user must configure the parameter, but, for security purposes, the actual unencrypted key value is not provided.

hash — Specifies the key is entered in an encrypted form. If the **hash** parameter is not used, the key is assumed to be in a non-encrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** parameter specified.

hash2 — Specifies the key is entered in a more complex encrypted form. If the **hash2** parameter is not used, the less encrypted **hash** form is assumed.

cluster

Syntax **cluster** *cluster-id*
no cluster

Context config>subscr-mgmt>bgp-prng-plcy

Description This command configures the cluster ID for a route reflector server.

Route reflectors are used to reduce the number of IBGP sessions required within an AS. Normally, all BGP speakers within an AS must have a BGP peering with every other BGP speaker in an AS. A route reflector and its clients form a cluster. Peers that are not part of the cluster are considered to be non-clients.

When a route reflector receives a route, first it must select the best path from all the paths received. If the route was received from a non-client peer, then the route reflector sends the route to all clients in the cluster. If the route came from a client peer, the route reflector sends the route to all non-client peers and to all client peers except the originator.

For redundancy, a cluster can have multiple route reflectors.

Confederations can also be used to remove the full IBGP mesh requirement within an AS.

The **no** form of the command deletes the cluster ID and effectively disables the Route Reflection for the given group.

Default no cluster — No cluster ID is defined.

Parameters *cluster-id* — The route reflector cluster ID is expressed in dot decimal notation.

Values Any 32 bit number in dot decimal notation. (0.0.0.1 — 255.255.255.255)

connect-retry

Syntax **connect-retry** *seconds*
no connect-retry

Context config>subscr-mgmt>bgp-prng-plcy

Description This command configures the BGP connect retry timer value in seconds.

When this timer expires, BGP tries to reconnect to the configured peer.

The **no** form of the command used at the global level reverts to the default value.

Default 120 seconds

Parameters *seconds* — The BGP Connect Retry timer value in seconds, expressed as a decimal integer.

Values 1 — 65535

damping

Syntax [**no**] **damping**

Context config>subscr-mgmt>bgp-prng-plcy

Description This command enables BGP route damping for learned routes which are defined within the route policy. Use damping to reduce the number of update messages sent between BGP peers and reduce the load on peers without affecting the route convergence time for stable routes. Damping parameters are set via route policy definition.

The **no** form of the command used at the global level disables route damping.

When damping is enabled and the route policy does not specify a damping profile, the default damping profile is used. This profile is always present and consists of the following parameters:

Half-life: 15 minutes
Max-suppress: 60 minutes
Suppress-threshold:3000
Reuse-threshold 750

Default no damping — Learned route damping is disabled.

disable-4byte-asn

Syntax [**no**] **disable-4byte-asn**

Context config>subscr-mgmt>bgp-prng-plcy

Description This command disables the use of 4-byte ASNs. It can be configured at all 3 level of the hierarchy so it can be specified down to the per peer basis.

If this command is enabled 4-byte ASN support should not be negotiated with the associated remote peer(s).

The **no** form of the command resets the behavior to the default which is to enable the use of 4-byte ASN.

disable-client-reflect

Syntax [**no**] **disable-client-reflect**

Context config>subscr-mgmt>bgp-prng-plcy

Description This command disables the reflection of routes by the route reflector to the group or neighbor. This only disables the reflection of routes from other client peers. Routes learned from non-client peers are still reflected to all clients.

The **no** form re-enables client reflection of routes.

Default no disable-client-reflect — Client routes are reflected to all client peers.

disable-communities

Syntax **disable-communities** [**standard**] [**extended**]
no disable-communities

Context config>subscr-mgmt>bgp-prng-plcy

Description This command configures BGP to disable sending communities.

Parameters **standard —** Specifies standard communities that existed before VPRNs or 2547.

extended — Specifies BGP communities used were expanded after the concept of 2547 was introduced, to include handling the VRF target.

disable-fast-external-failover

Syntax [**no**] **disable-fast-external-failover**

Context config>subscr-mgmt>bgp-prng-plcy

Description This command configures BGP fast external failover.

export

Syntax **export** *policy* [*policy...*]
no export

Context config>subscr-mgmt>bgp-prng-plcy

Description This command specifies the export policies to be used to control routes advertised to BGP neighbors.

When multiple policy names are specified, the policies are evaluated in the order they are specified. A maximum of five (5) policy names can be configured. The first policy that matches is applied.

Note that if a non-existent route policy is applied to a VPRN instance, the CLI generates a warning message. This message is only generated at an interactive CLI session and the route policy association is made. No warning message is generated when a non-existent route policy is applied to a VPRN instance in a configuration file or when SNMP is used.

The **no** form of this command removes all route policy names from the export list.

Default no export — BGP advertises routes from other BGP routes but does not advertise any routes from other protocols unless directed by an export policy.

Parameters *policy* — A route policy statement name.

hold-time

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

Context config>subscr-mgmt>bgp-prng-plcy

Description This command configures the BGP hold time, expressed in seconds.

The BGP hold time specifies the maximum time BGP waits between successive messages (either keepalive or update) from its peer, before closing the connection.

Even though the router OS implementation allows setting the keepalive time separately, the configured keepalive timer is overridden by the hold-time value under the following circumstances:

1. If the specified hold-time is less than the configured keepalive time, then the operational keepalive time is set to a third of the hold-time; the configured keepalive time is not changed.
2. If the hold-time is set to zero, then the operational value of the keepalive time is set to zero; the configured keepalive time is not changed. This means that the connection with the peer is up permanently and no keepalive packets are sent to the peer.

The **no** form of the command used at the global level reverts to the default value.

Default 90 seconds

Parameters *seconds* — The hold-time, in seconds, expressed as a decimal integer. A value of 0 indicates the connection to the peer is up permanently.

Values 0, 3 — 65535

import

Syntax **import** *policy* [*policy...*]
no import

Context config>subscr-mgmt>bgp-prng-plcy

Description This command specifies the import policies to be used to control routes advertised to BGP neighbors. Route policies are configured in the **config>router>policy-options** context. When multiple policy names are specified, the policies are evaluated in the order they are specified. A maximum of five (5) policy names can be specified. The first policy that matches is applied.

The **no** form of this command removes all route policy names from the import list.

Default no import — BGP accepts all routes from configured BGP neighbors. Import policies can be used to limit or modify the routes accepted and their corresponding parameters and metrics.

Parameters *policy* — A route policy statement name.

keepalive

Syntax **keepalive** *seconds*
no keepalive

Context config>subscr-mgmt>bgp-prng-plcy

Description This command configures the BGP keepalive timer. A keepalive message is sent every time this timer expires.

The keepalive value is generally one-third of the hold-time interval. Even though the OS implementation allows the keepalive value and the hold-time interval to be independently set, under the following circumstances, the configured keepalive value is overridden by the hold-time value:

If the specified keepalive value is greater than the configured hold-time, then the specified value is ignored, and the keepalive is set to one third of the current hold-time value.

If the specified hold-time interval is less than the configured keepalive value, then the keepalive value is reset to one third of the specified hold-time interval.

If the hold-time interval is set to zero, then the configured value of the keepalive value is ignored. This means that the connection with the peer is up permanently and no keepalive packets are sent to the peer.

The **no** form of the command used at the global level reverts to the default value..

Default 30 seconds

Parameters *seconds* — The keepalive timer in seconds, expressed as a decimal integer.

Values 0 — 21845

local-address

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

Context config>subscr-mgmt>bgp-prng-plcy

Description Configures the local IP address used by the group or neighbor when communicating with BGP peers.

Outgoing connections use the **local-address** as the source of the TCP connection when initiating connections with a peer.

When a local address is not specified, the 7750 SR OS uses the system IP address when communicating with IBGP peers and uses the interface address for directly connected EBGP peers. This command is used at the neighbor level to revert to the value defined under the group level.

The **no** form of the command removes the configured local-address for BGP.
The **no** form of the command used at the group level reverts to the value defined at the global level.
The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default **no local-address** — The router ID is used when communicating with IBGP peers and the interface address is used for directly connected EBGP peers.

ip-address — The local address expressed in dotted decimal notation. Allowed values are a valid routable IP address on the router, either an interface or system IP address.

local-as

Syntax **local-as** *as-number* [**private**]
no local-as

Context config>subscr-mgmt>bgp-prng-plcy

Description This command configures a BGP virtual autonomous system (AS) number.

In addition to the AS number configured for BGP in the `config>router>autonomous-system` context, a virtual (local) AS number is configured.The virtual AS number is added to the as-path message before the router's AS number makes the virtual AS the second AS in the as-path.

This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). Thus, by specifying this at each neighbor level, it is possible to have a separate as-number per EBGP session.

When a command is entered multiple times for the same AS, the last command entered is used in the configuration. The **private** attribute can be added or removed dynamically by reissuing the command.

Changing the local AS at the global level in an active BGP instance causes the BGP instance to restart with the new local AS number. Changing the local AS at the global level in an active BGP instance causes BGP to re-establish the peer relationships with all peers in the group with the new local AS number. Changing the local AS at the neighbor level in an active BGP instance causes BGP to re-establish the peer relationship with the new local AS number.

This is an optional command and can be used in the following circumstance:

Provider router P is moved from AS1 to AS2. The customer router that is connected to P, however, is configured to belong to AS1. To avoid reconfiguring the customer router, the **local-as** value on router P can be set to AS1. Thus, router P adds AS1 to the as-path message for routes it advertises to the customer router.

The **no** form of the command used at the global level will remove any virtual AS number configured. The **no** form of the command used at the group level reverts to the value defined at the global level. The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default no local-as

Parameters *as-number* — The virtual autonomous system number, expressed as a decimal integer.

Values 1 — 65535

private — Specifies the local-as is hidden in paths learned from the peering.

local-preference

Syntax **local-preference** *local-preference*
no local-preference

Context config>subscr-mgmt>bgp-prng-plcy

Description This command enables setting the BGP local-preference attribute in incoming routes if not specified and configures the default value for the attribute. This value is used if the BGP route arrives from a BGP peer without the **local-preference** integer set.

The specified value can be overridden by any value set via a route policy.

The **no** form of the command at the global level specifies that incoming routes with local-preference set are not overridden and routes arriving without local-preference set are interpreted as if the route had local-preference value of 100.

Default **no local-preference —** Does not override the local-preference value set in arriving routes and analyze routes without local preference with value of 100.

Parameters *local-preference* — The local preference value to be used as the override value, expressed as a decimal integer.

Values 0 — 4294967295

loop-detect

Syntax **loop-detect** {**drop-peer** | **discard-route** | **ignore-loop**| **off**}
no loop-detect

Context config>subscr-mgmt>bgp-prng-plcy

Description This command configures how the BGP peer session handles loop detection in the AS path.

Note that dynamic configuration changes of **loop-detect** are not recognized.

The **no** form of the command used at the global level reverts to default, which is **loop-detect ignore-loop**.

Default loop-detect ignore-loop

Parameters **drop-peer —** Sends a notification to the remote peer and drops the session.

discard-route — Discards routes received with loops in the AS path.

ignore-loop — Ignores routes with loops in the AS path but maintains peering.

off — Disables loop detection.

med-out

Syntax **med-out** {**number** | **igp-cost**}
no med-out

Context config>subscr-mgmt>bgp-prng-plcy

Description This command enables advertising the Multi-Exit Discriminator (MED) and assigns the value used for the path attribute for the MED advertised to BGP peers if the MED is not already set.

The specified value can be overridden by any value set via a route policy.

The **no** form of the command used at the global level reverts to default where the MED is not advertised.

no med-out

Parameters *number* — The MED path attribute value, expressed as a decimal integer.

Values 0 — 4294967295

igp-cost — The MED is set to the IGP cost of the given IP prefix.

min-as-origination

Syntax **min-as-origination** *seconds*
no min-as-origination

Context config>subscr-mgmt>bgp-prng-plcy

Description This command configures the minimum interval, in seconds, at which a path attribute, originated by the local router, can be advertised to a peer.

The **no** form of the command used at the global level reverts to default.

Default 15 seconds

Parameters *seconds* — The minimum path attribute advertising interval in seconds, expressed as a decimal integer.

Values 2 — 255

min-route-advertisement

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

Context config>subscr-mgmt>bgp-prng-plcy

Description This command configures the minimum interval, in seconds, at which a prefix can be advertised to a peer.

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

Default 30 seconds

Parameters *seconds* — The minimum route advertising interval, in seconds, expressed as a decimal integer.

Values 1— 255

multihop

Syntax **multihop** *ttl-value*
no multihop

Context config>subscr-mgmt>bgp-prng-plcy

Description This command configures the time to live (TTL) value entered in the IP header of packets sent to an EBGP peer multiple hops away.

This parameter is meaningful only when configuring EBGP peers. It is ignored if set for an IBGP peer.

The **no** form of the command is used to convey to the BGP instance that the EBGP peers are directly connected.
The **no** form of the command reverts to default values.

Default **1** — EBGP peers are directly connected.

64 — IBGP

Parameters *ttl-value* — The TTL value, expressed as a decimal integer.

Values 1 — 255

next-hop-self

Syntax [**no**] **next-hop-self**

Context config>subscr-mgmt>bgp-prng-plcy

Description This command configures the neighbor to always set the NEXTHOP path attribute to its own physical interface when advertising to a peer.

The no form of the command disables the command.

Default no next-hop-self

passive

Syntax [**no**] **passive**

Context config>subscr-mgmt>bgp-prng-plcy

Description This command enables the passive mode for the BGP neighbors.

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

Default no passive

peer-as

Syntax **peer-as** *as-number*
no peer-as

Context config>subscr-mgmt>bgp-prng-plcy

Description This command configures the autonomous system number for the remote peer. The peer AS number must be configured for each configured peer.

The no form of the command removes the *as-number* from the configuration.

Default No AS numbers are defined.

Parameters *as-number* — Specifies the AS number for the remote peer.

Values 1 — 4294967295

preference

Syntax [**no**] **preference** *preference*

Context config>subscr-mgmt>bgp-prng-plcy

Description This command configures the route preference for routes learned from the configured peer(s).

The lower the preference the higher the chance of the route being the active route. The OS assigns BGP routes highest default preference compared to routes that are direct, static or learned via MPLS or OSPF.

The **no** form of the command used at the global level reverts to default value.

Default 170

Parameters *preference* — The route preference, expressed as a decimal integer.

Values 1 — 255

prefix-limit

Syntax **prefix-limit** *limit* [**log-only**] [**threshold** *percent*]
no prefix-limit

Context config>subscr-mgmt>bgp-prng-plcy

Description This command configures the maximum number of routes BGP can learn from a peer.

When the number of routes reaches 90% of this limit, an SNMP trap is sent. When the limit is exceeded, the BGP peering is dropped and disabled.

The **no** form of the command removes the **prefix-limit**.

Parameters **log-only —** Enables the warning message to be sent at the specified threshold percentage, and also when the limit is exceeded. However, the BGP peering is not dropped.

percent — The threshold value (as a percentage) that triggers a warning message to be sent.

Default no prefix-limit

Parameters *limit* — The number of routes that can be learned from a peer, expressed as a decimal integer.

Values 1 — 4294967295

remove-private

Syntax [**no**] **remove-private**

Context config>subscr-mgmt>bgp-prng-plcy

Description This command allows private AS numbers to be removed from the AS path before advertising them to BGP peers.

The OS software recognizes the set of AS numbers that are defined by IANA as private. These are AS numbers in the range 64512 through 65535, inclusive.

The **no** form of the command used at the global level reverts to default value.

Default **no remove-private** — Private AS numbers will be included in the AS path attribute.

type

Syntax [**no**] **type** {**internal** | **external**}

Context config>subscr-mgmt>bgp-prng-plcy

Description This command designates the BGP peer as type internal or external.

The type of **internal** indicates the peer is an IBGP peer while the type of external indicates that the peer is an EBGP peer.

By default, the OS derives the type of neighbor based on the local AS specified. If the local AS specified is the same as the AS of the router, the peer is considered **internal**. If the local AS is different, then the peer is considered **external**.

The **no** form of the command used at the group level reverts to the default value.

Default **no type** — Type of neighbor is derived on the local AS specified.

Parameters **internal** — Configures the peer as internal.

external — Configures the peer as external.

ttl-security

Syntax **ttl-security** *min-ttl-value*
no ttl-security

Context config>subscr-mgmt>bgp-prng-plcy

Description Configure TTL security parameters for incoming packets.

Parameters *min-ttl-value* — Specify the minimum TTL value for an incoming BGP packet.

Values 1 — 255

RADIUS Route Download Commands

route-downloader

Syntax **route-downloader** *name* [**create**]
no route-downloader *name*

Context config>aaa

Description This command creates or enters the configuration of a route-downloader instance. The route-downloader is a process that uses radius access-request messages to a particular server. The server returns either an access-accept or access-deny message. Access-accept messages also contain the prefixes (in the form of static blackhole routes in various formats)

The **no** form of the command removes the name from the configuration. The object must be shutdown prior to deletion. No prefix is needed to delete an existing route-download object.

Default None. Only a single route-downloader object can be created.

Parameters *name* — Specifies the name of this RADIUS route downloader.

create — This keyword is mandatory while creating an instance of the route-download object.

base-user-name

Syntax **base-user-name** *user-name*
no base-user-name

Context config>aaa>route-downloader

Description This command sets the prefix for the user name that shall be used as access requests. The actual name used will be a concatenation of this string, the "-" (dash) character and a monotonically increasing integer.

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

Default The system's configured name (system-name).

Parameters *user-name* — Specifies the prefix of the username that is used in the RADIUS access requests. The username used in the RADIUS access requests is a concatenation of this string, the dash character and an increasing integer.

default-metric

Syntax **default-metric** *metric*
no default-metric

Context config>aaa>route-downloader

Description This command sets the default metric that routes imported by the RTM will acquire.

The no form of the command removes the metric

Default 2

Parameters *metric* — Specifies the default metric of the routes imported.

Values 0 — 254

default-tag

Syntax **default-tag** *tag*
no default-tag

Context config>aaa>route-downloader

Description This command sets the default tag that routes processed by the AAA route downloader will take. Note that any route received with a specific tag retains the specific tag. The tag value is passed to the Route Table Manager and is available as match condition on the export statement of other routing protocols.

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

Default 0

Parameters *tag* — Specifies the default tag of the routes imported.

Values 0 — 4294967295

download-interval

Syntax **download-interval** *minutes*
no download-interval

Context config>aaa>route-downloader

Description This command sets the time interval, in minutes, that the system waits for between two consecutive runs of the route-download process. The time is counted from the start-time of the run, thus, if an route-download process is still ongoing by the time the timer expires, the process will restart from count=1.

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

Default 720

Parameters *minutes* — Specifies the time interval, in minutes, between the start of the last route downloader run and the start of the next route downloader run.

Values 1 — 1440

max-routes

max-routes *routes*
no max-routes

Context config>aaa>route-downloader

Description This command determines the upper limits for total number of routes to be received and accepted by the system. The total number is inclusive of both IPv4 and IPv6 addresses and no differentiation is needed across protocols. It includes the sum of both. Once this limit is reached, the download process stops sending new access-requests until the next download-interval expires.

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

Default 200000

Parameters *routes* — Specifies the maximum number of the routes imported.

Values 1 — 200000

password

password *password* [**hash**|**hash2**]
no password

Context config>aaa>route-downloader

Description This command specifies the password that is used in the RADIUS access requests.It shall be specified as a string of up to 32 characters in length.

The **no** form of the command resets the password to its default of **ALU** and will be stored using hash/hash2 encryption.

Default ALU

Parameters *password* — Specifies a password string up to 32 characters in length.

hash — Specifies the key is entered in an encrypted form. If the **hash** parameter is not used, the key is assumed to be in a non-encrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** parameter specified.

hash2 — Specifies the key is entered in a more complex encrypted form. If the **hash2** parameter is not used, the less encrypted **hash** form is assumed.

radius-server-policy

Syntax **radius-server-policy** *policy-name*
no radius-server-policy

Context config>aaa>route-downloader

Description This command references an existing radius-server-policy (available under the **config>aaa** context). The server (or servers) referenced by the policy will be used as the targets for the access-request message.

The **no** form of the command removes the policy name from the route-downloader configuration.

Default none

Parameters *policy-name* — Specifies the RADIUS server policy.

retry-interval

Syntax **retry-interval min** *minimum* **max** *maximum*
no retry-interval

Context config>aaa>route-downloader

Description This command sets the duration, in minutes, of the retry interval. The retry interval is the interval meant for the system to retry sending an Access Request message after the previous one was unanswered (not with an access reject but rather just a RADIUS failure or ICMP port unreachable). This timer is actually an exponential backoff timer that starts at **min** and is capped at **max** minutes.

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

Default retry-interval min 10 max 20

Parameters **min** *minimum* **—** Specifies the duration, in minutes, of the retry interval. This duration grows exponentially after each sequential failure.

Values 1 — 1440

Default 10

max *maximum* **—** Specifies the maximum duration, in minutes, of the retry interval.

Values 1 — 1440

Default 20

Multi-Chassis Redundancy Commands

redundancy

Syntax **redundancy**

Context config

Description This command allows the user to perform redundancy operations.

Parameters **force-switchover** — Forces a switchover to the standby CPM card

Values **now** keyword - switch to standby CPM)

NOTE: Switching to the standby displays the following message.

```
WARNING: Configuration and/or Boot options may have changed since the last save.
Are you sure you want to switchover (y/n)?
```

synchronize — Synchronizes the secondary CPM.

Values **boot-env**|**config** : keywords

synchronize

Syntax **synchronize** {**boot-env** | **config**}

Context config>redundancy

Description This command performs a synchronization of the standby CPM's images and/or config files to the active CPM. Either the **boot-env** or **config** parameter must be specified.

In the **config>redundancy** context, this command performs an automatically triggered standby CPM synchronization.

When the standby CPM takes over operation following a failure or reset of the active CPM, it is important to ensure that the active and standby CPMs have identical operational parameters. This includes the saved configuration, CPM and IOM images.

The active CPM ensures that the active configuration is maintained on the standby CPM. However, to ensure smooth operation under all circumstances, runtime images and system initialization configurations must also be automatically synchronized between the active and standby CPM.

If synchronization fails, alarms and log messages that indicate the type of error that caused the failure of the synchronization operation are generated. When the error condition ceases to exist, the alarm is cleared.

Only files stored on the router are synchronized. If a configuration file or image is stored in a location other than on a local compact flash, the file is not synchronized (for example, storing a configuration file on an FTP server).

Default enabled

Parameters **boot-env —** Synchronizes all files required for the boot process (loader, BOF, images, and configuration files.

config — Synchronize only the primary, secondary, and tertiary configuration files.

Default config

multi-chassis

Syntax **multi-chassis**

Context config>redundancy

Description This command enables the context to configure multi-chassis parameters.

peer

Syntax **[no] peer *ip-address***

Context config>redundancy>multi-chassis

Description This command configures a multi-chassis redundancy peer.

Parameters *ip-address* — Specifies a peer IP address. Multicast address are not allowed.

authentication-key

Syntax **authentication-key** [*authentication-key* | *hash-key*] [**hash** | **hash2**]
no authentication-key

Context config>redundancy>multi-chassis>peer

Description This command configures the authentication key used between this node and the multi-chassis peer. The authentication key can be any combination of letters or numbers.

Parameters *authentication-key* — Specifies the authentication key. Allowed values are any string up to 20 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

hash-key — The hash key. The key can be any combination of ASCII characters up to 33 (hash1-key) or 55 (hash2-key) characters in length (encrypted). If spaces are used in the string, enclose the entire string in quotation marks (“ ”).

hash — Specifies the key is entered in an encrypted form. If the hash or hash2 parameter is not used, the key is assumed to be in a non-encrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the hash or hash2 parameter specified.

hash2 — Specifies the key is entered in a more complex encrypted form that involves more variables then the key value alone, this means that hash2 encrypted variable cannot be copied and pasted. If the hash or hash2 parameter is not used, the key is assumed to be in a non-encrypted, clear text

form. For security, all keys are stored in encrypted form in the configuration file with the hash or hash2 parameter specified.

mc-ipsec

Syntax **mc-ipsec**

Context config>redundancy>multi-chassis>peer

Description This command enters the configuration context of multi-chassis IPsec.

discovery-interval

Syntax **discovery-interval** *interval-1* [**boot** *interval-2*]
no discovery-interval

Context config>redundancy>multi-chassis>peer>mc-ipsec

Description This command specifies the time interval of tunnel-group stays in "Discovery" state. Interval-1 is used as discovery-interval when a new tunnel-group is added to multi-chassis redundancy (mp-ipsec); interval-2 is used as discovery-interval at system boot-up. It is optional and when it is not specified, interval-1 will be used.

Default 300

Parameters *interval-1/2* — Specifies the interval in seconds.

Values 1..1800 seconds

keep-alive-interval

Syntax **keep-alive-interval** *time-interval*
no keep-alive-interval

Context config>redundancy>multi-chassis>peer>mc-ipsec

Description This command specifies the time interval of the mastership election protocol sending the keep-alive packet.

Default 10

Parameters *time-interval* — Specifies the time interval in tenths of a second.

Values 5..500

hold-on-neighbor-failure

Syntax **hold-on-neighbor-failure** *multiplier*
no hold-on-neighbor-failure

Context config>redundancy>multi-chassis>peer>mc-ipsec

Description This command specifies the number of keep-alive failures before the peer is considered down.

Default 3

Parameters *multiplier* — Specifies the multiplier.

Values 2..25

bfd-enable

Syntax **bfd-enable service** *service-id* **interface** *interface-name* **dst-ip** *ip-address*
no bfd-enable

Context config>redundancy>multi-chassis>peer>mc-ipsec

Description This command enables tracking a central BFD session. If the BFD session goes down, then the system considers the peer down and changes the mc-ipsec status of the configured tunnel-group accordingly.

The BFD session uses the specified loopback interface (in the specified service) address as the source address and uses the specified dst-ip as the destination address. Other BFD parameters are configured with the "bfd" command on the specified interface.

Parameters *interface-name* — Specifies the name of the loopback interface.

service-id — Specifies the ID of the service.

dst-id — Specifies the destination address of the BFD packet.

tunnel-group

Syntax **tunnel-group** *group-id* [**create**]
no tunnel-group *group-id*

Context config>redundancy>multi-chassis>peer>mc-ipsec

Description This command enables multi-chassis redundancy for the specified tunnel-group or enters an already configured tunnel-group context. The configured tunnel-group could failover independently.

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

Values 1..16

create — Enables multi-chassis redundancy for the specified tunnel-group.

peer-group

Syntax **peer-group** *group-id*
no peer-group

Context config>redundancy>multi-chassis>peer>mc-ipsec>tunnel-group

Description This command specifies the corresponding tunnel-group ID on the peer node. The peer tunnel-group ID does not necessarily equal the local tunnel-group ID.

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

Values 1..16

priority

Syntax **priority** *priority*
no priority

Context config>redundancy>multi-chassis>peer>mc-ipsec>tunnel-group

Description This command specifies the local priority of the tunnel-group. This is used to elect the master (higher number is the master). If priorities are the same, then the peer with the more active ISA becomes the master. If the priority and the number of active ISAs are the same, then the peer with the higher IP address is the master.

Parameters *priority* — Specifies the priority of the tunnel-group.

Values 0..255

preempt

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

Context config>redundancy>multi-chassis>peer>mc-ipsec>tunnel-group

Description This command enables the preempt behavior of local node.

mc-lag

Syntax [**no**] **mc-lag**

Context config>redundancy>multi-chassis>peer>mc-lag

Description This command enables the context to configure multi-chassis LAG operations and related parameters.

The **no** form of this command administratively disables multi-chassis LAG. MC-LAG can only be issued only when mc-lag is shutdown.

hold-on-neighbor-failure

Syntax **hold-on-neighbor-failure** *multiplier*
no hold-on-neighbor-failure

Context config>redundancy>multi-chassis>peer>mc-lag

Description This command specifies the interval that the standby node will wait for packets from the active node before assuming a redundant-neighbor node failure. This delay in switch-over operation is required to accommodate different factors influencing node failure detection rate, such as IGP convergence, or HA switch-over times and to prevent the standby node to take action prematurely.

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

Default 3

Parameters *multiplier* — The time interval that the standby node will wait for packets from the active node before assuming a redundant-neighbor node failure.

Values 2 — 25

keep-alive-interval

Syntax **keep-alive-interval** *interval*
no keep-alive-interval

Context config>redundancy>multi-chassis>peer>mc-lag

Description This command sets the interval at which keep-alive messages are exchanged between two systems participating in MC-LAG. These keep-alive messages are used to determine remote-node failure and the interval is set in deci-seconds.

The no form of this command sets the interval to default value

Default 1s (10 hundreds of milliseconds means interval value of 10)

Parameters *interval* — The time interval expressed in deci-seconds

Values 5 — 500

lag

Syntax **lag** *lag-id* **lacp-key** *admin-key* **system-id** *system-id* [**remote-lag** *lag-id*] **system-priority** *system-priority*
no lag *lag-id*

Context config>redundancy>multi-chassis>peer>mc-lag

Description This command defines a LAG which is forming a redundant-pair for MC-LAG with a LAG configured on the given peer. The same LAG group can be defined only in the scope of 1 peer.

The same **lacp-key**, **system-id**, and **system-priority** must be configured on both nodes of the redundant pair in order to MC-LAG to become operational. In order MC-LAG to become operational, all

parameters (**lacp-key**, **system-id**, **system-priority**) must be configured the same on both nodes of the same redundant pair.

The partner system (the system connected to all links forming MC-LAG) will consider all ports using the same **lacp-key**, **system-id**, **system-priority** as the part of the same LAG. In order to achieve this in MC operation, both redundant-pair nodes have to be configured with the same values. In case of the mismatch, MC-LAG is kept in oper-down status.

Default none

Parameters *lag-id* — The LAG identifier, expressed as a decimal integer. Specifying the *lag-id* allows the mismatch between lag-id on redundant-pair. If no **lag-id** is specified it is assumed that neighbor system uses the same *lag-id* as a part of the given MC-LAG. If no matching MC-LAG group can be found between neighbor systems, the individual LAGs will operate as usual (no MC-LAG operation is established.).

Values 1 — 200

lacp-key ***admin-key*** **—** Specifies a 16 bit key that needs to be configured in the same manner on both sides of the MC-LAG in order for the MC-LAG to come up.

Values 1 — 65535

system-id ***system-id*** **—** Specifies a 6 byte value expressed in the same notation as MAC address

Values xx:xx:xx:xx:xx:xx - xx [00..FF]

remote-lag ***lag-id*** **—** Specifies the LAG ID on the remote system.

Values 1 — 200

system-priority ***system-priority*** **—** Specifies the system priority to be used in the context of the MC-LAG. The partner system will consider all ports using the same **lacp-key**, **system-id**, and **system-priority** as part of the same LAG.

Values 1 — 65535

source-address

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

Context config>redundancy>multi-chassis>peer

Description This command specifies the source address used to communicate with the multi-chassis peer.

Parameters *ip-address* — Specifies the source address used to communicate with the multi-chassis peer.

sync

Syntax [**no**] **sync**

Context config>redundancy>multi-chassis>peer

Description This command enables the context to configure synchronization parameters.

igmp

Syntax [**no**] **igmp**

Context config>redundancy>multi-chassis>peer>sync

Description This command specifies whether IGMP protocol information should be synchronized with the multi-chassis peer.

Default no igmp

igmp-snooping

Syntax [**no**] **igmp-snooping**

Context config>redundancy>multi-chassis>peer>sync

Description This command specifies whether IGMP snooping information should be synchronized with the multi-chassis peer.

Default no igmp-snooping

local-dhcp-server

Syntax [**no**] **local-dhcp-server**

Context config>redundancy>multi-chassis>peer>sync

Description This command synchronizes DHCP server information.

mc-ring

Syntax **[no] mc-ring**

Context config>redundancy>multi-chassis>peer>sync

Description This command synchronizes mc-ring information.

mld-snooping

Syntax **[no] mld-snooping**

Context config>redundancy>multi-chassis>peer>sync

Description This command synchronizes MLD snooping information.

port

Syntax **port** [*port-id* | *lag-id*] [**sync-tag** *sync-tag*] [**create**]
no port [*port-id* | *lag-id*]

Context config>redundancy>multi-chassis>peer>sync

Description This command specifies the port to be synchronized with the multi-chassis peer and a synchronization tag to be used while synchronizing this port with the multi-chassis peer.

Parameters *port-id* — Specifies the port to be synchronized with the multi-chassis peer.

lag-id — Specifies the LAG ID to be synchronized with the multi-chassis peer.

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

range

Syntax **range** *encap-range* **sync-tag** *sync-tag*
no range *encap-range*

Context config>redundancy>multi-chassis>peer>sync>port

Description This command configures a range of encapsulation values.

Parameters *encap-range* — Specifies a range of encapsulation values on a port to be synchronized with a multi-chassis peer.

Values		
	Dot1Q	*start-vlan-end-vlan*
	QinQ	Q1.*start-vlan*-Q1.*end-vlan*

sync-tag *sync-tag* — specifies a synchronization tag up to 32 characters in length to be used while synchronizing this encapsulation value range with the multi-chassis peer.

srrp

Syntax [**no**] **srrp**

Context config>redundancy>multi-chassis>peer>sync

Description This command specifies whether subscriber routed redundancy protocol (SRRP) information should be synchronized with the multi-chassis peer.

Default no srrp

sub-host-trk

Syntax [**no**] **sub-host-trk**

Context config>redundancy>multi-chassis>peer>sync

Description This command synchronizes subscriber host tracking information.

sub-mgmt

Syntax **sub-mgmt** [**ipoe** | **pppoe**]
[**no**] **sub-mgmt**

Context config>redundancy>multi-chassis>peer>sync

Description This command will enable synchronization of subscriber states between chassis. Synchronization will be enabled per protocol type (IPoE or PPPoE).

The keywords (**ipoe**, **pppoe**) must match on both nodes. If not, subscriber synchronization will fail.

For example if one node is configured with:

configure>multi-chassis>peer>sync>sub-mgmt ipoe

but the other node is configured with:

configure>multi-chassis>peer>sync>sub-mgmt ipoe pppoe

synchronization will fail even for ipoe application.

Default no sub-mgmt

Parameters **ipoe —** ipoe subscribers will be synchronized

pppoe — pppoe subscribers will be synchronized

tunnel-group

Syntax **tunnel-group** *tunnel-group-id* **sync-tag** *tag-name* [**create**]
no tunnel-group

Context config>redundancy>multi-chassis>peer>sync

Description This command enables multi-chassis synchronization of IPsec states of a specified tunnel-group with its peer. Sync-tag is used to match corresponding tunnel-groups on both peers. IPsec states will be synchronized between tunnel-groups with the same sync-tag.

Parameters *tunnel-group-id* — Specifies the ID of the tunnel-group

tag-name — Specifies the name of sync-tag.

ipsec

Syntax [**no**] **ipsec**

Context config>redundancy>multi-chassis>peer>sync

Description This command enables multi-chassis synchronization of IPsec states on system level.

mc-ring

Syntax **mc-ring**

Context config>redundancy>multi-chassis>peer

Description This command enables the context to configure the multi-chassis ring parameters.

Default mc-ring

ring

Syntax [**no**] **ring** *sync-tag* [**create**]

Context config>redundancy>multi-chassis>peer>mcr

Description This command configures a multi-chassis ring.

The **no** form of the command removes the sync-tag from the configuration.

Default none

l3-ring

Syntax [**no**] **l3-ring name** [**create**]

Context config>redundancy>multi-chassis>peer>mcr

Parameters This command configures a layer 3 multi-chassis ring.

in-band-control-path

Syntax **in-band-control-path**

Context config>redundancy>multi-chassis>peer>mcr>ring
config>redundancy>multi-chassis>peer>mc>l3-ring

Description This command enables the context to configure control path parameters.

Default none

debounce

Syntax [**no**] **debounce**

Context config>redundancy>multi-chassis>peer>mcr>ring>in-band-control-path
config>redundancy>multi-chassis>peer>mc>l3-ring>in-band-control-path

Description This command enables the inband control path debouncing. The **no** form of the command disables inband control path debouncing.

dst-ip

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

Context config>redundancy>multi-chassis>peer>mcr>ring>in-band-control-path
config>redundancy>multi-chassis>peer>mc>l3-ring>in-band-control-path

Description This command specifies the destination IP address used in the inband control connection.

If the destination IP address is not configured, the ring cannot become operational.

Default none

Parameters *ip-address* — The destination IP address.

interface

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

Context config>redundancy>multi-chassis>peer>mcr>ring>in-band-control-path
config>redundancy>multi-chassis>peer>mc>l3-ring>in-band-control-path

Description This command specifies the name of the IP interface used for the inband control connection.

If an interface name is not configured, the ring cannot become operational.

Parameters *ip-int-name* — Specifies an interface name up to 32 characters in length.

max-debounce-time

Syntax **max-debounce-time** *max-debounce-time*
no max-debounce-time

Context config>redundancy>multi-chassis>peer>mcr>ring>in-band-control-path
config>redundancy>multi-chassis>peer>mc>l3-ring>in-band-control-path

Description This command configures the inband control path maximum debounce time.

Parameters *max-debounce-time* — Specifies the maximum debounce time on the transition of the operational state of the inband control connection.

Values 5 — 200 seconds

service-id

Syntax **service-id** *service-id*
no service-id

Context config>redundancy>multi-chassis>peer>mcr>ring>in-band-control-path
config>redundancy>multi-chassis>peer>mc>l3-ring>in-band-control-path

Description This command configures the service ID of the SAP used for the Ring-Node Connectivity Verification of this ring node.

Parameters *service-id* — [Specifies an existing service ID or service name.

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

path-b

Syntax [**no**] **path-b**

Context config>redundancy>multi-chassis>peer>mcr>ring

Description This command specifies the set of upper-VLAN IDs associated with the SAPs that belong to path B with respect to load-sharing. All other SAPs belong to path A.

Default If not specified, the default is an empty set.

range

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

Context config>redundancy>multi-chassis>peer>mcr>ring>path-b
config>redundancy>multi-chassis>peer>mcr>ring>path-excl

Description This command specifies the set of VLAN IDs associated with the SAPs that are controlled by the remote peer. It is a bitmap that associates bit i with VLAN ID i, with i in [0..4094]. Setting the value to the empty string is equivalent to setting it to 512 zeroes.

ring-node

Syntax [**no**] **ring-node** *ring-node-name*

Context config>redundancy>mc>peer>mcr>ring

Description This command specifies the unique name of a multi-chassis ring access node.

path-excl

Syntax [**no**] **path-excl**

Context config>redundancy>multi-chassis>peer>mcr>ring

Description This command specifies the set of upper-VLAN IDs associated with the SAPs that are to be excluded from control by the multi-chassis ring.

Default If not specified, the default is an empty set.

connectivity-verify

Syntax **connectivity-verify**

Context config>redundancy>multi-chassis>peer>mcr>ring
config>redundancy>multi-chassis>peer>mc>l3-ring

Description This command configures the node connectivity check.

interval

Syntax **interval** *interval*

Context config>redundancy>multi-chassis>peer>mcr>ring>>connectivity-verify
config>redundancy>multi-chassis>peer>mc>l3-ring>connectivity-verify

Description This command specifies the polling interval of the ring-node connectivity verification of this ring node.

Parameters *interval* — Specifies the polling interval of the ring-node connectivity verification of this ring node.

Values 1 — 6000

service-id

Syntax **service-id** *service-id*
no service-id

Context config>redundancy>mc>peer>mcr>ring-node>connect-verify
config>redundancy>multi-chassis>peer>mc>l3-ring>connectivity-verify

Description This command specifies the service ID of the SAP used for ring-node connectivity verification of this ring node.

Parameters *service-id* — Specifies the service ID or service name.

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

src-ip

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

Context config>redundancy>mc>peer>mcr>ring-node>connect-verify
config>redundancy>multi-chassis>peer>mc>l3-ring>connectivity-verify

Description This command specifies the source IP address used in ring-node connectivity verification of this ring node.

Parameters *ip-address* — Specifies the source IP address used in ring-node connectivity verification of this ring node.

src-mac

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

Context config>redundancy>mc>peer>mcr>ring-node>connect-verify
config>redundancy>multi-chassis>peer>mc>l3-ring>connectivity-verify

Description This command specifies the source MAC address used for the Ring-Node Connectivity Verification of this ring node.

If all zeros are specified, then the MAC address of the system management processor (CPM) is used.

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

vlan

Syntax **vlan** [**0**..**4094**]

Context config>redundancy>mc>peer>mcr>ring-node>connect-verify
config>redundancy>mc>peer>mcr>l3ring>node>cv

Description This command specifies the VLAN tag of the SAP used for ring-node connectivity verification of this ring node. It is only meaningful if the value of is not zero.

srrp-instance

Syntax **[no] srrp-instance *srrp-id***

Context config>redundancy>multi-chassis>peer>mc>l3-ring

Description This command configures an SRRP instance for layer 3 ring.

Parameters *srrp-id* — Specifies the SRRP ID of this SRRP instance.

Values 1 — 4294967295

Log Commands

accounting-policy

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

Context config>log

Description This command creates an access or network accounting policy. An accounting policy defines the accounting records that are created.

Parameters *policy-id* — The policy ID that uniquely identifies the accounting policy, expressed as a decimal integer.

Values 1 — 99

collection-interval

Syntax **collection-interval** *minutes*
no collection-interval

Context config>log>accounting-policy

Description This command configures the accounting collection interval.

Parameters *minutes* — indicates how frequently, in minutes, the statistics are collected and written to their destination.

Values 5 — 120

default

Syntax [**no**] **default**

Context config>log>accounting-policy

This command contains the description for the accounting policy to be used with all SAPs that do not have an accounting policy.

record

Syntax **record** *record-name*
no record

Context config>log>accounting-policy

Description This command configures the accounting record type to be forwarded to the configured accounting file.

Parameters *record-name* — service-ingress-octets | service-egress-octets | service-ingress-packets | service-egress-packets | network-ingress-octets | network-egress-octets | network-ingress-packets | network-egress-packets | compact-service-ingress-octets | combined-service-ingress | combined-network-ing-egr-octets | combined-service-ing-egr-octets complete-service-ingress-egress | complete-sdp-ingress-egress |combined-sdp-ingress-egress | complete-subscriber-ingress-egress | aa-protocol | aa-application | aa-app-group | aa-subscriber-protocol | aa-subscriber-application | aa-subscriber-app-group

shutdown

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

Context config>log>accounting-policy

Description This command administratively enables or disables the accounting policy.

to

Syntax **to file** *log-file-id*

Context config>log>accounting-policy

Description This command configures the log destination.

Parameters *log-file-id* — Specifies the log ID.

Values 1 — 99

event-control

Syntax **event-control** *application-id* [*event-name* | *event-number*] [**generate** [*severity-level*] [**throttle**]
event-control *application-id* [*event-name* | *event-number*] **suppress**
no event-control *application* [*event-name* | *event-number*]

Context config>log

Description This command is used to specify that a particular event or all events associated with an application is either generated or suppressed.

Events are generated by an application and contain an event number and description explaining the cause of the event. Each event has a default designation which directs it to be generated or suppressed.

Events are generated with a default severity level that can be modified by using the *severity-level* option.

Events that are suppressed by default are typically used for debugging purposes. Events are suppressed at the time the application requests the event's generation. No event log entry is generated regardless of the destination. While this feature can save processor resources, there may be a negative effect on the ability to troubleshoot problems if the logging entries are squelched. In reverse, indiscriminate application may cause excessive overhead.

The **no** form of the command reverts the parameters to the default setting for events for the application or a specific event within the application. The severity, generate, and suppress options will also be reset to the initial values.

Default Each event has a default suppress or generate state. To display a list of all events and the current configuration use the **event-control** command.

Parameters *application-id* — The application whose events are affected by this event control filter.

Default None, this parameter must be explicitly specified.

Values A valid application name. To display a list of valid application names, use the **applications** command. Valid applications are:

aps, atm, bgp, cflowd, chassis, debug, dhcp, efm_oam, filter, gsmp, igmp, igmp_snooping, ip, isis, lag, ldp, logger, mc_redundancy, mirror, mpls, ntp, oam, ospf, pim, port, ppp, rip, route_policy, rsvp, security, snmp, stp, svcmgr, system, user, vrrp, vrtr

event-name | *event-number* — To generate, suppress, or revert to default for a single event, enter the specific number or event short name. If no event number or name is specified, the command applies to all events in the application. To display a list of all event short names use the **event-control** command.

Default none

Values A valid event name or event number.

generate serverity-level — This keyword indicates that the specified events will generate logging events.

Default generate

severity-name — An ASCII string representing the severity level to associate with the specified generated events

Default The system assigned severity name

Values One of: cleared, indeterminate, critical, major, minor, warning.

suppress — This keyword indicates that the specified events will not be logged. If the **suppress** keyword is not specified then the events are generated by default.

Default generate

throttle — Specifies whether or not events of this type will be throttled.
By default, event throttling is off for each specific event type. It must be explicitly enabled for each event type where throttling is desired. This makes backwards compatability easier to manage.

Subscriber Identification Policy Commands

sub-ident-policy

Syntax [**no**] **sub-ident-policy** *sub-ident-policy-name*

Context config>subscr-mgmt

Description This command configures a subscriber identification policy. Each subscriber identification policy can have a default subscriber profile defined. The subscriber identification policy default subscriber profile overrides the system default and the subscriber SAP default subscriber profiles. Defining a subscriber identification policy default subscriber profile is optional.

The subscriber identification policy default subscriber profile cannot be defined with the subscriber profile name default.

Defining a subscriber profile as a subscriber identification policy default subscriber profile will cause all active subscribers currently associated with a subscriber SAP using the policy and associated with a subscriber policy through the system default or subscriber SAP default subscriber profiles to be reassigned to the subscriber policy defined as default on the subscriber identification policy.

Attempting to delete a subscriber profile that is currently defined as a default for a subscriber identification policy will fail.

When attempting to remove a subscriber identification policy default subscriber profile definition, the system will evaluate each active subscriber on all subscriber SAPs the subscriber identification policy is currently associated with that are using the default definition to determine whether the active subscriber can be either reassigned to a subscriber SAP default or the system default subscriber profile. If all active subscribers cannot be reassigned, the removal attempt will fail.

Parameters *sub-ident-policy-name* — Specifies the name of the subscriber identification policy.

app-profile-map

Syntax **app-profile-map**

Context config>subscr-mgmt>sub-ident-pol

Description This command enables the context to configure an application profile mapping.

entry

Syntax **entry key** *app-profile-string* **app-profile** *app-profile-name*
no entry key *app-profile-string*

Context config>subscr-mgmt>sub-ident-pol>app-profile-map

Description This command configures an application profile string.

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

Parameters *app-profile-string* — Specifies the application profile string.

app-profile-name — Specifies the application profile name.

use-direct-map-as-default

Syntax [**no**] **use-direct-map-as-default**

Context config>subscr-mgmt>sub-ident-pol>app-profile-map
config>subscr-mgmt>sub-ident-pol>sla-profile-map

Description This command enables direct mapping of application profile as default. With this flag, a script returned string will be used as the named profile. If the named profiled cannot be found, the default profile will be used.

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

Default no use-direct-map-as-default

primary

Syntax **primary**

Context config>subscr-mgmt>sub-ident-pol

Description This command configures a primary identification script.

script-url

Syntax **script-url** *dhcp-script-url*

Context config>subscr-mgmt>sub-ident-pol>primary
config>subscr-mgmt>sub-ident-pol>secondary
config>subscr-mgmt>sub-ident-pol>tertiary

Description This command specifies the URL of the identification scripts.

Parameters *dhcp-primary-script-url* — Specifies the URL of the primary identification script.

dhcp-secondary-script-url — Specifies the URL of the secondary identification script.

dhcp-tertiary-script-url — Specifies the URL of the tertiary identification script.

secondary

Syntax **secondary**

Context config>subscr-mgmt>sub-ident-pol

Description This command configures a secondary identification script.

sla-profile-map

Syntax **sla-profile-map**

Context config>subscr-mgmt>sub-ident-pol

Description This command configures an SLA profile mapping.

sub-profile-map

Syntax **sla-profile-map**

Context config>subscr-mgmt>sub-ident-pol

Description This command configures a subscriber profile mapping.

entry

Syntax **entry key** *sla-profile-string* **sla-profile** *sla-profile-name*
no entry key *sla-profile-string*

Context config>subscr-mgmt>sub-ident-pol>sla-profile-map

Description This command configures an SLA profile string. Each subscriber identification string can be provisioned into a subscriber mapping table providing an explicit mapping of the string to a specific subscriber profile. This allows certain subscribers to be directly mapped to the appropriate subscriber profile in the event that the default mappings are not desired for the subscriber.

An explicit mapping of a subscriber identification string to a subscriber profile cannot be defined with the subscriber profile name default. It is possible for the subscriber identification string to be entered in the mapping table without a defined subscriber profile which can result in the explicitly defined subscriber to be associated with the subscriber profile named default.

Explicitly mapping a subscriber identification string to a subscriber profile will cause an existing active subscriber associated with the string to be reassigned to the newly mapped subscriber profile. An explicit mapping overrides all default subscriber profile definitions.

Attempting to delete a subscriber profile that is currently defined as in an explicit subscriber identification string mapping will fail.

The system will fail the removal attempt of an explicit subscriber identification string mapping to a subscriber profile definition when an active subscriber is using the mapping and cannot be reassigned to a defined default non-provisioned subscriber profile.

Parameters *sla-profile-string* — Identifies the SLA profile string.

Values 16 characters maximum

sla-profile-name — Identifies the SLA profile name.

Values 32 characters maximum

entry

Syntax **entry key** *sub-profile-string* **sub-profile** *sub-profile-name*
no entry key *sub-profile-string*

Context config>subscr-mgmt>sub-ident-pol>sub-profile-map

Description This command configures a subscriber profile string.

Parameters *sub-profile-string* — Specifies the subscriber profile string.

Values 16 characters maximum

sub-profile-name — Specifies the subscriber profile name.

Values 32 characters maximum

tertiary

Syntax **tertiary**

Context config>subscr-mgmt>sub-ident-pol

Description This command configures a tertiary identification script.

Auto-Generated Subscriber Identification Key Commands

auto-sub-id-key

Syntax **auto-sub-id-key**

Context config>subscr-mgmt

ipoe-sub-id-key

Syntax **ipoe-sub-id-key** *sub-id-key* [*sub-id-key*...(up to 4 max)]
no ipoe-sub-id-key

Context config>subscr-mgmt>>auto-sub-id-key

Description This command enables certain fields to become the base for auto-generation of the default sub-id name. The sub-id name will be auto generated if there is not a more specific method available. Such more specific methods would be a default sub-id name as a sap-id, a preconfigured static string or explicit mappings based on RADIUS/LUDB returned strings.

In case that a more specific sub-id name generation method is not available AND the auto-id keyword is defined under the def-sub-id hierarchy, the sub-id name will be generated by concatenating fields defined in this command separated by a "|" character.

The maximum sub-id name length is 32 characters while the concatenation of subscriber identification fields can easily exceed 32 characters. Subscriber host instantiation will fail in case that the sub-id name is based on subscriber identification fields whose concatenated length exceeds 32 characters. Failing the host creation rather than truncating sub-id name on a 32 character boundary will prevent collision of sub-ids (subscriber name duplication).

In case that a more specific sub-id name generation method is not available AND the auto-id keyword is NOT defined under the def-sub-id hierarchy, the sub-id name will be a random 10 character encoded string based on the fields defined under this command.

There is only one set of identification fields allowed per host type (IPoE or PPP) per chassis.

Parameters *sub-id-key* — Specifies the auto-generated sub-id keys for IPoE hosts.

Values **mac** — The MAC address can be combined with other subscriber host identification fields (circuit-id, remote-id, session-id or sap-id) to form a sub-id name in a user readable format or as a random 10 character encoded value.

In case that the mac address is used as a concatenation field in the sub-id name, then its format becomes a string xx:xx:xx:xx:xx:xx with the length 17B.

The MAC address as the subscriber host identification field is not applicable to PPPoA hosts or static hosts.

circuit-id — The circuit-id can be combined with other subscriber host identification fields (mac, remote-id, session-id or sap-id) to form a sub-id name in a user readable format or as a random 10 character encoded value.

In case that the circuit-id is used as a concatenation field in the sub-id name, then its format becomes access-node-id eth slot/port:[vlan-id] or access-node-id atm slot/port:vpi.vci with a variable length.

Note that if circuit-id contains any non printable ASCI characters, the entire circuit-id string will be formatted in hex in the sub-id name output. Otherwise all characters in circuit-id will be converted to ASCII. ASCII printable characters contain bytes in range 0x20..0x7E.

The circuit-id as the subscriber identification field is not applicable to PPPoA hosts, ARP hosts or static hosts.

remote-id — The remote-id can be combined with other subscriber host identification fields (mac, circuit-id, session-id or sap-id) to form a sub-id name in a user readable format or as a random 10 character encoded value.

In case that the remote-id is used as a concatenation field in the sub-id name, then its format becomes a remote-id string with a variable length.

Note that if remote-id contains any non printable ASCI characters, the entire remote-id string will be formatted in hex in the sub-id name output. Otherwise all characters in remote-id will be converted to ASCII. ASCII printable characters contain bytes in range 0x20..0x7E.

The remote-id as the subscriber identification field is not applicable to PPPoA hosts, ARP hosts or static hosts.

sap-id — The sap-id can be combined with other subscriber host identification fields (mac, circuit-id, remote-id, or session-id) to form a sub-id name in a user readable format or as a random 10 character encoded value.

In case that the circuit-id is used as a concatenation field in the sub-id name, then its format becomes : slot/mda:[outer-vlan].[inner-vlan] with a variable length.

The sap-id as the subscriber identification field is applicable to all hosts types with exception of static hosts.

Default ipoe-sub-id-key mac sap-id

ppp-sub-id-key

Syntax **ppp-sub-id-key** *sub-id-key* [*sub-id-key*...(up to 5 max)]
no ppp-sub-id-key

Context config>subscr-mgmt>>auto-sub-id-key

Description This command enable certain fields to become the base for auto-generation of default sub-id name. The sub-id name will be auto-generated if there is not a more specific method available. Examples of these specific methods would be a default sub-id name as a sap-id, a preconfigured static string or explicit mappings based on RADIUS/LUDB returned strings.

In case that a more specific sub-id name generation method is not available and the **auto-id** keyword is defined under the def-sub-id hierarchy, the sub-id name will be generated by concatenating fields defined in this command separated by a "|" character.

The maximum sub-id name length is 32 characters while the concatenation of subscriber identification fields can easily exceed 32 characters. The subscriber host instantiation will fail if the sub-id name is based on subscriber identification fields whose concatenated length exceeds 32 characters. Failing the host creation rather than truncating sub-id name on a 32 character boundary will prevent collision of sub-ids (subscriber name duplication).

In case that a more specific sub-id name generation method is not available and the **auto-id** keyword is not defined under the def-sub-id hierarchy, the sub-id name will be a random 10 character encoded string based on the fields defined under this command.

There is only one set of identification fields allowed per host type (IPoE or PPP) per chassis.

Parameters *sub-id-key* — Specifies the auto-generated sub-id keys for PPP hosts.

Values **mac** — The MAC address can be combined with other subscriber host identification fields (circuit-id, remote-id, session-id or sap-id) to form a sub-id name in a user readable format or as a random 10 character encoded value.

In case that the mac address is used as a concatenation field in the sub-id name, then its format becomes a string xx:xx:xx:xx:xx:xx with the length 17B.

The MAC address as the subscriber host identification field is not applicable to PPPoA hosts or static hosts.

circuit-id — The circuit-id can be combined with other subscriber host identification fields (mac, remote-id, session-id or sap-id) to form a sub-id name in a user readable format or as a random 10 character encoded value.

In case that the circuit-id is used as a concatenation field in the sub-id name, then its format becomes access-node-id eth slot/port:[vlan-id] or access-node-id atm slot/port:vpi.vci with a variable length.

Note that if circuit-id contains any non printable ASCI characters, the entire circuit-id string will be formatted in hex in the sub-id name output. Otherwise all characters in circuit-id will be converted to ASCII. ASCII printable characters contain bytes in range 0x20..0x7E.

.The circuit-id as the subscriber identification field is not applicable to PPPoA hosts, ARP hosts or static hosts.

remote-id — The remote-id can be combined with other subscriber host identification fields (mac, circuit-id, session-id or sap-id) to form a sub-id name in a user readable format or as a random 10 character encoded value.

In case that the remote-id is used as a concatenation field in the sub-id name, then its format becomes a remote-id string with a variable length.

Please note that if remote-id contains any non printable ASCI characters, the entire remote-id string will be formatted in hex in the sub-id name output. Otherwise all characters in remote-id will be converted to ASCII. ASCII printable characters contain bytes in range 0x20..0x7E.

The remote-id as the subscriber identification field is not applicable to PPPoA hosts, ARP hosts or static hosts.

sap-id — The sap-id can be combined with other subscriber host identification fields (mac, circuit-id, remote-id, or session-id) to form a sub-id name in a user readable format or as a random 10 character encoded value.

In case that the circuit-id is used as a concatenation field in the sub-id name, then its format becomes : slot/mda:[outer-vlan].[inner-vlan] with a variable length.

The sap-id as the subscriber identification field is applicable to all hosts types with exception of static hosts.

session-id — The session-id can be combined with other subscriber host identification fields (mac, circuit-id, remote-id, or sap-id) to form a sub-id name in a user readable format or as a random 10 character encoded value.

In case that the circuit-id is used as a concatenation field in the sub-id name, then its format becomes a decimal number with variable length.

The session-id as the subscriber identification field is applicable only to PPPoE/ PPPoEoA type hosts.

Default ppp-sub-id-key mac sap-id session-id

Subscriber Profile Commands

sub-profile

Syntax [**no**] sub-profile ***subscriber-profile-name***

Context config>subscr-mgmt

Description This command enables the context to configure a subscriber profile. A subscriber profile is a template used to define the aggregate QoS for all hosts within a subscriber context. This is done through the definition of the egress and ingress scheduler policies that govern the aggregate SLA for subscribers using the subscriber profile. Subscriber profiles also allow for specific SLA profile definitions when the default definitions from the subscriber identification policy must be overridden.

Subscribers are either explicitly mapped to a subscriber profile template or are dynamically associated by one of various non-provisioned subscriber profile definitions.

A subscriber host can be associated with a subscriber profile in the following ways, listed from lowest to highest precedence:

1. The subscriber profile named default.
2. The subscriber profile defined as the subscriber SAP default.
3. The subscriber profile found by the subscriber identification policy sub-profile-map.
4. The subscriber profile found by the subscriber identification policy explicit map.

In the event that no defaults are defined and the subscriber identification string is not explicitly provisioned to map to a subscriber profile, either the static subscriber host creation will fail or the dynamic subscriber host DHCP ACK will be discarded.

Default Subscriber profile:

When a subscriber profile is created with the *subscriber-profile-name* default, it will be used when no other subscriber profile is associated with the subscriber host by the system. Creating a subscriber profile with the *subscriber-profile-name* default is optional. If a default subscriber profile is not created, all subscriber hosts subscriber identification strings must match either a non-provisioned default or be provisioned as an explicit match to a subscriber profile.

The default profile has no effect on existing active subscriber on the system as they exist due to higher precedence mappings.

Attempting to delete any subscriber profile (including the profile named default) while in use by existing active subscribers will fail.

Parameters *subscriber-profile-name* — Specify the name of the subscriber profile.

Values 32 characters maximum, default

accounting-policy

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

Context config>subscr-mgmt>sub-prof

Description This command specifies the policy to use to collect accounting statistics on this subscriber profile.

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

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

Default no accounting policy

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

Values 1 — 99

collect-stats

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

Context config>subscr-mgmt>sub-prof

Description When enabled, the agent collects non-RADIUS accounting statistics.

When the **no collect-stats** command is issued the statistics are still accumulated by the IOM cards. However, the CPU will not obtain the results and write them to the billing file. If a subsequent **collect-stats** command is issued then the counters written to the billing file include all the traffic while the **no collect-stats** command was in effect.

Default collect-stats

agg-rate-limit

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

Context config>subscr-mgmt>sub-prof>egress

Description This command define a subscriber aggregate limit when the subscriber profile is directly associated with an egress port based scheduler instead of a scheduler policy. The optional queue-frame-based-accounting keyword allows the subscriber queues to operate in the frame based accounting mode.

Once egress frame based accounting is enabled on the subscriber profile, all queues associated with the subscriber (created through the sla-profile associated with each subscriber host) will have their rate and CIR values interpreted as frame based values. When shaping, the queues will include the 12 byte Inter-Frame Gap (IFG) and 8 byte preamble for each packet scheduled out the queue. The profiling CIR threshold will also include the 20 byte frame encapsulation overhead. Statistics associated with the queue do not include the frame encapsulation overhead.

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

Enabling queue-frame-based-accounting will not cause statistics for queues associated with the subscriber to be cleared.

The **no** form of the command removes both an egress aggregate rate limit and egress frame based accounting for all subscribers associated with the sub-profile. If a subscriber's accounting mode is changed, the subscriber's queue statistics are cleared.

Parameters {**max** | *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 egress-agg-rate-limit command.

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

kilobits-per-second — The kilobits-per-second parameter defines an actual egress aggregate rate to which all queues associated with the sub-profile will be limited. The limit will be managed per subscriber associated with the sub-profile. The value must be defined as an integer and is representative of increments of 1000 bits per second.

Values 1 to 40000000

Default max

queue-frame-based-accounting — The optional queue-frame-based-accounting keyword enables frame based accounting on all queues associated with the subscriber profile. If frame based accounting is required when a subscriber aggregate limit is not necessary, the max keyword should precede the queue-frame-based-accounting keyword. If frame based accounting must be disabled, execute egress-agg-rate-limit without the queue-frame-based-accounting keyword present.

Default Frame based accounting is disabled by default

queue-frame-based-accounting — Specifies whether to use frame-based accounting when evaluating the aggregation rate limit for the egress queues for this SAP.

avg-frame-size

Syntax **avg-frame-size** *bytes*
no avg-frame-size

Context config>subscriber-managemet>sub-profile>egress

Description This command specifies the average frame size used in the calculation of the fixed and variable encapsulation offset when the command encap-offset is enabled in the egress context of a subscriber profile.

If the user does not explicitly configure a value for the avg-frame-size parameter, then it will also be assumed the offset is zero.

The **no** form of the command removes the avg-frame-size parameter from the subscriber profile.

Default 0

Parameters *bytes* — specifies the average frame size value to be used in the adjustment of the subscriber aggregate rate to account for the per packet variable expansion of the last mile for the specific session used by the subscriber host.

Values 64 — 4096

encap-offset

Syntax **encap-offset** [**type** *type*]
no encap-offset

Context config>subscriber-managemet>sub-profile>egress

Description This command enables the adjustment of the queue and subscriber aggregate rate based on the last mile Ethernet or ATM encapsulation.

In R9.0, the data path computes the adjusted frame size real-time for each serviced packet from a queue by adding the actual packet size to the fixed offset provided by CPM for this queue and variable AAL5 padding.

When this command is enabled, the fixed packet offset is derived from the encapsulation type value signaled in the Access-loop-encapsulation sub-TLV in the Vendor-Specific PPPoE Tags or DHCP Relay Options as per RFC 4679. If the user specifies an encapsulation type with the command, this value is used as the default value for all hosts of this subscriber until a host session signaled a valid value. The signaled value is applied to this host only and the remaining hosts of this subscriber continue to use the user entered default type value if configured, or no offset is applied. Note that however, hosts of the same subscriber using the same SLA profile and which are on the same SAP will share the same instance of FC queues. In this case, the last valid encapsulation value signaled by a host of that same instance of the SAP egress QoS policy will override any previous signaled or configured value.

If the user manually applied a constant byte offset to each packet serviced by the queue by configuring the packet-byte-offset, it will have no effect on the net offset computed for the packet. This net offset is stored in the subscriber host table.

The procedures for handling signaling changes or configuration changes affecting the subscriber profile are as follows:

1. The avg-frame-size parameter in the subscriber profile is ignored.
2. If the user specifies an encapsulation type with the command, this value is used as the default value for all hosts of this subscriber until a host session signaled a valid value. The signaled value is applied to this host and other hosts of the same subscriber sharing the same SLA profile and which are on the same SAP. The remaining hosts of this subscriber continue to use the user entered default type value if configured, or no offset is applied.
3. If the user enables/disables the encap-offset option, or changes the parameter value of the encap-offset option, CPM immediately triggers a re-evaluation of subscribers hosts using the corresponding subscriber profile and an update the IOM with the new fixed offset value.
4. If a subscriber has a static host or an ARP host, the subscriber host continues to use the user-configured default encapsulation type value or the last valid encapsulation value signaled in the

PPPoE tags or DHCP relay options by other hosts of the same subscriber which use the same SLA profile instance. If none was signaled or configured, then no rate adjustment is applied.

When the encap-offset option is configured in the subscriber profile, the subscriber host queue rates, that is, CLI and operational PIR and CIR as well as queue bucket updates, the queue statistics, that is, forwarded, dropped, and HQoS offered counters use the last-mile frame-over-the-wire format. The scheduler policy CLI and operational rates also use LM-FoW format. The port scheduler max-rate and the priority level rates and weights, if a Weighted Scheduler Group is used, are always entered in CLI and interpreted as local port frame-over-the-wire rates. The same is true for an agg-rate-limit applied to a vport. Finally the subscriber agg-rate-limit is entered in CLI as last-mile frame-over-the-wire rate. The system maintains a running average frame expansion ratio for each queue to convert queue rates between these two formats.

Parameters **type** *type* — The name of the default encapsulation used for all host queues of a subscriber in the absence of a valid value signaled in the PPPoE tags.

Values pppoa-llc|pppoa-null|pppoeoa-llc|pppoeoa-llc-fcs|pppoeoa-llc-tagged|pppoeoa-llc-tagged-fcs|pppoeoa-null|pppoeoa-null-fcs|pppoeoa-null-tagged|pppoeoa-null-tagged-fcs|ipoa-llc|ipoa-null|ipoeoa-llc|ipoeoa-llc-fcs|ipoeoa-llc-tagged|ipoeoa-llc-tagged-fcs|ipoeoa-null|ipoeoa-null-fcs|ipoeoa-null-tagged|ipoeoa-null-tagged-fcs|pppoe|pppoe-tagged|ipoe|ipoe-tagged

scheduler

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

Context config>subscr-mgmt>sub-prof>egr-sched-pol
config>subscr-mgmt>sub-prof>ingr-sched-pol

Description This command provides a way to override parameters of the existing scheduler associated with the egress or ingress scheduler policy. A scheduler defines bandwidth controls that limit each child (other schedulers and queues) associated with the scheduler. Scheduler objects are created within the hierarchical tiers of the policy. It is assumed that each scheduler created will have queues or other schedulers defined as child associations. The scheduler can be a child (take bandwidth from a scheduler in a higher tier, except for schedulers created in tier 1).

Parameters **scheduler** *scheduler-policy-name* — Specify an existing scheduler policy name.

pir-rate — Defines the administrative PIR rate, in kilobits, for the queue. When the **rate** command is executed, a valid PIR setting must be explicitly defined. When the **rate** command has not been executed, the default PIR of **max** is assumed.
Fractional values are not allowed and must be given as a positive integer.

The actual PIR rate is dependent on the queues **adaptation-rule** parameters and the actual hardware where the queue is provisioned.

Values 1 — 100000000, max

Default max

cir-rate — The **cir** parameter overrides the default administrative CIR used by the queue. When the **rate** command is executed, a CIR setting is optional. When the **rate** command has not been

executed or the **cir** parameter is not explicitly specified, the default CIR (0) is assumed. Fractional values are not allowed and must be given as a positive integer.

Values 0 — 100000000, **max**

Default 0

host-lockout-policy

Syntax **host-lockout-policy** *policy-name* [**create**]
no host-lockout-policy *policy-name*

Context config>subscriber-mgmt

Description This command creates a host lockout policy. The policy contains set of host lockout configuration parameters. It is applied to SAP or MSAPs (by a MSAP-policy). Any change does not impact existing locked-out hosts, but only new incoming hosts that enter lockout.

The **no** form of the command removes the policy name from the configuration. The policy must not be associated with any entity.

Default none

Parameters *policy-name* — Specifies an existing host lockout policy to associate with the SAP.

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

host-lockout-policy

Syntax **host-lockout-policy** *policy-name*
no host-lockout-policy

Context config>service>ies>interface>sap
config>service>ies>subscriber-interface>sap
config>service>vpls>sap
config>service>vprn>interface>sap
config>service>vprn>subscriber-interface>sap

Description This command selects an existing host lockout policy. The **host-lockout-policy** *policy-name* is created in the **config>subscriber-mgmt** context.

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

Default none

Parameters *policy-name* — Specifies an existing host lockout policy to associate with the SAP.

lockout-time

Syntax **lockout-time** [**min** *seconds*] [**max** *seconds*]
no lockout-time

Context config>subscriber-mgmt>host-lockout-policy

Description This command configures the time for which a client stays in the lockout state during which authentication and ESM host creation is suppressed. The range for the min and max lockout times is 1 second to 86400 seconds. The min time defaults to 10 seconds, and max time defaults to 3600 seconds.

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

Parameters **min** *seconds* — specifies the minimum lockout-time for this host lockout policy.

Values 1 — 86400

Default 10 seconds

max *seconds* — specifies the maximum lockout-time for this host lockout policy.

Values 1 — 86400

Default 3600 seconds

lockout-reset-time

Syntax **lockout-reset-time** *seconds*
no lockout-reset-time

Context config>subscriber-mgmt>host-lockout-policy

Description This command configures the time that needs to elapse from the point a client enters lockout to when the client's lockout time can be reset to the configured minimum value. The range is 1 sec

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

Parameters *seconds* — Specifies the lockout reset time in seconds.

Values 1 — 86400

Default 60 seconds

max-lockout-hosts

Syntax **max-lockout-hosts** *hosts*
no max-lockout-hosts

Context config>subscriber-mgmt>host-lockout-policy

Description When a client enters lockout, authentication and ESM host creation is suppressed. A lightweight context maintains the lockout state and the timeouts for the client in lockout. This command allows the number of lockout contexts to be configured per SAP. If the number of existing contexts reaches the

configured count, incoming hosts that fail authentication or creation are not subject to lockout, and are retired as normal.

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

Parameters *hosts* — Specifies the maximum number of lockout hosts.

Values 1 — 1000

Default 100

host-tracking-policy

Syntax **host-tracking-policy** *policy-name* [**create**]
no host-tracking-policy *policy-name*

Context config>subscr-mgmt
config>subscr-mgmt>sub-prof

Description This command configures a host tracking policy. IGMP host tracking is an option in the subscriber profile that allows the factoring in of a subscriber's (multicast) video traffic by reducing the unicast operational egress aggregate rate or the rate of the scheduler specified in the ANCP policy to account for a subscriber's multicast traffic. If no ANCP policy is defined, the egress aggregate rate configured in the subscriber profile is reduced. If an ANCP policy is defined, the "rate-modify" parameter in the policy specifies whether the egress aggregate rate or the rate of the egress policer specified in the policy is to be reduced to account for the subscriber's multicast traffic.

Default disabled

egress-rate-modify

Syntax **egress-rate-modify agg-rate-limit**
egress-rate-modify scheduler *scheduler-name*
no egress-rate-modify

Context config>subscr-mgmt>trk-plcy

Description This command specifies the egress-rate modification that is to be applied.

agg-rate-limit — Specifies the egress rate limit.

scheduler *scheduler-name* **—** Specifies the scheduler name to use.

scheduler-policy

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

Context config>subscriber-mgmt>sub-profile>egress
config>subscriber-mgmt>sub-profile>ingress

Description This command specifies a scheduler policy to associate to the subscriber profile. Scheduler policies are configured in the **configure>qos>scheduler>policy** context. Each scheduler policy is divided up into groups of schedulers based on the tier each scheduler is created under. A tier is used to give structure to the schedulers within a policy and define rules for parent scheduler associations. The policy defines the hierarchy and operating parameters for virtual schedulers.

Parameters *scheduler-policy-name* — Specify an existing scheduler policy name.

policer-control-policy

Syntax **policer-control-policy** *policy-name* [**create**]
no policer-control-policy

Context config>subscr-mgmt>sub-prof>ingress
config>subscr-mgmt>sub-prof>egress

Description This command is used to create, delete, or modify policer control policies. The **policer-control-policy** controls the aggregate bandwidth available to a set of child policers. Once created, the policy can be applied to ingress or egress SAPs. The policy can also be applied to the ingress or egress context of a sub-profile.

Default no policer-control-policy

Parameters *policy-name* — Each policer-control-policy must be created with a unique policy name. The name must given as *policy-name* must adhere to the system policy ASCII naming requirements. If the defined policy-name already exists, the system will enter that policy's context for editing purposes. If policy-name does not exist, the system will attempt to create a policy with the specified name. Creating a policy may require use of the create parameter when the system is configured for explicit object creation mode.

Default None

create — The **create** keyword is required when a new policy is being created and the system is configured for explicit object creation mode.

max-rate

Syntax **max-rate** {*kilobits-per-second* | **max**}
no max-rate

Context config>subscr-mgmt>sub-prof>ingress>policer-control-policy
config>subscr-mgmt>sub-prof>egress>policer-control-policy

Description The **max-rate** command defines the parent policer's PIR leaky bucket's decrement rate. A parent policer is created for each time the policer-control-policy is applied to either a SAP or subscriber instance. Packets that are not discarded by the child policers associated with the SAP or subscriber instance are evaluated against the parent policer's PIR leaky bucket.

For each packet, the bucket is first decremented by the correct amount based on the decrement rate to derive the current bucket depth. The current depth is then compared to one of two discard thresholds associated with the packet. The first discard threshold (discard-unfair) is applied if the FIR (Fair Information Rate) leaky bucket in the packet's child policer is in the confirming state. The second discard threshold (discard-all) is applied if the child policer's FIR leaky bucket is in the exceed state. Only one of the two thresholds is applied per packet. If the current depth of the parent policer PIR bucket is less than the threshold value, the parent PIR bucket is in the conform state for that particular packet. If the depth is equal to or greater than the applied threshold, the bucket is in the violate state for the packet.

If the result is "conform," the bucket depth is increased by the size of the packet (plus or minus the per-packet-offset setting in the child policer) and the packet is not discarded by the parent policer. If the result is "violate," the bucket depth is not increased and the packet is discarded by the parent policer. When the parent policer discards a packet, any bucket depth increases (PIR, CIR and FIR) in the parent policer caused by the packet are canceled. This prevents packets that are discarded by the parent policer from consuming the child policers PIR, CIR and FIR bandwidth.

The **policer-control-policy root max-rate** setting may be overridden on each SAP or sub-profile where the policy is applied.

Default max

Parameters *kilobits-per-second* — Defining a kilobits-per-second value is mutually exclusive with the max parameter. The kilobits-per-second value must be defined as an integer that represents the number of kilobytes that the parent policer will be decremented per second. The actual decrement is performed per packet based on the time that has elapsed since the last packet associated with the parent policer.

Values Integer 0 – 20,000,000

max — The **max** parameter is mutually exclusive with defining a **kilobits-per-second** value. When max is specified, the parent policer does not enforce a maximum rate on the aggregate throughput of the child policers. This is the default setting when the **policer-control-policy** is first created and is the value that the parent policer returns to when no max-rate is executed. In order for the parent policer to be effective, a kilobits-per-second value should be specified.

no max-rate — The **no max-rate** command returns the policer-control-policy's parent policer maximum rate to max.

priority-mbs-thresholds

Syntax **priority-mbs-thresholds**

Context config>subscr-mgmt>sub-prof>ingress>policer-control-policy
config>subscr-mgmt>sub-prof>egress>policer-control-policy

Description The **priority-mbs-thresholds** command contains the root arbiter parent policer's **min-thresh-separation** command and each priority level's **mbs-contribution** command that is used to internally derive each priority level's shared-portion and fair-portion values. The system uses each priority

level's shared-portion and fair-portion value to calculate each priority level's discard-unfair and discard-all MBS thresholds that enforce priority sensitive rate-based discards within the root arbiter's parent policer.

The **priority-mbs-thresholds** CLI node always exists and does not need to be created.

Default None.

min-thresh-separation

Syntax **min-thresh-separation** *size* [**bytes** | **kilobytes**]
no min-thresh-separation

Context config>subscr-mgmt>sub-prof>ingress>policer-control-policy>priority-mbs-thresholds
config>subscr-mgmt>sub-prof>egress>policer-control-policy>priority-mbs-thresholds

Description The **min-thresh-separation** command defines the minimum required separation between each in-use discard threshold maintained for each parent policer context associated with the policer-control-policy. The min-thresh-separation value may be overridden on each SAP or sub-profile to which the policy is applied.

The system uses the default or specified min-thresh-separation value in order to determine the minimum separation required between each of the of the parent policer discard thresholds. The system enforces the minimum separation based on the following behavior in two ways. The first is determining the size of the shared-portion for each priority level (when the **mbs-contribution** command's optional fixed keyword is not specified):

- When a parent policer instance's priority level has less than two child policers associated, the shared-portion for the level will be zero.
- When a parent policer instance's priority level has two or more child policers associated, the shared-portion for the level will be equal to the current value of **min-thresh-separation**.

The second function the system uses the **min-thresh-separation** value for is determining the value per priority level for the fair-portion:

- When a parent policer instance's priority level has no child policers associated, the fair-portion for the level will be zero.
- When a parent policer instance's priority level has one child policer associated, the fair-portion will be equal to the maximum of the min-thresh-separation value and the priority level's mbs-contribution value.
- When a parent policer instance's priority level has two or more child policers associated, the fair-portion will be equal to the maximum of the following:

  –**min-thresh-separation** value

  –The priority level's **mbs-contribution** value less **min-thresh-separation** value

When the **mbs-contribution** command's optional fixed keyword is defined for a priority level within the policy, the system will treat the defined **mbs-contribution** value as an explicit definition of the priority level's MBS. While the system will continue to track child policer associations with the parent policer priority levels, the association counters will have no effect. Instead the following rules will be used to determine a fixed priority level's shared-portion and fair-portion:

- If a fixed priority level's **mbs-contribution** value is set to zero, both the shared-portion and fair-portion will be set to zero
- If the **mbs-contribution** value is not set to zero:

  –The shared-portion will be set to the current **min-thresh-separation** value

  –The fair-portion will be set to the maximum of the following:

  min-thresh-separation value

  mbs-contribution value less **min-thresh-separation value**

Each time the **min-thresh-separation** value is modified, the thresholds for all instances of the parent policer created through association with this **policer-control-policy** are reevaluated.

Determining the Correct Value for the Minimum Threshold Separation Value

The minimum value for **min-thresh-separation** should be set equal to the maximum size packet that will be handled by the parent policer. This ensures that when a lower priority packet is incrementing the bucket, the size of the increment will not cause the bucket's depth to equal or exceed a higher priority threshold. It also ensures that an unfair packet within a priority level cannot cause the PIR bucket to increment to the discard-all threshold within the priority.

When evaluating maximum packet size, each child policer's per-packet-offset setting should be taken into consideration. If the maximum size packet is 1518 bytes and a per-packet-offset parameter is configured to add 20 bytes per packet, min-thresh-separation should be set to 1538 due to the fact that the parent policer will increment its PIR bucket using the extra 20 bytes.

In most circumstances, a value larger than the maximum packet size is not necessary. Management of priority level aggregate burst tolerance is intended to be implemented using the priority level **mbs-contribution** command. Setting a value larger than the maximum packet size will not adversely affect the policer performance, but it may increase the aggregate burst tolerance for each priority level.

NOTE: One thing to note is that a priority level's shared-portion of the parent policer's PIR bucket depth is only necessary to provide some separation between a lower priority's discard-all threshold and this priority's discard-unfair threshold. It is expected that the burst tolerance for the unfair packets is relatively minimal since the child policers feeding the parent policer priority level all have some amount of fair burst before entering into an FIR exceed or unfair state. The fair burst amount for a priority level is defined using the mbs-contribution command.

The **no** form of this command returns the policy's **min-thresh-separation** value to the default value.

Default **no min-thresh-separation**

Parameters *size* [**bytes** | **kilobytes**] — The size parameter is required when executing the **min-thresh-separation** command. It is expressed as an integer and specifies the shared portion in bytes or kilobytes that is selected by the trailing bytes or kilobytes keywords. If both bytes and kilobytes are missing, kilobytes is the assumed value. Setting this value has no effect on parent policer instances where the **min-thresh-separation** value has been overridden.

Values 0 – 4194304

Default none

[**bytes** | **kilobytes**] — The **bytes** keyword is optional and is mutually exclusive with the **kilobytes** keyword. When specified, size is interpreted as specifying the size of **min-thresh-separation** in bytes.

The **kilobytes** keyword is optional and is mutually exclusive with the **bytes** keyword. When specified, size is interpreted as specifying the size of **min-thresh-separation** in kilobytes.

Values **bytes** or **kilobytes**

Default **kilobytes**

priority

Syntax **priority** *level*

Context config>subscr-mgmt>sub-prof>ingress>policer-control-policy>priority-mbs-thresholds
config>subscr-mgmt>sub-prof>egress>policer-control-policy>priority-mbs-thresholds

Description The **priority** level command contains the **mbs-contribution** configuration command for a given strict priority level. Eight levels are supported numbered 1 through 8 with 8 being the highest strict priority.

Each of the eight priority CLI nodes always exists and do not need to be created. While parameters exist for each priority level, the parameters are only applied when the priority level within a parent policer instance is currently supporting child policers.

Default None.

mbs-contribution

Syntax **mbs-contribution** *size* [**bytes | kilobytes**] [**fixed**]
no mbs-contribution

Context config>subscr-mgmt>sub-prof>ingress>policer-control-policy>priority-mbs-thresholds>priority
config>subscr-mgmt>sub-prof>egress>policer-control-policy>priority-mbs-thresholds>priority

Description The **mbs-contribution** command is used to configure the policy-based burst tolerance for a parent policer instance created when the policy is applied to a SAP or subscriber context. The system uses the parent policer's **min-thresh-separation** value, the priority level's **mbs-contribution** value and the number of child policers currently attached to the priority level to derive the priority level's shared-portion and fair-portion of burst tolerance within the local priority level. The shared-portion and fair-portions for each priority level are then used by the system to calculate each priority level's discard-unfair threshold and discard-all threshold.

The value for a priority level's **mbs-contribution** within the policer-control-policy may be overridden on the SAP or subscriber sub-profile where the policy is applied in order to allow fine tuning of the discard-unfair and discard-all thresholds relevant to the needs of the local child policers on the object.

Accumulative Nature of Burst Tolerance for a Parent Policer Priority Level

When defining **mbs-contribution**, the specified size may only be a portion of the burst tolerance associated with the priority level. The packets associated with the priority level share the burst tolerance of lower within the parent policer. As the parent policer PIR bucket depth increases during con-

gestion, the lower priority packets eventually experience discard based on each priority's discard-unfair and discard-all thresholds. Assuming congestion continues once all the lower priority packets have been prevented from consuming bucket depth, the burst tolerance for the priority level will be consumed by its own packets and any packets associated with higher priorities.

The Effect of Fair and Unfair Child Policer Traffic at a Parent Policer Priority Level

The system continually monitors the offered rate of each child policer on each parent policer priority level and detects when the policer is in a congested state (the aggregate offered load is greater than the decrement rate defined on the parent policer). As previously stated, the result of congestion is that the parent policer's bucket depth will increase until it eventually hovers around either a discard-unfair or discard-all threshold belonging to one of the priority levels. This threshold is the point where enough packets are being discarded that the increment rate and decrement rate begin to even out. If only a single child policer is associated to the priority level, the discard-unfair threshold is not used since fairness is only applicable when multiple child policers are competing at the same priority level.

When multiple child policers are sharing the congested priority level, the system uses the offered rates and the parenting parameters of each child to determine the fair rate per child when the parent policer is unable to meet the bandwidth needs of each child. The fair rate represents the amount of bandwidth that each child at the priority level should receive relative to the other children at the same level according to the policer control policy instance managing the child policers. This fair rate is applied as the decrement rate for each child's FIR bucket. Changing a child's FIR rate does not modify the amount of packets forwarded by the parent policer for the child's priority level. It simply modifies the forwarded ratio between the children on that priority level. Since each child FIR bucket has some level of burst tolerance before marking its packets as unfair, the current parent policer bucket depth may at times rise above the discard-unfair threshold. The mbs-contribution value provides a means to define how much separation is provided between the priority level's discard-unfair and discard-all threshold to allow the parent policer to absorb some amount of FIR burst before reaching the priority's discard-all threshold.

This level of fair aggregate burst tolerance is based on the decrement rate of the parent policer's PIR bucket while the individual fair bursts making up the aggregate are based on each child's FIR decrement rate. The aggregate fair rate of the priority level is managed by the system with consideration of the current rate of traffic in higher priority levels. In essence, the system ensures that for each iteration of the child FIR rate calculation, the sum of the child FIR decrement rates plus the sum of the higher priority traffic increment rates equals the parent policers decrement rate. This means that dynamic amounts of higher priority traffic can be ignored when sizing a lower priority's fair aggregate burst tolerance. Consider the following:

- The parent policer decrement rate is set to 20 Mbps (max-rate 20,000).
- A priority level's fair burst size is set to 30 Kbytes (mbs-contribution 30 kilobytes).
- Higher priority traffic is currently taking 12 Mbps.
- The priority level has three child policers attached.
- Each child's PIR MBS is set to 10 Kbytes, which makes each child's FIR MBS 10 Kbytes.
- The children want 10 Mbps, but only 8 Mbps is available,
- Based on weights, the children's FIR rates are set as follows:

	FIR Rate	FIR MBS
Child 1	4 Mbps	10 Kbytes
Child 2	3 Mbps	10 Kbytes
Child 3	1 Mbps	10 Kbytes

The 12 Mbps of the higher priority traffic and the 8 Mbps of fair traffic equal the 20 Mbps decrement rate of the parent policer.

It is clear that the higher priority traffic is consuming 12 Mbps of the parent policer's decrement rate, leaving 8 Mbps of decrement rate for the lower priority's fair traffic.

- The burst tolerance of child 1 is based on 10 Kbytes above 4 Mbps,
- The burst tolerance of child 2 is based on 10 Kbytes above 3 Mbps,
- The burst tolerance of child 3 is based on 10 Kbytes above 1 Mbps.

If all three children burst simultaneously (unlikely), they will consume 30 Kbytes above 8 Mbps. This is the same as the remaining decrement rate after the higher priority traffic.

Parent Policer Total Burst Tolerance and Downstream Buffering

The highest in-use priority level's discard-all threshold is the total burst tolerance of the parent policer. In some cases the parent policer represents downstream bandwidth capacity and the max-rate of the parent policer is set to prevent overrunning the downstream bandwidth. The burst tolerance of the parent policer defines how much more traffic may be sent beyond the downstream scheduling capacity. In the worst case scenario, when the downstream buffering is insufficient to handle the total possible burst from the parent policer, downstream discards based on lack of buffering may occur. However, in all likelihood, this is not the case.

In most cases, lower priority traffic in the policer will be responsible for the greater part of congestion above the parent policer rate. Since this traffic is discarded with a lower threshold, this lowers the effective burst tolerance even while the highest priority traffic is present.

Configuring a Priority Level's MBS Contribution Value

In the most conservative case, a priority level's **mbs-contribution** value may be set to be greater than the sum of child policer's mbs and one max-size-frame per child policer. This ensures that even in the absolute worst case where all the lower priority levels are simultaneously bursting to the maximum capacity of each child, enough burst tolerance for the priority's children will exist if they also burst to their maximum capacity.

Since simply adding up all the child policer's PIR MBS values may result in large overall burst tolerances that are not ever likely to be needed, you should consider some level of burst oversubscription when configuring the **mbs-contribution** value for each priority level. The amount of oversubscription should be determined based on the needs of each priority level.

Using the Fixed Keyword to Create Deterministic Parent Policer Discard Thresholds

In the default behavior, the system ignores the **mbs-contribution** values for a priority level on a subscriber or SAP parent policer when a child policer is not currently associated with the level. This prevents additional burst tolerance from being added to higher priority traffic within the parent policer.

This does cause fluctuations in the defined threshold values when child policers are added or removed from a parent policer instance. If this behavior is undesirable, the fixed keyword may be used which causes the **mbs-contribution** value to always be included in the calculation of parent policer's dis-

card thresholds. The defined **mbs-contribution** value may be overridden on a subscriber sla-profile or on a SAP instance, but the fixed nature of the contribution cannot be overridden.

If the defined **mbs-contribution** value for the priority level is zero, the priority level will have no effect on the parent policer's defined discard thresholds. A packet associated with the priority level will use the next lower priority level's discard-unfair and discard-all thresholds.

Parameters *size* [**bytes** | **kilobytes**] — The size parameter is required when executing the **mbs-contribution** command. It is expressed as an integer and specifies the priority's specific portion amount of accumulative MBS for the priority level in bytes or kilobytes which is selected by the trailing **bytes** or **kilobytes** keywords. If both **bytes** and **kilobytes** are missing, **kilobytes** is assumed. Setting this value has no effect on parent policer instances where the priority level's **mbs-contribution** value has been overridden.

Values 0 — 4194304

Default none

bytes | **kilobytes**: — The **bytes** keyword is optional and is mutually exclusive with the **kilobytes** keyword. When specified, size is interpreted as specifying the size of **min-thresh-separation** in bytes.

The **kilobytes** keyword is optional and is mutually exclusive with the **bytes** keyword. When specified, size is interpreted as specifying the size of min-thresh-separation in kilobytes.

Default **kilobytes**

fixed — The optional fixed keyword is used to force the inclusion of the defined **mbs-contribution** value in the parent policer's discard threshold calculations. If the **mbs-contribution** command is executed without the **fixed** keyword, the fixed calculation behavior for the priority level is removed.

Default **no mbs-contribution**

The **no mbs-contribution** command returns the policy's priority level's MBS contribution to the default value. When changed, the thresholds for the priority level and all higher priority levels for all instances of the parent policer will be recalculated.

radius-accounting-policy

Syntax **radius-accounting-policy** *acct-policy-name* [**duplicate** *acct-policy-name*]
no radius-accounting-policy

Context config>subscr-mgmt>sub-prof

Description This command specifies an existing RADIUS accounting policy to use to collect accounting statistics on this subscriber profile by RADIUS. This command is used independently of the **collect-stats** command.

Parameters *acct-policy-name* — Specifies an existing RADIUS based accounting policy.

duplicate *acct-policy-name* **—** Specifies the RADIUS accounting policy to be used to generate duplicate accounting information.

sla-profile-map

Syntax **sla-profile-map**

Context config>subscr-mgmt>sub-prof

Description This command enables the context to configure SLA profile mapping.

entry

Syntax **entry key** *sub-profile-string* **sub-profile** *sub-profile-name*
no entry key *sub-profile-string*

Context config>subscr-mgmt>sub-prof>sla-prof-map

Description This command configures SLA profile string mappings.

Parameters *sub-profile-string* — Specifies the subscriber profile string.

Values 16 characters maximum

sub-profile-name — Specifies the subscriber profile name.

Values 32 characters maximum

use-direct-map-as-default

Syntax [**no**] **use-direct-map-as-default**

Context config>subscr-mgmt>sub-prof>sla-prof-map

Description This command enables direct mapping of the SLA profile as default.

The **no** form of the command disables direct mapping,

sub-mcac-policy

Syntax **sub-mcac-policy** *policy-name*
no sub-mcac-policy

Context config>subscr-mgmt>sub-prof

Description This command references the policy template in which the mcac bandwidth limits are defined. Mcac for the subscriber is effectively enabled with this command when the sub-profile is applied to the subscriber. The bandwidth of the channels is defined in a different policy (under the **configure>router>mcac** context) and this policy is applied on the interface level as follows:

- For group-interfaces under the **configure>service>vrf>igmp>group-interface>mcac** context
- For regular interfaces under the **configure>service/router>igmp>interface>mcac** context

In case of HQoS Adjustment, it is mandatory that the sub-mcac-policy be created and applied to the subscriber. The sub-mac-policy does not have to contain any bandwidth constrains, but it has to be in a no shutdown state in order for HQoS Adjustment to work.

Default none

Parameters *policy-name* — Specifies the policy name configured in the config>subscr-mgmt>sub-mcac-policy context.

igmp-policy

Syntax **igmp-policy** *policy-name*
no igmp-policy

Context config>subscr-mgmt>sub-prof

Description This command will enable IGMP processing per subscriber host. Without this command IGMP states will not be maintained per subscriber hosts. The referenced policy is defined under the **configure>subscr-mgmt** context and can be only applied via the sub-profile.

The referenced policy contains entries such as:

- description statement
- import statement — IGMP filters
- egress-rate-modify statement—HQoS Adjustment
- mcast-redirection statement—redirection to alternate interface
- static statement—definition of static IGMP groups
- version statement —IGMP version
- fast-leave statement
- max-num-groups statement—t max number of multicast groups allowed

Parameters *policy-name* — Name of the IGMP policy for the subscriber. The policy itself is defined under the **configure>sub-mgmt** context.

hsmda

Syntax **hsmda**

Context config>subscr-mgmt>sub-prof

Description This command enables the context to configure egress and ingress HSMDA queue parameters.

egress-qos

Syntax **egress-queues**

Context config>subscr-mgmt>sub-prof>hsmda

Description This command enables the context to configure SAP egress QOS policy for the HSMDA egress queue.

ingress-qos

Syntax **ingress-queues**

Context config>subscr-mgmt>sub-prof>hsmda>egress-queues

Description This command enables the context to configure SAP egress QOS policy for the HSMDA ingress queue

agg-rate-limit

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

Context configure>port>ethernet>access>egress>vport

Description This command configures an aggregate rate for the vport. This command is mutually exclusive with the port-scheduler-policy command.

Default agg-rate-integer

Parameters *agg-rate-integer* — Kilobits per second (1,000 bits per second).

Values Values — 1 — 100,000,000.

agg-rate-limit

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

Context config>subscr-mgmt>sub-prof>hsmda>egress-qos

Description This command defines a maximum total rate for all subscriber egress queues for each subscriber associated with the sub-profile. The egress-agg-rate-limit command is mutually exclusive with the egress-scheduler-policy. When an egress-scheduler-policy is defined on the sub-profile, the egress-agg-rate-limit command will fail. If the egress-agg-rate-limit command is specified, at attempt to bind an egress-scheduler-policy to the sub-profile will fail.

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

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

The **no** form of the command removes the aggregate rate limit from the sub-profile.

Default no agg-rate-limit

Parameters *agg-rate* — Defines the maximum aggregate rate the egress queues associated with the subscriber profile may operate. The value is specified in kilobits per second in a base 10 context. A value of 1 indicates a rate of 1000 bits per second.

Values 1 — 40000000, max Kbps

qos

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

Context config>subscr-mgmt>sub-prof>hsmda>egress-qos

Description This command assigns a SAP egress QOS policy to the HSMDA egress queue.

Parameters *policy-id* — Specifies the policy ID of an existing QoS SAP egress policy.

Values 1 — 65535

qos

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

Context config>subscr-mgmt>sub-prof>hsmda>ingress-qos

Description This command assigns a SAP ingress QOS policy to the HSMDA ingress queue.

Parameters *policy-id* — Specifies the policy ID of an existing QoS SAP egress policy.

Values 1 — 65535

packet-byte-offset

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

Context config>subscr-mgmt>sub-prof>hsmda>egress-qos
config>subscr-mgmt>sub-prof>hsmda>ingress-qos>policer

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

The accounting functions affected include:

- Offered High Priority / In-Profile Octet Counter

- Offered Low Priority / Out-of-Profile Octet Counter
- Discarded High Priority / In-Profile Octet Counter
- Discarded Low Priority / Out-of-Profile Octet Counter
- Forwarded In-Profile Octet Counter
- Forwarded Out-of-Profile Octet Counter
- Peak Information Rate (PIR) Leaky Bucket Updates
- Committed Information Rate (CIR) Leaky Bucket Updates
- Queue Group Aggregate Rate Limit Leaky Bucket Updates

The secondary shaper leaky bucket, scheduler priority level leaky bucket and the port maximum rate updates are not affected by the configured packet-byte-offset. Each of these accounting functions are frame based and always include the preamble, DLC header, payload and the CRC regardless of the configured byte offset.

The packet-byte-offset command accepts either add or subtract as valid keywords which define whether bytes are being added or removed from each packet traversing the queue. Up to 31 bytes may be added to the packet and up to 32 bytes may be removed from the packet. An example use case for subtracting bytes from each packet is an IP based accounting function. Given a Dot1Q encapsulation, the command packet-byte-offset subtract 14 would remove the DLC header and the Dot1Q header from the size of each packet for accounting functions only. The 14 bytes are not actually removed from the packet, only the accounting size of the packet is affected.

As inferred above, the variable accounting size offered by the packet-byte-offset command is targeted at the queue and queue group level. The packet-byte-offset, when set, applies to all queues in the queue group. The accounting size of the packet is ignored by the secondary shapers, the scheduling priority level shapers and the scheduler maximum rate. The actual on-the-wire frame size is used for these functions to allow an accurate representation of the behavior of the subscriber's packets on an Ethernet aggregation network.

The packet-byte-offset value may be overridden at the queue-group level.

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

Values 0 — 31

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

Values 1 — 32

queue

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

Context config>subscr-mgmt>sub-prof>hsmda>ingress-qos>qos

Description This command specifies the HSMDA queue mapping for all packets in point-to-point services and unicast destined packets in multipoint services. Point-to-point services include epipe and other VLL type services. Multipoint services include IES, VPLS and VPRN services. The queue command does not apply to multicast, broadcast or unknown unicast packets within multipoint services (the multicast, broadcast and unknown commands must be used to define the queue mapping for non-unicast packets within a forwarding class). For Epipe services, the **queue** *queue-id* mapping applies to all packets, regardless of the packets destination MAC address.

Each forwarding class has a default queue ID based on the intrinsic hierarchy between the forwarding classes. Executing the queue command within the HSMDA context of a forwarding class with a different queue ID than the default overrides the default mapping. Multiple forwarding classes may be mapped to the same HSMDA queue ID.

The **no** form of the command returns the HSMDA queue mapping for queue to the default mapping for the forwarding class.

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

Values 1 — 8

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

rate

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

Context config>subscr-mgmt>sub-prof>hsmda>egress-qos>qos>queue
config>subscr-mgmt>sub-prof>hsmda>ingress-qos>queue
config>subscr-mgmt>sub-prof>hsmda>ingress-qos>policer

Description This command can be used to override specific attributes of the specified queue's Peak Information Rate (PIR) and the Committed Information Rate (CIR) parameters. The PIR defines the maximum rate that the queue can transmit packets out an egress interface (for SAP egress queues). Defining a PIR does not necessarily guarantee that the queue can transmit at the intended rate. The actual rate sustained by the queue can be limited by oversubscription factors or available egress bandwidth.

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

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

The **rate** command can be executed at any time, altering the PIR and CIR rates for all queues created through the association of the SAP egress QoS policy with the *queue-id*.

The **no** form of the command returns all queues created with the *queue-id* by association with the QoS policy to the default PIR and CIR parameters (**max**, 0).

Default **rate max cir 0 —** The **max** default specifies the amount of bandwidth in kilobits per second (thousand bits per second). The **max** value is mutually exclusive to the **pir-rate** value.

Parameters *pir-rate* — Defines the administrative PIR rate, in kilobits, for the queue. When the **rate** command is executed, a valid PIR setting must be explicitly defined. When the **rate** command has not been executed, the default PIR of **max** is assumed.
Fractional values are not allowed and must be given as a positive integer.

The actual PIR rate is dependent on the queue's **adaptation-rule** parameters and the actual hardware where the queue is provisioned.

Values 1 — 100000000

Default max

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

Values 0 — 100000000, **max**, **sum**

Default 0

slope-policy

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

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

Context config>subscr-mgmt>sub-prof>hsmda>egress-qos>qos>queue

Description This command specifies an existing slope policy name. The policy contains the Maximum Buffer Size (MBS) that will be applied to the queue and the high and low priority RED slope definitions. The function of the MBS and RED slopes is to provide congestion control for an HSMDA queue. The MBS parameter defines the maximum depth a queue may reach when accepting packets. The low and high priority RED slopes provides for random early detection of congestion and slope based discards based on queue depth.

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

Default HSMDA Slope Policy

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

The no form of the command removes the slope policy from the subscriber profile HSMDA configuration.

wrr-weight

Syntax **wrr-weight** *value*
no wrr-weight

Context config>subscr-mgmt>sub-prof>hsmda>egress-qos>qos>queue

Description This command assigns the weight value to the HSMDA queue.

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

Parameters *percentage* — Specifies the weight for the HSMDA queue.

Values 1— 32

wrr-policy

Syntax **wrr-policy** *hsmda-wrr-policy-name*
no wrr-policy

Context config>subscr-mgmt>sub-prof>hsmda>egress-qos>qos

Description This command associates an existing HSMDA weighted-round-robin (WRR) scheduling loop policy to the HSMDA queue.

Parameters *hsmda-wrr-policy-name* — Specifies the existing HSMDA WRR policy name to associate to the queue.

SLA Profile Commands

sla-profile

Syntax **sla-profile** *sla-profile-name*

Context config>subscr-mgmt

Description This command configures an SLA profile mapping. Hosts associated with a subscriber are subdivided into Service Level Agreement (SLA) profiles. For each subscriber host an SLA profile can be specified. For a subscriber host, the SLA profile determines:

- The QoS-policies to use

  –The classification

  –The queues

  –The queue mapping

- The IP filters to use

The SLA profile also has the attribute host-limit which limits the total number of hosts (belonging to the same subscriber) on a certain SAP that can be using this SLA profile.

Default none

Parameters *sla-profile-name* — Specifies the name of the SLA profile.

egress

Syntax **egress**

Context config>subscr-mgmt>sla-profile

Description This command configures egress parameters for the SLA profile.

ingress

Syntax **ingress**

Context config>subscr-mgmt>sla-profile

Description This command configures ingress parameters for the SLA profile.

host-limit

Syntax **host-limit** *max-number-of-hosts* [**remove-oldest**]
no host-limit

Context config>subscr-mgmt>sla-profile

Description This command specifies the maximum number of hosts for this SLA profile.

Parameters *max-number-of-hosts* — Specifies the host limit for this SLA profile.

Values 1 — 100

remove-oldest — When the keywords **remove-oldest** are specified, the oldest subscriber host will be removed when the host limit is reached.

ip-filter

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

Context config>subscr-mgmt>sla-profile>egress
config>subscr-mgmt>sla-profile>ingress

Description This command configures an egress or ingress IP filter.

Parameters *filter-id* — Specify an existing IP filter policy ID.

Values 1 — 65535

SLA Profile QoS Commands

qos

Syntax **qos** *sap-egress-policy-id* [*vport-scheduler*|*port-scheduler*] [**force**]
no qos

Context config>subscr-mgmt>sla-prof>egress

Description This command specifies the egress QoS policy applicable to this SLA profile. The policy must already be defined in the **configure>qos>sap-egress** context.

Default 1

Parameters *sap-egress-policy-id* — Specifies the egress policy to be applied to the egress SLA profile.

Values 1 — 65535

vport-scheduler | *port-scheduler* — Specifies if a host queue with the port-parent option enabled should be scheduled within the context of a vport port scheduler policy or a the port's port scheduler policy.

force — Forces a policy change.

qos

Syntax **qos** *policy-id* [**shared-queuing** | **multipoint-shared** | **service-queuing**] [**force**]
no qos

Context config>subscr-mgmt>sla-prof>ingress

Description This command specifies the ingress QoS policy applicable to this SLA profile. The policy must already be defined in the **configure>qos>sap-ingress** context.

Default qos 1

Parameters *sap-ingress-policy-id* — Specifies the policy to be applied to the ingress SLA profile.

Values 1 — 65535

shared-queuing — This keyword is mutually exclusive with the **multipoint-shared** and **service-queuing** keywords to specify the policy used by this SAP. When the value of this object is null it means that the SAP will use individual ingress QoS queues, instead of the shared ones.

multipoint-shared — This keyword is mutually exclusive with the **shared-queuing** and **service-queuing** keywords. When multipoint-shared is specified, the ingress forwarding plane will conserve hardware queues by performing two tier queuing on ingress unicast and multipoint packets through the SAP. Unicast service queues defined in the SAP ingress QoS policy are created for the SAP on the ingress forwarding plane without regard for the switch fabric destinations to which the SAP may need to forward (other destinations in the VPLS context). The multipoint queues defined in the SAP ingress QoS policy are not created for the SAP. Instead, all multipoint traffic is mapped to the unicast queues based on forwarding class in the first pass. In

the second pass the unicast packets will be mapped to the unicast shared queues while the multipoint traffic will be mapped to the multipoint shared queues.

service-queuing — This keyword is mutually exclusive with the **multipoint-shared** and **shared-queuing** keywords to state that service queueing is needed.

force — Forces a policy change.

queue

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

Context config>subscr-mgmt>sla-prof>egress>qos
config>subscr-mgmt>sla-prof>ingress>qos

Description This command configures the context to configure egress or ingress queue parameters. Parameters defined in the **config>qos>sap-egress** *policy-id* or the **config>qos>sap-ingress** *policy-id* context are overridden by parameters specified in the subscriber management SLA profile context.

The classification and the queue mapping are shared by all the hosts on the same complex that use the same QoS policy (specified in the **sla-profile** SAP egress and SAP ingress policy IDs).

The queues are shared by all the hosts (of the same subscriber) on the same SAP that are using the same SLA profile. Queues are instantiated when, on a given SAP, a host of a subscriber is the first to use a certain SLA profile. This instantiation is referred to as an SLA profile instance.

The **no** form of the command removes the *queue-id* from the SLA profile.

Default none

Parameters *queue-id* — Specifies the *queue-id* for the SAP egress or ingress queue, expressed as a decimal integer. The *queue-id* uniquely identifies the queue within the profile.

Default none

avg-frame-overhead

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

Context config>subscr-mgmt>sla-prof>egress>qos>queue

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

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

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

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

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

Port scheduler operation using frame transformed rates — The port scheduler uses the frame based rates to determine the maximum rates that each queue may receive during the within-cir and above-cir

bandwidth allocation passes. During the within-cir pass, a queue may receive up to its frame based within-cir offered-load. The maximum it may receive during the above-cir pass is the difference between the frame based within-pir offered load and the amount of actual bandwidth allocated during the within-cir pass.

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

The **no** form of this command restores the average frame overhead parameter for the queue to the default value of 0 percent. When set to 0, the system uses the packet based queue statistics for calculating port scheduler priority bandwidth allocation. If the no avg-frame-overhead command is executed in a queue-override queue id context, the avg-frame-overhead setting for the queue within the sap-egress QoS policy takes effect.

Default 0

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

Values 0 — 100

burst-limit

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

Context config>subscr-mgmt>sla-prof>egress>qos>queue
config>subscr-mgmt>sla-prof>ingress>qos>queue

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

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

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

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

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

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

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

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

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

cbs

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

Context config>subscr-mgmt>sla-prof>egress>qos>queue
config>subscr-mgmt>sla-prof>ingress>qos>queue

Description This command can be used to override specific attributes of the specified queue's CBS parameters. It is permissible, and possibly desirable, to oversubscribe the total CBS reserved buffers for a given access port egress buffer pool. Oversubscription may be desirable due to the potential large number of service queues and the economy of statistical multiplexing the individual queues' CBS settings into the defined reserved total.

When oversubscribing the reserved total, it is possible for a queue depth to be lower than its CBS setting and still not receive a buffer from the buffer pool for an ingress frame. As more queues are using their CBS buffers and the total in use exceeds the defined reserved total, essentially the buffers are being removed from the shared portion of the pool without the shared in use average and total counts being decremented. This can affect the operation of the high and low priority RED slopes on the pool, causing them to miscalculate when to start randomly drop packets.

If the CBS value is larger than the MBS value, an error will occur, preventing the CBS change.

The **no** form of this command returns the CBS size to the size as configured in the QoS policy.

Default no cbs

Parameters *size-in-kbytes* — The size parameter is an integer expression of the number of kilobytes reserved for the queue. If a value of 10KBytes is desired, enter the value 10. A value of 0 specifies that no reserved buffers are required by the queue (a minimal reserved size can still be applied for scheduling purposes).

Values 0 — 131072 or default

high-prio-only

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

Context config>subscr-mgmt>sla-prof>egress>qos>queue
config>subscr-mgmt>sla-prof>ingress>qos>queue

Description This command configures the value of the percentage of buffer space for the queue, used exclusively by high priority packets. The specified value overrides the default value for the context.

The priority of a packet can only be set in the SAP ingress QoS policy and is only applicable on the ingress queues for a SAP. The **high-prio-only** parameter is used to override the default value derived from the **network-queue** command.

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

The **no** form of this command returns high-prio-only to the size as configured in the QoS policy.

Default no high-prio-only

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

Values 0 — 100 | default

mbs

Syntax **mbs** *size-in-kbytes*
no mbs

Context config>subscr-mgmt>sla-prof>egress>qos>queue

Description This command configures the maximum size for the queue.

The sum of the MBS for all queues on an egress access port can oversubscribe the total amount of buffering available. When congestion occurs and buffers become scarce, access to buffers is controlled by the RED slope a packet is associated with. A queue that has not exceeded its MBS size is not guaranteed that a buffer will be available when needed or that the packet's RED slope will not force the discard of the packet. Setting proper CBS parameters and controlling CBS oversubscription is one major safeguard to queue starvation (when a queue does not receive its fair share of buffers). Another is properly setting the RED slope parameters for the needs of services on this port or channel.

If the CBS value is larger than the MBS value, an error will occur, preventing the MBS change.

The **no** form of this command returns the MBS size to the size as configured in the QoS policy.

Default no mbs

Parameters *size-in-kbytes* — The size parameter is an integer expression of the maximum number of kilobytes of buffering allowed for the queue.For a value of 100 kbps, enter the value 100. A value of 0 causes the queue to discard all packets.

Values 0 — 131072 or default

mbs

Syntax **mbs** *size* [**bytes** | **kilobytes**]
no mbs

Context config>subscr-mgmt>sla-prof>ingress>qos>queue

Description The Maximum Burst Size (MBS) command configures the explicit definition of the maximum amount of buffers allowed for a specific queue.

The MBS value is used by a queue to determine whether it has exhausted all of its buffers while enqueuing packets. Once the queue has exceeded the amount of buffers allowed by MBS, all packets are discarded until packets have been drained from the queue.

The sap-ingress context for mbs provides a mechanism for overriding the default maximum size for the queue.

The sum of the MBS for all queues on an ingress access port can oversubscribe the total amount of buffering available. When congestion occurs and buffers become scarce, access to buffers is controlled by the RED slope a packet is associated with. A queue that has not exceeded its MBS size is not guaranteed that a buffer will be available when needed or that the packet's RED slope will not force the discard of the packet. Setting proper CBS parameters and controlling CBS oversubscription is one major safeguard to queue starvation (when a queue does not receive its fair share of buffers). Another is properly setting the RED slope parameters for the needs of services on this port or channel.

If the CBS value is larger than the MBS value, an error will occur, preventing the MBS change.

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

The **no** form of this command returns the MBS size to the size as configured in the QoS policy.

Default no mbs

Parameters *size* [**bytes** | **kilobytes**] — The size parameter is an integer expression of the maximum number of bytes or kilobytes of buffering allowed for the queue. For a value of 100 kbps enter the value 100 and specify the **kilobytes** parameter. A value of 0 causes the queue to discard all packets.

Values 0 — 131072 or default

rate

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

Context **config>subscr-mgmt>sla-prof>egress>qos>queue**

Description This command defines the administrative Peak Information Rate (PIR) and the administrative Committed Information Rate (CIR) parameters for the queue. The PIR defines the maximum rate that the queue can transmit packets out an egress interface (for SAP egress queues). Defining a PIR does not necessarily guarantee that the queue can transmit at the intended rate. The actual rate sustained by the queue can be limited by oversubscription factors or available egress bandwidth.

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

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

The **rate** command can be executed at anytime, altering the PIR and CIR rates for all queues created through the association of the SAP egress QoS policy with the *queue-id*.

The **no** form of the command returns all queues created with the *queue-id* by association with the QoS policy to the default PIR and CIR parameters (**max**, 0).

Default no rate

Parameters *pir-rate* — Defines the administrative PIR rate, in kilobits, for the queue. When the **rate** command is executed, a valid PIR setting must be explicitly defined. When the **rate** command has not been executed, the default PIR of **max** is assumed.
Fractional values are not allowed and must be given as a positive integer.

The actual PIR rate is dependent on the queues **adaptation-rule** parameters and the actual hardware where the queue is provisioned.

Values 1 — 100000000, max

Default **max**

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

Values 0 — 100000000, **max**

Default 0

qos-marking-from-sap

Syntax [**no**] **qos-marking-from-sap**

Context configure>subscr-mgmt>sla-profile>egress

Description This command sets the QoS policy from which the egress QoS marking rules are applied. Note that if applied to a managed SAP, the default SAP-egress qos-policy (sap-egress 1) cannot be changed.

The **no** form of the command reverts to the egress QoS marking defined in SAP-egress policy defined at sla-profile level.

Default qos-marking-from-sap

report-rate

Syntax **report-rate agg-rate-limit**
report-rate scheduler *scheduler-name*
report-rate pppoe-actual-rate
report-rate rfc5515-actual-rate
no report-rate

Context config>subscr-mgmt>sla-prof>ingress
config>subscr-mgmt>sla-prof>egress

Description This command configures the source for Tx and Rx connect speeds in AVP 38 (Rx Connect Speed) and AVP 24 (Tx Connect Speed) of an L2TP session established on a LAC.

Default no report-rate – Rates takes from the physical port speed.

Parameters **agg-rate-limit** — (egress only) rate taken from:

1. The agg-rate RADIUS override (RADIUS VSA "Alc-Subscriber-QoS-Override" in a RADIUS Access-Accept message) if present.
2. The configured agg-rate-limit in the **config>subscr-mgmt>sub-prof>egr** context.
3. Fall back to the default (no report-rate).

scheduler *scheduler-name* — Specifies the rate taken from the **scheduler** *scheduler-name*. If the **scheduler** *scheduler-name* is not present in the scheduler-policy configured in the **config>subscr-mgmt>sub-prof>egr** context, fall back to the default (no report-rate)

pppoe-actual-rate — Specifies rates taken from the "DSL Line characteristics" PPPoE tags (Actual Data Rate Upstream/Downstream) if present; otherwise fall back to the default (no report-rate).

report-rate rfc5515-actual-rate — Puts the same value as the transmitted Actual-Data-Rate-Upstream AVP in the Rx-Connect-Speed AVP, and the same value as the transmitted Actual-Data-Rate-Downstream AVP in the Tx-Connect-Speed AVP.

use-ingress-l2tp-dscp

Syntax [**no**] **use-ingress-l2tp-dscp**

Context config>subscr-mgmt>sla-prof>egress

Description This command enables the use of the DSCP marking taken from the L2TP header received on an L2TP Access Concentrator (LAC) for egress classification for the subscriber host using the associated sla-profile.

This command is ignored if the ingress packet is not identified as an L2TP packet.

Default no use-ingress-l2tp-dscp

one-time-http-redirection

Syntax **one-time-http-redirection** *filter-id*
[no] one-time-http-redirection

Context config>subscr-mgmt>sla-prof

Description This command specify the one-time http redirection filter id. This filter will apply to the host when host is created, and will be replaced by the sla-profile ingress filter (configured in the **config>subscr-mgmt>sla-prof>ingress** context) after first HTTP request from host has been redirected.

Note: system does not check if the configured filter include http-redirection entry. If the filter does not include the http-redirection then it will not be replaced in future.

If 7750 receives filter insertion via CoA or access-accept when one-time redirection filter is still active then the received filter entries will only be applied to the sla-profile ingress filter. And after 1st http redirection, the original sla-profile ingress filter + received filter will replace the redirection filter.

Default no

Parameters *filter-id* — Specifies the id of filter that is used for HTTP redirection.

rate

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

Context config>subscr-mgmt>sla-prof>ingress>qos>queue

Description This command defines the administrative Peak Information Rate (PIR) and the administrative Committed Information Rate (CIR) parameters for the queue. The PIR defines the maximum rate that the queue can transmit packets through the switch fabric (for SAP ingress queues). Defining a PIR does not necessarily guarantee that the queue can transmit at the intended rate. The actual rate sustained by the queue can be limited by oversubscription factors or available egress bandwidth.

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

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

The **rate** command can be executed at anytime, altering the PIR and CIR rates for all queues created through the association of the SAP ingress or SAP egress QoS policy with the *queue-id*.

The **no** form of the command returns all queues created with the *queue-id* by association with the QoS policy to the default PIR and CIR parameters (**max**, 0).

Default no rate

Parameters *pir-rate* — Defines the administrative PIR rate, in kilobits, for the queue. When the **rate** command is executed, a valid PIR setting must be explicitly defined. When the **rate** command has not been

executed, the default PIR of **max** is assumed.
Fractional values are not allowed and must be given as a positive integer.

The actual PIR rate is dependent on the queues **adaptation-rule** parameters and the actual hardware where the queue is provisioned.

Values 1 — 100000000, max

Default **max**

cir-rate — Specifies the **cir** parameter used by the queue. When the **rate** command is executed, a CIR setting is optional. When the **rate** command has not been executed or the **cir** parameter is not explicitly specified, the default CIR (0) is assumed.
Fractional values are not allowed and must be given as a positive integer.

Values 0 — 100000000, **max**

Default 0

policer

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

Context config>subscr-mgmt>sla-prof>ingress>qos
config>subscr-mgmt>sla-prof>egress>qos
config>subscr-mgmt>sub-profile>hsmda>ingress-qos>qos>policer

Description This command is used in the sap-ingress and sap-egress QoS policies to create, modify or delete a policer. Policers are created and used in a similar manner to queues. The policer ID space is separate from the queue ID space, allowing both a queue and a policer to share the same ID. The sap-ingress policy may have up to 32 policers (numbered 1 through 32) may be defined while the sap-egress QoS policy supports a maximum of 8 (numbered 1 through 8). While a policer may be defined within a QoS policy, it is not actually created on SAPs or subscribers associated with the policy until a forwarding class is mapped to the policer's ID.

All policers must be created within the QoS policies. A default policer is not created when a sap-ingress or sap-egress QoS policy is created.

Once a policer is created, the policer's metering rate and profiling rates may be defined as well as the policer's maximum and committed burst sizes (MBS and CBS respectively). Unlike queues which have dedicated counters, policers allow various stat-mode settings that define the counters that will be associated with the policer. Another supported feature—packet-byte-offset—provides a policer with the ability to modify the size of each packet based on a defined number of bytes.

Once a policer is created, it cannot be deleted from the QoS policy unless any forwarding classes that are mapped to the policer are first moved to other policers or queues.

The system will allow a policer to be created on a SAP QoS policy regardless of the ability to support policers on objects where the policy is currently applied. The system only scans the current objects for policer support and sufficient resources to create the policer when a forwarding class is first mapped to the policer ID. If the policer cannot be created due to one or more instances of the policy not supporting policing or having insufficient resources to create the policer, the forwarding class mapping will fail.

The **no** form of this command is used to delete a policer from a sap-ingress or sap-egress QoS policy. The specified policer cannot currently have any forwarding class mappings for the removal of the policer to succeed. It is not necessary to actually delete the policer ID for the policer instances to be removed from SAPs or subscribers associated with the QoS policy once all forwarding classes have been moved away from the policer. It is automatically deleted from each policing instance although it still appears in the QoS policy.

Parameters *policer-id* — The *policer-id* must be specified when executing the policer command. If the specified ID already exists, the system enters that policer's context to allow the policer's parameters to be modified. If the ID does not exist and is within the allowed range for the QoS policy type, a context for the policer ID will be created (depending on the system's current create keyword requirements which may require the create keyword to actually add the new policer ID to the QoS policy) and the system will enter that new policer's context for possible parameter modification.

Values 1—32

cbs

Syntax **cbs** {*size* [**bytes** | **kilobytes**] | **default**}
no cbs

Context config>subscr-mgmt>sla-prof>ingress>qos>policer
config>subscr-mgmt>sla-prof>egress>qos>policer
config>subscr-mgmt>sub-profile>hsmda>ingress-qos>qos>queue
config>subscr-mgmt>sub-profile>hsmda>ingress-qos>qos>policer

Description This command is used to configure the policer's CIR leaky bucket's exceed threshold. The CIR bucket's exceed threshold represents the committed burst tolerance allowed by the policer. If the policer's forwarding rate is equal to or less than the policer's defined CIR, the CIR bucket depth hovers around the 0 depth with spikes up to the maximum packet size in the offered load. If the forwarding rate increases beyond the profiling rate, the amount of data allowed to be in-profile above the rate is capped by the threshold.

The policer's **cbs** size defined in the QoS policy may be overridden on an **sla-profile** or SAP where the policy is applied.

The no form of this command returns the policer to its default CBS size.

Default **none**

Parameters *size* [**bytes** | **kilobytes**] — The *size* parameter is required when specifying **cbs** and is expressed as an integer representing the required size in either bytes or kilobytes. The default is kilobytes. The optional **byte** and **kilobyte** keywords are mutually exclusive and are used to explicitly define whether size represents bytes or kilobytes.

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

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

Values 1 — 4194304

Default **kilobyte**

mbs

Syntax **mbs** {*size* [**bytes** | **kilobytes**] | **default**}
no mbs

Context config>subscr-mgmt>sla-prof>ingress>qos>policer
config>subscr-mgmt>sla-prof>egress>qos>policer
config>subscr-mgmt>sub-profile>hsmda>egress-qos>qos>queue
config>subscr-mgmt>sub-profile>hsmda>ingress-qos>qos>queue
config>subscr-mgmt>sub-profile>hsmda>ingress-qos>qos>policer

Description This command is used to configure the policer's PIR leaky bucket's high priority violate threshold. The **high-prio-only** command is applied to the MBS value to derive the bucket's low priority violate threshold. For ingress, trusted in-profile packets and un-trusted high priority packets use the policer's high priority violate threshold while trusted out-of-profile and un-trusted low priority packets use the policer's low priority violate threshold. At egress, in-profile packets use the policer's high priority violate threshold and out-of-profile packets use the policer's low priority violate threshold.

The PIR bucket's violate threshold represent the maximum burst tolerance allowed by the policer. If the policer's offered rate is equal to or less than the policer's defined rate, the PIR bucket depth hovers around the 0 depth with spikes up to the maximum packet size in the offered load. If the offered rate increases beyond the metering rate, the amount of data allowed above the rate is capped by the threshold. The low priority violate threshold provides a smaller burst size for the lower priority traffic associated with the policer. Since all lower priority traffic is discarded at the lower burst tolerance size, the remaining burst tolerance defined by **high-prio-only** is available for the higher priority traffic.

The policer's mbs size defined in the QoS policy may be overridden on an sla-profile or SAP where the policy is applied.

The no form of this command returns the policer to its default MBS size.

Default None

Parameters *size* [**bytes** | **kilobytes**] — The *size* parameter is required when specifying **mbs** and is expressed as an integer representing the required size in either bytes or kilobytes. The default is kilobytes. The optional **byte** and **kilobyte** keywords are mutually exclusive and are used to explicitly define whether size represents bytes or kilobytes.

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

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

Values 1 — 4194304

Default **kilobyte**

packet-byte-offset

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

Context config>subscr-mgmt>sla-prof>ingress>qos>policer
config>subscr-mgmt>sla-prof>egress>qos>policer

Description This command is used to modify the size of each packet handled by the policer by adding or subtracting a number of bytes. The actual packet size is not modified; only the size used to determine the bucket depth impact is changed. The **packet-byte-offset** command is meant to be an arbitrary mechanism the can be used to either add downstream frame encapsulation or remove portions of packet headers. Both the policing metering and profiling throughput is affected by the offset as well as the stats associated with the policer.

When child policers are adding to or subtracting from the size of each packet, the parent policer's **min-thresh-separation** value should also need to be modified by the same amount.

The policer's **packet-byte-offset** defined in the QoS policy may be overridden on an **sla-profile** or SAP where the policy is applied.

The **no** version of this command is used to remove per packet size modifications from the policer.

Parameters **add** *bytes* — The **add** keyword is mutually exclusive to the **subtract** keyword. Either **add** or **subtract** must be specified. When **add** is defined the corresponding bytes parameter specifies the number of bytes that is added to the size each packet associated with the policer for rate metering, profiling and accounting purposes. From the policer's perspective, the maximum packet size is increased by the amount being added to the size of each packet.

Values 1 — 32

Default None

subtract *bytes* — The **subtract** keyword is mutually exclusive to the **add** keyword. Either **add** or **subtract** must be specified. When b is defined the corresponding bytes parameter specifies the number of bytes that is subtracted from the size of each packet associated with the policer for rate metering, profiling and accounting purposes. From the policer's perspective, the maximum packet size is reduced by the amount being subtracted from the size of each packet.

Values 0—31

Default None

rate

Syntax **rate** {**max** | **kilobits-per-second**} [**cir** {**max** | **kilobits-per-second**}]
no rate

Context config>subscr-mgmt>sla-prof>ingress>qos>policer
config>subscr-mgmt>sla-prof>egress>qos>policer

Description This command is used to configure the policer's metering and optional profiling rates. The metering rate is used by the system to configure the policer's PIR leaky bucket's decrement rate while the profiling rate configures the policer's CIR leaky bucket's decrement rate. The decrement function empties the bucket while packets applied to the bucket attempt to fill it based on the each packets size. If

the bucket fills faster than how much is decremented per packet, the bucket's depth eventually reaches it's exceed (CIR) or violate (PIR) threshold. The **cbs**, **mbs**, and **high-prio-only** commands are used to configure the policer's PIR and CIR thresholds.

If a packet arrives at the policer while the bucket's depth is less than the threshold associated with the packet, the packet is considered to be conforming to the bucket's rate. If the bucket depth is equal to or greater than the threshold, the packet is considered to be in the exception state. For the CIR bucket, the exception state is exceeding the CIR rate while the PIR bucket's exception state is violating the PIR bucket rate. If the packet is violating the PIR, the packet is marked red and will be discarded. If the packet is not red, it may be green or yellow based on the conforming or exceeding state from the CIR bucket.

When a packet is red neither the PIR or CIR bucket depths are incremented by the packets size. When the packet is yellow the PIR bucket is incremented by the packet size, but the CIR bucket is not. When the packet is green, both the PIR and CIR buckets are incremented by the packet size. This ensures that conforming packets impact the bucket depth while exceeding or violating packets do not.

The policer's **adaptation-rule** command settings are used by the system to convert the specified rates into hardware timers and decrement values for the policer's buckets.

By default, the policer's metering rate is **max** and the profiling rate is 0 Kbps (all packets out-of-profile).

The **rate** settings defined for the policer in the QoS policy may be overridden on an **sla-profile** or SAP where the policy is applied.

The **no** form of this command is used to restore the default metering and profiling rate to a policer.

Parameters {**max** | *kilobits-per-second*} — Specifying the keyword **max** or an explicit *kilobits-per-second* parameter directly following the rate command is required and identifies the policer's metering rate for the PIR leaky bucket. When the policer is first created, the metering rate defaults to max. The *kilobits-per-second* value must be expressed as an integer and defines the rate in kilobits-per-second. The integer value is multiplied by 1,000 to derive the actual rate in bits-per-second.

Values **max** or 0—20,000,000

cir {**max** | *kilobits-per-second*} — The optional **cir** keyword is used to override the default CIR rate of the policer. Specifying the keyword max or an explicit *kilobits-per-second* parameter directly following the cir keyword is required and identifies the policer's profiling rate for the CIR leaky bucket. When the policer is first created, the profiling rate defaults to 0 Kbps. The *kilobits-per-second* value must be expressed as an integer and defines the rate in kilobits-per-second. The integer value is multiplied by 1,000 to derive the actual rate in bits-per-second.

Values **max** or 0—20,000,000

stat-mode

Syntax **stat-mode** {**no-stats** | **minimal** | **offered-profile-no-cir** | **offered-priority-no-cir** | **offered-limited-profile-cir** | **offered-profile-cir** | **offered-priority-cir** | **offered-total-cir**}
no stat mode

Context config>subscr-mgmt>sla-prof>ingress>qos>policer
config>subscr-mgmt>sla-prof>egress>qos>policer
config>subscr-mgmt>sub-profile>hsmda>ingress-qos>qos>policer

Description This command is used to configure the forwarding plane counters that allow offered, output and discard accounting to occur for the policer. An ingress policer has multiple types of offered packets (explicit in-profile, explicit out-of-profile, high priority or low priority) and each of these offered types is interacting with the policer's metering and profiling functions resulting in colored output packets (green, yellow and red). Due to the large number of policers, it is not economical to allocate counters in the forwarding plane for all possible offered packet types and output conditions. Many policers will not be configured with a CIR profiling rate and not all policers will receive explicitly profiled offered packets. The **stat-mode** command allows provisioning of the number of counters each policer requires and how the offered packet types and output conditions should be mapped to the counters.

While a **no-stats** mode is supported which prevents any packet accounting, the use of the policer's **parent** command requires at the policer's **stat-mode** to be set at least to the **minimal** setting so that offered stats are available for the policer's Fair Information Rate (FIR) to be calculated. Once a policer has been made a child to a parent policer, the **stat-mode** cannot be changed to **no-stats** unless the policer parenting is first removed.

Each time the policer's **stat-mode** is changed, any previous counter values are lost and any new counters are set to zero.

Each mode uses a certain number of counters per policer instance that are allocated from the forwarding plane's policer counter resources. You can view the total/allocated/free stats by using the **tools dump system-resources** command. If insufficient counters exist to implement a mode on any policer instance, the **stat-mode** change will fail and the previous mode will continue unaffected for all instances of the policer.

The default **stat-mode** when a policer is created within the policy is **minimal**.

The **stat-mode** setting defined for the policer in the QoS policy may be overridden on an **sla-profile** or SAP where the policy is applied. If insufficient policer counter resources exist to implement the override, the **stat-mode** override command will fail. The previous **stat-mode** setting active for the policer will continue to be used by the policer.

The **no** form of this command attempts to return the policer's stat-mode setting to minimal. The command will fail if insufficient policer counter resources exist to implement minimal where the QoS policer is currently applied and has a forwarding class mapping.

Parameters **no-stats —** Counter resource allocation:0

The policer does not have any forwarding plane counters allocated and cannot provide offered, discard and forward statistics. A policer using no-stats cannot be a child to a parent policer and the policer's parent command will fail.

When **collect-stats** is enabled, the lack of counters causes the system to generate the following statistics:

a. offered-in = 0

b. offered-out = 0

c. discard-in = 0

d. discard-out = 0

e. forward-in = 0

f. forward-out= 0

Counter 0 indicates that the accounting statistic returns a value of zero.

minimal — Counter resource allocation:1

The default **stat-mode** for a policer is **minimal**. The **minimal** mode allocates 1 forwarding plane offered counter and one traffic manager discard counter. The forwarding counter is derived by subtracting the discard counter from the offered counter. The counters do not differentiate possible offered types (profile or priority) and do not count green or yellow output. This does not prevent the policer from supporting different offered packet types and does not prevent the policer from supporting a CIR rate.

This counter mode is useful when only the most basic accounting information is required.

The counters are used in the following manner:

1. 'offered = profile in/out, priority high/low

2. 'discarded = Same as 1

3. 'forwarded = Derived from 1 - 2

When **collect-stats** is enabled, the counters are used by the system to generate the following statistics:

a. offered-in = 1

b. offered-out= 0

c. discard-in = 2

d. discard-out= 0

e. forward-in = 3

f. 'orward-out= 0

Counter 0 indicates that the accounting statistic returns a value of zero.

With **minimal** enabled as the policer **stat-mode**, the SAP offered stats for the policer returned via MIB query and CLI show commands will return the following values:

i. offered-in = 1

ii. offered-out= 0

iii. offered-undefined= 0

iv. offered-managed= 0(IMPM managed packets are not redirected from the policer)

Counter 0 indicates that the SAP policer statistic returns a value of zero.

offered-profile-no-cir — Counter resource allocation:2

The **offered-profile-no-cir** mode allocates two forwarding plane offered counters and two traffic manager discard counters.

The **offered-profile-no-cir** mode is most useful when the policer is receiving only in-profile and out-of-profile pre-marked (and trusted) packets. It is expected that in this instance a CIR rate will not be defined since all packet are already pre-marked. This mode does not prevent the policer from receiving un-trusted (color undefined) nor does it prevent the policer from being configured with a CIR rate.

The counters are used in the following manner:

1. offered-in = profile in

2. offered-out= profile out, priority high/low

3. dropped-in= Same as 1

4. dropped-out= Same as 2

5. forwarded-in= Derived from 1 - 3

6. forwarded-out= Derived from 2 - 4

When **collect-stats** is enabled, the counters are used by the system to generate the following statistics:

a. offered-in = 1

b. offered-out= 2

c. discard-in = 3

d. discard-out= 4

e. forward-in = 5

f. forward-out= 6

With **offered-profile-no-cir** enabled as the policer **stat-mode**, the SAP offered stats for the policer returned via MIB query and CLI show commands will return the following values:

i. offered-in = 1

ii. offered-out= 2

iii. offered-undefined= 0

iv. offered-managed= 0(IMPM managed packets are not redirected from the policer)

Counter 0 indicates that the SAP policer statistic returns a value of zero.

offered-priority-no-cir — Counter resource allocation:2

The **offered-priority-no-cir** mode allocates two forwarding plane offered counters and two traffic manager discard counters.

The **offered-priority-no-cir** mode is most useful when the policer is receiving only un-trusted packets and the ingress priority high and priority low classification options are being used without a CIR profiling rate defined. This mode does not prevent the policer from receiving trusted packets that are pre-marked in-profile or out-of-profile nor does it prevent the policer from being configured with a CIR rate.

The counters are used in the following manner:

1. offered-high = profile in, priority high

2. offered-low= profile out, priority low

3. dropped-high= Same as 1

4. dropped-low= Same as 2

5. forwarded-high= Derived from 1 - 3

6. forwarded-low= Derived from 2 - 4

When **collect-stats** is enabled, the counters are used by the system to generate the following statistics:

a. offered-high= 1

b. offered-low= 2

c. discard-high= 3

d. discard-low= 4

e. forward-high= 5

f. forward-low= 6

With **offered-priority-no-cir** enabled as the policer **stat-mode**, the SAP offered stats for the policer returned via MIB query and CLI show commands will return the following values:

i. offered-high= 1

ii. offered-low= 2

iii. offered-undefined= 0

iv. offered-managed= 0(IMPM managed packets are not redirected from the policer)

Counter 0 indicates that the SAP policer statistic returns a value of zero.

offered-limited-profile-cir — Counter resource allocation:3

The **offered-limitied-profile-cir** mode allocates three forwarding plane offered counters and three traffic manager discard counters.

The **offered-limited-profile-cir** mode is most useful when the policer is receiving trusted out-of-profile (profile out but no profile in) traffic and un-trusted packets are being applied to a defined CIR profiling rate. This mode does not prevent the policer from receiving trusted in-profile packets.

The counters are used in the following manner:

1. offered-undefined-that-turned-green= profile in, priority high/low

2. offered-undefined-that-turned-yellow-or-red= priority high/low

3. offered-out-that-stayed-yellow-or-turned-red= profile out

4. dropped-undefined-that-turned-green= Same as 1

5. dropped-undefined-that-turned-yellow-or-red= Same as 2

6. dropped-out-that-turned-yellow-or-red= Same as 3

7. forwarded-undefined-that-turned-green= Derived from 1 - 4

8. forwarded-undefined-that-turned-yellow= Derived from 2 - 5

9. forwarded-out-that-turned-yellow= Derived from 3 - 6

When **collect-stats** is enabled, the counters are used by the system to generate the following statistics:

a. offered-in = 0

b. offered-out= 1 + 2 + 3

c. discard-in = 0

d. discard-out= 4 + 5 + 6

e. forward-in = 7

f. 'orward-out= 8 + 9

With **offered-limited-profile-cir** enabled as the policer **stat-mode**, the SAP offered stats for the policer returned via MIB query and CLI show commands will return the following values:

i. offered-in = 0

ii.'offered-out= 3

iii.'offered-undefined= 1 + 2

iv. offered-managed= 0(IMPM managed packets are not redirected from the policer)

Counter 0 indicates that the SAP policer statistic returns a value of zero.

offered-profile-cir — Counter resource allocation:4

The **offered-profile-cir** mode allocates four forwarding plane offered counters and four traffic manager discard counters.

The **offered-profile-cir** mode is most useful when the policer is receiving trusted out-of-profile and in-profile traffic and is also receiving un-trusted packets that are being applied to a defined CIR profiling rate. This mode differs from **offered-limited-profile-cir** mode in that it expects both trusted in-profile and out-of-profile packets while still performing CIR profiling on packets with un-trusted markings. It is expected that in most cases where both trusted and un-trusted packets are received, the predominate case will not include trusted in-profile packets making the offered-limited-profile-cir accounting mode acceptable.

The counters are used in the following manner:

1. offered-in-that-stayed-green-or-turned-red= profile in

2. offered-undefined-that-turned-green= priority high/low

3. offered-undefined-that-turned-yellow-or-red= priority high/low

4. offered-out-that-stayed-yellow-or-turned-red= profile out

5. dropped-in-that-stayed-green-or-turned-red= Same as 1

6. dropped-undefined-that-turned-green= Same as 2

7. dropped-undefined-that-turned-yellow-or-red= Same as 3

8. dropped-out-that-turned-yellow-or-red= Same as 4

9. forwarded-in-that-stayed-green= Derived from 1 - 5

10. forwarded-undefined-that-turned-green= Derived from 2 - 6

11. forwarded-undefined-that-turned-yellow= Derived from 3 - 7

12. forwarded-out-that-turned-yellow= Derived from 4 - 8

When **collect-stats** is enabled, the counters are used by the system to generate the following statistics:

a. offered-in = 1

b. offered-out= 2 + 3 + 4

c. discard-in = 5 + 6

d. discard-out= 7 + 8

e. forward-in = 9 + 10

f. forward-out= 11 + 12

With **offered-profile-cir** enabled as the policer **stat-mode**, the SAP offered stats for the policer returned via MIB query and CLI show commands will return the following values:

i. offered-high= 1

ii. offered-low= 4

iii. offered-undefined= 2 + 3

iv. offered-managed= 0 (IMPM managed packets are not redirected from the policer)

Counter 0 indicates that the SAP policer statistic returns a value of zero.

offered-priority-cir — Counter resource allocation:4

The **offered-priority-cir** mode allocates four forwarding plane offered counters and four traffic manager discard counters.

The **offered-priority-cir** mode is most useful when the policer is receiving only un-trusted packets that are being classified as high priority or low priority and are being applied to a defined CIR profiling rate. This mode differs from **offered-profile-cir** mode in that it does not expect trusted in-profile and out-of-profile packets but does not exclude the ability of the policer to receive them.

The counters are used in the following manner:

1. offered-high-that-turned-green= profile in, priority high

2. offered-high-that-turned-yellow-or-red= profile in, priority high

3. offered-low-that-turned-green= profile out, priority low

4. offered-low-that-turned-yellow-or-red= profile out, priority low

5. dropped-high-that-turned-green= Same as 1

6. dropped-high-that-turned-yellow-or-red= Same as 2

7. dropped-low-that-turned-green= Same as 3

8. dropped-low-that-turned-yellow-or-red= Same as 4

9. forwarded-high-that-turned-green= Derived from 1 - 5

10. forwarded-high-that-turned-yellow= Derived from 2 - 6

11. forwarded-low-that-turned-green= Derived from 3 - 7

12. forwarded-low-that-turned-yellow= Derived from 4 - 8

When **collect-stats** is enabled, the counters are used by the system to generate the following statistics:

a. offered-high= 1 + 2

b. offered-low= 3 + 4

c. discard-in = 5 + 7

d. discard-out= 6 + 8

e. forward-in = 9 + 11

f. forward-out= 10 + 12

With **offered-priority-cir** enabled as the policer **stat-mode**, the SAP offered stats for the policer returned via MIB query and CLI show commands will return the following values:

i. offered-high= 1 + 2

ii. offered-low= 3 + 4

iii. offered-undefined= 0

iv. offered-managed= 0 (IMPM managed packets are not redirected from the policer)

Counter 0 indicates that the SAP policer statistic returns a value of zero.

offered-total-cir — Counter resource allocation:2

The **offered-total-cir** mode allocates two forwarding plane offered counters and two traffic manager discard counters.

The **offered-total-cir** mode is most useful when the policer is not receiving trusted in-profile or out-of-profile traffic and both high and low priority classifications are not being used on the un-trusted packets and the offered packets are being applied to a defined CIR profiling rate. This mode does not prevent the policer from receiving trusted in-profile or out-of-profile packets and does not prevent the use of priority high or low classifications on the un-trusted packets.

The counters are used in the following manner:

1. offered-that-turned-green= profile in/out, priority high/low

2. offered- that-turned-yellow-or-red= profile in/out, priority high/low

3. dropped-offered-that-turned-green= Same as 1

4. dropped-offered-that-turned-yellow-or-red= Same as 2

5. forwarded-offered-that-turned-green= Derived from 1 - 3

6. forwarded-offered-that-turned-yellow= Derived from 2 - 4

When **collect-stats** is enabled, the counters are used by the system to generate the following statistics:

a. offered-in = 1 + 2

b. offered-out= 0

c. discard-in = 3

d. discard-out= 4

e. forward-in = 5

f. forward-out= 6

Counter 0 indicates that the accounting statistic returns a value of zero.

With **offered-total-cir** enabled as the policer **stat-mode**, the SAP offered stats for the policer returned via MIB query and CLI show commands will return the following values:

i. offered-high= 1 + 2

ii. offered-low= 0

iii. offered-undefined= 0

iv. offered-managed= 0 (IMPM managed packets are not redirected from the policer)

Counter 0 indicates that the SAP policer statistic returns a value of zero.

Explicit Subscriber Mapping Commands

explicit-sub-map

Syntax **explicit-sub-map**

Context config>subscr-mgmt

Description This command configures an explicit subscriber mapping

entry

Syntax **entry key** *sub-ident-string* [**sub-profile** *sub-profile-name*] [**alias** *sub-alias-strin*g] [**sla-profile** *sla-profile-name*]
no entry key *sub-profile-string*

Context config>subscr-mgmt>explicit-sub-map

Description This command configures a subscriber identification string.

Parameters *sub-ident-string* — Specifies the profile string.

Values 16 characters maximum

sub-profile-name — Specifies an existing subscriber profile name.

Values 32 characters maximum

alias *sub-alias-string* **—** Specifies an alias for the subscriber identification string.

sla-profile *sla-profile-name* **—** Specifies an existing SLA profile.

IGMP Policy commands

igmp-policy

Syntax **igmp-policy** *policy-name* [**create**]
no igmp-policy

Context config>sub-mgmt

Description This command configures an IGMP policy.

Parameters *policy-name* — Specifies the policy name.

Values 32 chars max

egress-rate-modify

Syntax **egress-rate-modify** [**egress-rate-limit | scheduler** *scheduler-name*]
no egress-rate-modify

Context configure>subscr-mgmt>igmp-policy

Description This command is used to apply HQoS Adjustment to a subscriber. HQoS Adjustment is needed when multicast traffic flow for the subscriber is dissociated from subscriber host queues. Multicast redirection is typical such case although it can be applied in direct IPoE subscriber per-sap replication mode.

The channel bandwidth definition policy is defined in the mcac policy under the configure>router>mcac>policy hierarchy. The policy is applied under the redirected interface or under the group-interface.

In order for HQoS Adjustment to take effect, sub-mcac-policy must be in a no shutdown mode and applied under the sub-profile even if mcac is not deployed.

Parameters **egress-rate-limit —** Subscriber's bandwidth is capped via the aggregate-rate-limit command in the sub-profile or via a Change of Authorization (CoA) request. This bandwidth cap will be dynamically adjusted according to the multicast channel definition and channel association with the host via IGMP.

scheduler *scheduler-name* — Subscriber's bandwidth is capped via the scheduling-policy in the sub-profile or via a Change of Authorization (CoA) request . HQoS Adjustment will modify the rate of the scheduler (scheduler-name) defined in the scheduling-policy or configured via CoA.

Default HQoS Adjustment is disabled.

import

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

Context config>sub-mgmt>igmp-policy

Description This command specifies the import policy to filter IGMP packets.

Parameters *policy-name* — Specifies the policy name.

Values 32 chars max

max-num-groups

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

Context config>sub-mgmt>igmp-policy

Description This command configures the max number of multicast groups.

Parameters *max-num-groups* — Specifies the maximum number of multicast groups.

Values 0 — 16000

max-num-sources

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

Context config>sub-mgmt>igmp-policy

Description This command configures the max number of multicast sources.

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

Default no max-num-sources

Parameters *max-num-sources* —

Values 1 — 1000

max-num-grp-sources

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

Context config>sub-mgmt>igmp-policy
config>sub-mgmt>msap-policy>igmp-host-tracking

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

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

Default no max-num-grp-sources

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

mcast-reporting

Syntax [**no**] **mcast-reporting**

Context config>sub-mgmt>igmp-policy

Description This command configures mcast reporting.

mcast-reporting-dest

Syntax **mcast-reporting-dest** *dest-name*
no mcast-reporting-dest

Context configure>subscriber-mgmt>igmp-policy>mcast-reporting>
configure>subscriber-mgmt>host-tracking-policy>mcast-reporting>

Description This command references Multicast Reporting Destination to which IGMP related events are exported.

The Multicast Reporting Destination is referenced with the subscriber itself or within the Host-Tracking-Policy.

Parameters *dest-name* — Name of the Multicast Reporting Destination.

Default no mcast-reporting-dest is referenced.

opt-reporting-fields

Syntax **opt-reporting-fields** [**host-mac**] [**pppoe-session-id**] [**svc-id**] [**sap-id**]
no opt-reporting-fields

Context configure>subscriber-mgmt>igmp-policy>mcast-reporting>
configure>subscriber-mgmt>host-tracking-policy>mcast-reporting>

Description This command will specify optional data relevant to the IGMP event that can be exported. This optional data includes:

- Host MAC address

- PPPoE session-ID
- Service ID
- SAP

Parameters **host-mac —** Specifies the host-mac optional field should be included into the multicast reporting messages.

pppoe-session-id — Specifies the pppoe-session-id optional field should be included into the multicast reporting messages.

svc-id — Specifies the svc-id optional field should be included into the multicast reporting messages.

sap-id — Specifies the sap-id optional field should be included into the multicast reporting messages.

Default Optional data is disabled.

Sample Output

```
configure
     system
          security
               source-address
                    application <app> <ip-int-name | ip-address>

<app>                 : cflowd|dns|ftp|ntp|ping|radius|snmptrap|sntp|ssh|
                        syslog|tacplus|telnet|traceroute|mcreporter
 <ip-int-name|ip-ad*> : ip-int-name    - 32 chars max
                        ip-address     - a.b.c.d
```

sub-mcac-policy

Syntax **sub-mcac-policy** *policy-name*
no sub-mcac-policy

Context configure>subscr-mgmt

Description This command will create a policy template with mcac bandwidth limits that will be applied to the subscriber.

Per interface mcac bandwidth limits will be set directly under the interface (regular interface or group-interface) and no such policy templates are needed.

The need for a separate policy template for subscribers is due to the fact that sub-groups of subscribers under the group-interface can share certain settings that can be configured via templates.

To summarize, the mcac bandwidth constraints for subscribers are defined in the sub-mcac-policy while the mcac bandwidth constraints for the interface are configured directly under the **igmp>interface>mcac** or **igmp>group-interface>mcac** context without the need for policy templates.

Note that the sub-mcac-policy only deals with the mcac bandwidth limits and not the channel bandwidth definitions. Channels bandwidth is defined in a different policy (under the configure>router>mcac hierarchy) and that policy is applied on the interface level as follows:

In case of HQoS Adjustment, it is mandatory that the sub-mcac-policy be created and applied to the subscriber. The sub-mac-policy does not have to contain any bandwidth constrains, but it has to be in a no shutdown state in order for HQoS Adjustment to work.

Parameters *policy-name* — Name of the policy.

Default No sub-mcac-policy is created.

sub-mcac-policy

Syntax **sub-mcac-policy** *policy-name*
no sub-mcac-policy

Context configure>subscr-mgmt>sub-profile

Description This command references the policy template in which the mcac bandwidth limits are defined. Mcac for the subscriber is effectively enabled with this command when the sub-profile is applied to the subscriber. The bandwidth of the channels is defined in a different policy (under the configure>router>mcac hierarchy) and this policy is applied on the interface level as follows:

for regular interfacs under the configure>service/router>igmp>interface>mcac hierarchy

In case of HQoS Adjustment, it is mandatory that the sub-mcac-policy be created and applied to the subscriber. The sub-mac-policy does not have to contain any bandwidth constrains, but it has to be in a no shutdown state in order for HQoS Adjustment to work.

Parameters *policy-name* — Name of the policy.

Default No policy is referenced.

version

Syntax **version** *version*
no version

Context config>sub-mgmt>igmp-policy

Description This command configures the version of IGMP.

Parameters *version* — Specifies the version of IGMP.

Values 1, 2 or 3

fast-leave

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

Context config>sub-mgmt>igmp-policy

Description This command enables/disables IGMP fast-leave processing.

Default fast-leave

static

Syntax **static**

Context config>sub-mgmt>igmp-policy

Description This command adds or removes IGMP static group membership.

per-host-replication

Syntax [**no**] **per-host-replication**

Context configure>subscr-mgmt>igmp-policy

Description This command enables per-host-replication in IPoE model. For PPPoX, per-host-replication is the only mode of operation. In the per-host-replication mode, multicast traffic is replicated per each host within the subscriber irrespective of the fact that some hosts may be subscribed to the same multicast stream. As a result, in case that multiple hosts within the subscriber are registered for the same multicast group, the multicast streams of that group will be generated. The destination MAC address of multicast streams will be changed to unicast so that each host receives its own copy of the stream. Multicast traffic in the per-host-replication mode can be classified via the existing QoS CLI structure. As such the multicast traffic will flow through the subscriber queues. HQoS Adjustment is not needed in this case.

The alternative behavior for multicast replication in IPoE environment is per-SAP- replication. In this model, only a single copy of the multicast stream is sent per SAP, irrespective of the number of hosts that are subscribed to the same multicast group. This behavior applies to 1:1 connectivity model as well as on 1:N connectivity model (SAP centric behavior as opposed to subscriber centric behavior).

In the per-SAP-replication model the destination MAC address is multicast (as opposed to unicast in the per-host-replication model). Multicast traffic is flowing via the SAP queue which is outside of the subscriber context. The consequence is that multicast traffic is not accounted in the subscriber HQoS. In addition, HQoS Adaptation is not supported in the per SAP replication model.

Default disabled

redirection-policy

Syntax **redirection-policy** *policy-name*
no redirection-policy

Context config>sub-mgmt>igmp-policy

Description This command will apply multicast redirection action to the subscriber. The redirection action along with the redirected interface (and possibly service id) is defined in the referenced policy-name. IGMP messages will be redirected to an alternate interface if that alternate interface has IGMP enabled. The alternate interface does not have to have any multicast groups registered via IGMP. Currently all

IGMP messages are redirected and there is no ability to selectively redirect IGMP messages based on match conditions (multicast-group address, source IP address, etc.). Multicast redirection is supported between VPRN services and also between interfaces within the Global Routing Context. Multicast Redirection is not supported between the VRPN services and the Global Routing Table (GRT).

IGMP state is maintained per subscriber host and per redirected interface. Traffic is however forwarded only on the redirected interface.

Default none

Parameters *policy-name* — This is a regular policy defined under the **configure>router>policy-option>policy-statement** context.

group

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

Context config>sub-mgmt>igmp-policy>static

Description This command adds or removes a static multicast group.

Parameters *ip-address* — Specifies the IP address.

Values a.b.c.d

Subscriber Management Service Commands

SAP Subscriber Management Commands

sub-sla-mgmt

Syntax [**no**] **sub-sla-mgmt**

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

Description This command enables the context to configure subscriber management parameters for this SAP.

Default no sub-sla-mgmt

host-connectivity-verify

Syntax **host-connectivity-verify** [**interval** *interval*] [**action** {**remove** | **alarm**}]

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

Description This command enables subscriber host connectivity verification on a given SAP within a VPLS or IES service. This tool will periodically scan all known hosts (from dhcp-state) and perform UC ARP requests. The subscriber host connectivity verification will maintain state (connected vs. not-connected) for all hosts.

Default no host-connectivity-verify

Parameters **interval *interval*** — The interval, in minutes, which specifies the time interval which all known sources should be verified. The actual rate is then dependent on the number of known hosts and interval.

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

action {**remove** | **alarm**} — Defines the action taken on a subscriber host connectivity verification failure for a given host. The **remove** keyword raises an alarm and removes dhcp-state and releases all allocated resources (queues, table entries and etc.). DHCP-RELEASE will be signaled to corresponding DHCP server. Static hosts will be never removed. The **alarm** keyword raises an alarm indicating that the host is disconnected.

def-sla-profile

Syntax **def-sla-profile** *default-sla-profile-name*
no def-sla-profile

Context config>service>vpls>sap>sub-sla-mgmt
config>service>ies>if>sap>sub-sla-mgmt
config>service>ies>sub-if>grp-if>sap>sub-sla-mgmt

Description This command specifies a default SLA profile for this SAP.

An SLA profile is a named group of QoS parameters used to define per service QoS for all subscriber hosts common to the same subscriber within a provider service offering. A single SLA profile may define the QoS parameters for multiple subscriber hosts.

The **no** form of the command removes the default SLA profile from the SAP configuration.

Default no def-sla-profile

Parameters *default-sla-profile-name* — Specifies a default SLA profile for this SAP.

def-sub-profile

Syntax **def-sub-profile** *default-subscriber-profile-name*

Context config>service>vpls>sap>sub-sla-mgmt
config>service>ies>if>sap>sub-sla-mgmt
config>service>ies>sub-if>grp-if>sap>sub-sla-mgmt

Description This command specifies a default subscriber profile for this SAP.

A subscriber profile defines the aggregate QoS for all hosts within a subscriber context. This is done through the definition of the egress and ingress scheduler policies that govern the aggregate SLA for subscriber using the subscriber profile.

The **no** form of the command removes the default SLA profile from the SAP configuration.

Parameters *default-sub-profile* — Specifies a default subscriber profile for this SAP.

sub-ident-policy

Syntax **sub-ident-policy** *sub-ident-policy-name*

Context config>service>vpls>sap>sub-sla-mgmt
config>service>ies>if>sap>sub-sla-mgmt
config>service>ies>sub-if>grp-if>sap>sub-sla-mgmt

Description This command associates a subscriber identification policy to this SAP.

Subscribers are managed by the system through the use of subscriber identification strings. A subscriber identification string uniquely identifies a subscriber. For static hosts, the subscriber identification string is explicitly defined with each static subscriber host.

For dynamic hosts, the subscriber identification string must be derived from the DHCP ACK message sent to the subscriber host. The default value for the string is the content of Option 82 CIRCUIT-ID and REMOTE-ID fields interpreted as an octet sting. As an option, the DHCP ACK message may be processed by a subscriber identification policy which has the capability to parse the message into an alternative ASCII or octet string value.

When multiple hosts on the same port are associated with the same subscriber identification string they are considered to be host members of the same subscriber.

The **no** form of the command removes the default subscriber identification policy from the SAP configuration.

Default no sub-ident-policy

Parameters *sub-ident-policy-name* — Specifies a subscriber identification policy for this SAP.

multi-sub-sap

Syntax **multi-sub-sap** *number-of-sub*
no multi-sub-sap

Context config>service>vpls>sap>sub-sla-mgmt
config>service>ies>if>sap>sub-sla-mgmt
config>service>ies>sub-if>grp-if>sap>sub-sla-mgmt

Description This command defines the maximum number of subscribers (dynamic + static) that can be simultaneously active on this SAP.

If the limit is reached, a new host will be denied access and the corresponding DHCP ACK will be dropped.

Default 1

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

Default no multi-sub-sap

Parameters *multi-sub-sap* — Specifies the maximum allowed.

single-sub-parameters

Syntax **single-sub-parameters**

Context config>service>vpls>sap>sub-sla-mgmt
config>service>ies>sub-if>grp-if>sap>sub-sla-mgmt
config>service>ies>if>sap>sub-sla-mgmt

Description This command configure single subscriber SAP parameters.

non-sub-traffic

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

Context config>service>vpls>sap>sub-sla-mgmt>single-sub
config>service>ies>if>sap>sub-sla-mgmt>single-sub
config>service>ies>sub-if>grp-if>sap>sub-sla-mgmt>single-sub

Description This command configures traffic profiles for non-IP traffic such as PPPoE.It is used in conjunction with the profiled-traffic-only on single subscriber SAPs and creates a subscriber host which is used to forward non-IP traffic through the single subscriber SAP without the need for SAP queues.

The **no** form of the command removes any configured profile.

Default no non-sub-traffic

Parameters *sub-profile-name* — Identifies the subscriber profile name.

Values 32 characters maximum

sla-profile-name — Identifies the SLA profile name.

Values 32 characters maximum

profiled-traffic-only

Syntax [**no**] **profiled-traffic-only**

Context config>service>vpls>sap>sub-sla-mgmt>single-sub-parameters
config>service>ies>if>sap>sub-sla-mgmt>single-sub
config>service>ies>sub-if>grp-if>sap>sub-sla-mgmt>single-sub

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

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

Default no profiled-traffic-only

srrp

Syntax [**no**] **srrp** *srrp-id*

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

Description This command creates an SRRP instance on a group IP interface. An SRRP instance manages all subscriber subnets within the group interfaces subscriber IP interface or other subscriber IP interfaces

that are associated through a wholesale/retail relationship. Only one unique SRRP instance can be configured per group interface.

The **no** form of the command removes an SRRP instance from a group IP interface. Once removed, the group interface ignores ARP requests for the SRRP gateway IP addresses that may exist on subscriber subnets associated with the group IP interface. Then the group interface stops routing using the redundant IP interface associated with the group IP interface and will stop routing with the SRRP gateway MAC address. Ingress packets destined to the SRRP gateway MAC will also be silently discarded. This is the same behavior as a group IP interface that is disabled (shutdown).

Default no srrp

Parameters *srrp-id* — Specifies a 32 bit instance ID that must be unique to the system. The instance ID must also match the instance ID used by the remote router that is participating in the same SRRP context. SRRP is intended to perform a function similar to VRRP where adjacent IP hosts within local subnets use a default gateway to access IP hosts on other subnets.

Values 1 — 4294967295

gw-mac

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

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

Description This command overrides the default SRRP gateway MAC address used by the SRRP instance. Unless specified, the system uses the same base MAC address for all SRRP instances with the last octet overridden by the lower 8 bits of the SRRP instance ID. The same SRRP gateway MAC address should be in-use by both the local and remote routers participating in the same SRRP context.

One reason to change the default SRRP gateway MAC address is if two SRRP instances sharing the same broadcast domain are using the same SRRP gateway MAC. The system will use the SRRP instance ID to separate the SRRP messages (by ignoring the messages that does not match the local instance ID), but a unique SRRP gateway MAC is essential to separate the routed packets for each gateway IP address.

The **no** form of the command removes the explicit SRRP gateway MAC address from the SRRP instance. The SRRP gateway MAC address can only be changed or removed when the SRRP instance is shutdown.

Parameters *mac-address* — Specifies a MAC address that is used to override the default SRRP base MAC address

Values Any MAC address except all zeros, broadcast or multicast addresses. The offset is expressed in normal Ethernet MAC address notation. The defined gw-mac cannot be 00:00:00:00:00:00, ff:ff:ff:ff:ff:ff or any multicast address.

If not specified, the system uses the default SRRP gateway MAC address with the last octet set to the 8 least significant bits of the SRRP instance ID.

keep-alive-interval

Syntax **keep-alive-interval** *interval*
no keep-alive-interval

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

Description This command defines the interval between SRRP advertisement messages sent when operating in the master state. The interval is also the basis for setting the master-down timer used to determine when the master is no longer sending. The system uses three times the keep-alive interval to set the timer. Every time an SRRP advertisement is seen that is better then the local priority, the timer is reset. If the timer expires, the SRRP instance assumes that a master does not exist and initiates the attempt to become master.

When in backup state, the SRRP instance takes the keep-alive interval of the master as represented in the masters SRRP advertisement message. Once in master state, the SRRP instance uses its own configured keep-alive interval.

The keep-alive-interval may be changed at anytime, but will have no effect until the SRRP instance is in the master state.

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

Parameters *interval* — Specifies the interval, in milliseconds, between SRRP advertisement messages sent when operating in the master state.

Values 1 — 100

Default 10 milliseconds

message-path

Syntax **message-path** *sap-id*
no message-path

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

Description This command defines a specific SAP for SRRP in-band messaging. A message-path SAP must be defined prior to activating the SRRP instance. The defined SAP must exist on the SRRP instances group IP interface for the command to succeed and cannot currently be associated with any dynamic or static subscriber hosts. Once a group IP interface SAP has been defined as the transmission path for SRRP Advertisement messages, it cannot be administratively shutdown, will not support static or dynamic subscriber hosts and cannot be removed from the group IP interface.

The SRRP instance message-path command may be executed at anytime on the SRRP instance. Changing the message SAP will fail if a dynamic or static subscriber host is associated with the new SAP. Once successfully changed, the SRRP instance will immediately disable anti-spoof on the SAP and start sending SRRP Advertisement messages if the SRRP instance is activated.

Changing the current SRRP message SAP on an active pair of routers should be done in the following manner:

1. Shutdown the backup SRRP instance.
2. Change the message SAP on the shutdown node.

3. Change the message SAP on the active master node.

4. Re-activate the shutdown SRRP instance.

Shutting down the backup SRRP instance prevents the SRRP instances from becoming master due to temporarily using differing message path SAPs.

If an MCS peering is operational between the redundant nodes and the SRRP instance has been associated with the peering, the designated message path SAP will be sent from each member.

The **no** form of the command can only be executed when the SRRP instance is shutdown. Executing no message-path allows the existing SAP to be used for subscriber management functions. A new message-path SAP must be defined prior to activating the SRRP instance.

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition. See Common Service Commands on page 1740 for *sap-id* command syntax.

policy

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

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

Description This command associates one or more VRRP policies with the SRRP instance. A VRRP policy is a collection of connectivity and verification tests used to manipulate the in-use priorities of VRRP and SRRP instances. A VRRP policy can test the link state of ports, ping IP hosts, discover the existence of routes in the routing table or the ability to reach L2 hosts. When one or more of these tests fail, the VRRP policy has the option of decrementing or setting an explicit value for the in-use priority of an SRRP instance.

More than one VRRP policy may be associated with an SRRP instance. When more than one VRRP policy is associated with an SRRP instance the delta decrement of the in-use priority is cumulative unless one or more test fail that have explicit priority values. When one or more explicit tests fail, the lowest priority value event takes effect for the SRRP instance. When the highest delta-in-use-limit is used to manage the lowest delta derived in-use priority for the SRRP instance.

VRRP policy associations may be added and removed at anytime. A maximum of two VRRP policies can be associated with a single SRRP instance.

The **no** form of the command removes the association with vrrp-policy-id from the SRRP instance.

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

Values 1 — 9999

priority

Syntax **priority** *priority*
no priority

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

Description This command overrides the default base priority for the SRRP instance. The SRRP instance priority is advertised by the SRRP instance to its neighbor router and is compared to the priority received

from the neighbor router. The router with the best (highest) priority enters the master state while the other router enters the backup state. If the priority of each router is the same, the router with the lowest source IP address in the SRRP advertisement message assumes the master state.

The base priority of an SRRP instance can be managed by VRRP policies. A VRRP policy defines a set of connectivity or verification tests which, when they fail, may lower an SRRP instances base priority (creating an in-use priority for the instance). Every time an SRRP instances in-use priority changes when in master state, it sends an SRRP advertisement message with the new priority. If the dynamic priority drops to zero or receives an SRRP Advertisement message with a better priority, the SRRP instance transitions to the *becoming backup* state.

When the priority command is not specified, or the no priority command is executed, the system uses a default base priority of 100. The priority command may be executed at anytime.

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

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

Values 1 — 254

Default 100

tos-marking-state

Syntax **tos-marking-state** {**trusted** | **untrusted**}
no tos-marking-state

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

Description This command is used to alter the default trusted state to a non-trusted state. When unset or reverted to the trusted default, the ToS field will not be remarked by egress network IP interfaces unless the egress network IP interface has the remark-trusted state set, in which case the egress network interface treats all VPRN and network IP interface as untrusted.

When the ingress interface is set to untrusted, all egress network IP interfaces will remark IP packets received on the network interface according to the egress marking definitions on each network interface. The egress network remarking rules also apply to the ToS field of IP packets routed using IGP shortcuts (tunneled to a remote next-hop). However, the tunnel QoS markings are always derived from the egress network QoS definitions.
Egress marking and remarking is based on the internal forwarding class and profile state of the packet once it reaches the egress interface. The forwarding class is derived from ingress classification functions. The profile of a packet is either derived from ingress classification or ingress policing.

The default marking state for network IP interfaces is trusted. This is equivalent to declaring no tos-marking-state on the network IP interface. When undefined or set to tos-marking-state trusted, the trusted state of the interface will not be displayed when using show config or show info unless the detail parameter is given. The **save config** command will not store the default tos-marking-state trusted state for network IP interfaces unless the detail parameter is also specified.

The **no** tos-marking-state command is used to restore the trusted state to a network IP interface. This is equivalent to executing the tos-marking-state trusted command.

Default trusted

Parameters **trusted —** The default prevents the ToS field to not be remarked by egress network IP interfaces unless the egress network IP interface has the remark-trusted state set.

untrusted — Specifies that all egress network IP interfaces will remark IP packets received on the network interface according to the egress marking definitions on each network interface.

mac-da-hashing

Syntax **mac-da-hashing**
no mac-da-hashing

Context config>service>vpls>sap>sub-sla-mgmt

Description This command specifies whether subscriber traffic egressing a LAG SAP has its egress LAG link selected by a function of the MAC destination address instead of the subscriber ID.

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

Default no mac-da-hashing

host

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

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

Description This command creates a static subscriber host for the SAP. Static subscriber hosts may be used by the system for various purposes. Applications within the system that make use of static host entries include anti-spoof, ARP reply agent and source MAC population into the VPLS forwarding database.

Multiple static hosts may be defined on the SAP. Each host is identified by either a source IP address, a source MAC address or both a source IP and source MAC address. Every static host definition must have at least one address defined, IP or MAC.

Static hosts can exist on the SAP even with anti-spoof and ARP reply agent features disabled. When enabled, each feature has different requirements for static hosts.

Use the **no** form of the command to remove a static entry from the system. The specified *ip-address* and *mac-address* must match the host's exact IP and MAC addresses as defined when it was created. When a static host is removed from the SAP, the corresponding anti-spoof filter entry and/or FDB entry is also removed.

Default none

Parameters **ip** *ip-address* — Specify this parameter to associate a subscriber with the static subscriber host. Only one static host can be configured on the SAP with a given IP address.

mac *mac-address* — Specify this optional parameter when defining a static host. The MAC address must be specified for **anti-spoof mac** and **anti-spoof ip-mac**. Multiple static hosts may be configured with the same MAC address given that each definition is distinguished by a unique IP address.

Every static host definition must have at least one address defined, IP or MAC.

subscriber *sub-ident-string* — Specify this parameter to configure an existing subscriber identification profile to be associated with the static subscriber host. The subscriber identification profile is configured in the **config>subscr-mgmt>sub-ident-policy** context. The subscriber information is used by the VPLS SAP arp-reply-agent to determine the proper handling of received ARP requests from subscribers.

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

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

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

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

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

sub-profile *sub-profile-name* — Specify this parameter to configure an existing subscriber profile name to be associated with the static subscriber host. The subscriber profile is configured in the **config>subscr-mgmt>sub-profile** context.

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

Note that if Enhanced Subscriber Management is enabled on a SAP using the **sub-sla-mgmt** command, the **sub-ident**, **sub-profile,** and **sla-profile** must be configured for all static hosts defined on this SAP.

Managed SAP Policy Commands

msap-policy

Syntax **msap-policy** *msap-policy-name* [**create**]
no msap-policy *msap-policy-name*

Context config>subscr-mgmt

Description This command configures a managed SAP policy. Managed SAPs allow the use of policies and a SAP template for the creation of a SAP.

Default none

Parameters *msap-policy-name* — Specifies the managed SAP policy name.

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

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

cpu-protection

Syntax **cpu-protection**

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

Description This command enables the context to configure CPU protection policies.

cpu-protection

Syntax **cpu-protection** *policy-id* [**mac-monitoring**]
no cpu-protection

Context config>subscr-mgmt>msap-policy [mac-monitoring]
config>service>ies>sub-if>grp-if>sap [mac-monitoring]
config>service>vpls>sap [mac-monitoring]
config>service>vprn>sub-if>grp-if>sap [mac-monitoring]

Description This command assigns an existing CPU protection policy to the SAP or interface.

CPU protection policies are configured in the **config>sys>security>cpu-protection** context.

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

Default none

Parameters *policy-id* — Specifies an existing CPU protection policy to assign to the SAP or interface.

mac-monitoring — Specifies that the per-source rate limit be applied.

cpu-protection

Syntax **cpu-protection** *policy-id*
no cpu-protection

Context config>router>if
config>service>ies>if
config>service>vprn>if

Description This command assigns an existing CPU protection policy to the SAP or interface.

CPU protection policies are configured in the **config>sys>security>cpu-protection** context.

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

Default none

Parameters *policy-id* — Specifies an existing CPU protection policy to assign to the SAP.

default-host

Syntax **default-host** *ip-address*/*mask* **next-hop** *next-hop-ip*
no default-host *ip-address*/*mask*

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

Description This command configures the default-host to be used. More than one default-host can be configured per SAP.

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

Parameters *ip-address/mask* — Assigns an IP address/IP subnet format to the interface.

next-hop *next-hop-ip* **—** Assigns the next hop IP address.

dist-cpu-protection

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

Context config>subscriber-management>msap-policy

Description This command assigns a Distributed CPU Protection (DCP) policy to the MSAP policy. The DCP policy will automatically get assigned to any MSAPs created with this policy. A non-existant DCP policy can be assigned to an msap-policy since an msap-policy is effectively a template that gets applied at some point in the future during msap creation. The existence of the DCP policy will be validated at the time that the msap is created, and the msap creation will be blocked (and an

appropriate log event created) if the DCP policy does not exist. Note that for other types of objects (for example, normal non-msap SAPs and network interfaces) the DCP policy must exist before it can be assigned to the SAP.

Default. no dist-cpu-protection

ies-vprn-only-sap-parameters

Syntax **ies-vprn-only-sap-parameters**

Context config>subscr-mgmt>msap-policy

Description This command configures Managed SAP IES and VPRN properties.

igmp-host-tracking

Syntax **igmp-host-tracking**

Context config>subscr-mgmt>msap-policy

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

expiry-time

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

Context config>subscr-mgmt>msap-policy>igmp-host-tracking

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

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

Default no expiry-time

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

Values 1 — 65535

import

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

Context config>subscr-mgmt>msap-policy>igmp-host-tracking

Description This command specifies the import routing policy to be used for IGMP packets to be used on this SAP or SDP. Only a single policy can be imported on a single SAP at any time.

The **no** form of the command removes the policy association from the SAP or SDP.

Default no import (No import policy is specified)

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

max-num-group

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

Context config>subscr-mgmt>msap-policy>igmp-host-tracking

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

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

Default no max-num-groups

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

Values 1 — 196607

max-num-sources

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

Context config>subscr-mgmt>msap-policy>igmp-host-tracking

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

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

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

Values 1 — 1000

max-num-grp-sources

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

Context config>subscr-mgmt>msap-policy>igmp-host-tracking

Description This command configures the maximum number of group sources for which IGMP can have local receiver information based on received IGMP reports on this interface. When this configuration is

changed dynamically to a value lower than currently accepted number of group sources, the group sources that are already accepted are not deleted. Only new group sources will not be allowed. When this object has a value of 0, there is no limit to the number of group sources.

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

Default no max-num-grp-sources

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

lag-link-map-profile

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

Context config>subscr-mgmt>msap-policy

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

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

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

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

sub-sla-mgmt

Syntax [**no**] **sub-sla-mgmt**

Context config>subscr-mgmt>msap-policy
config>service>vprn>sub-if>grp-if>sap>sub-sla-mgmt
config>service>vpls>sap>sub-sla-mgmt

Description This command enables the context to configure subscriber management parameters for an MSAP.

Default no sub-sla-mgmt

def-app-profile

Syntax **def-app-profile** *app-profile-name*
no def-app-profile

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

Description This command specifies the application profile to be used by a subscriber host.

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

Default no def-app-profile

Parameters *app-profile-name* — specifies an existing application profile to be mapped to the subscriber profile by default.

def-inter-dest-id

Syntax **def-inter-dest-id** {**string** *string* | **use-top-q**}
no def-inter-dest-id

Context config>subscr-mgmt>msap-policy>sub-sla-mgmt

Description This command specifies a default destination string for all subscribers associated with the SAP. The command also accepts the **use-top-q** flag that automatically derives the string based on the top most delineating Dot1Q tag from the SAP's encapsulation.

The **no** form of the command removes the default subscriber identification string from the configuration.

no def-sub-id

Default no def-inter-dest-id

Parameters **use-top-q —** Derives the string based on the top most delineating Dot1Q tag from the SAP's encapsulation.

string *string* **—** Specifies the subscriber identification applicable for a subscriber host.

def-sub-id

Syntax **def-sub-id use-auto-id**
def-sub-id use-sap-id
def-sub-id string *sub-id*
no def-sub-id

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

Description This command specifies the explicit default sub-id for dynamic subscriber hosts (including ARP hosts) in case that the sub-id string is NOT supplied through RADIUS or LUDB.

The sub-id is assigned to a new subscriber host in the following order of priority:

- RADIUS
- LUDB

- Explicit default – with the def-sub-id command we explicitly set the sub-id name of the host to be one of the following:
  - → The sap-id to which the new host is associated with
  - → Explicit string
  - → Auto-generated string consisting of the concatenated subscriber identification fields defined under the **subscr-mgmt>auto-sub-id-key** node. The fields are taken in the order in which they are configured and are separated by a '|'character. The subscriber host identification fields are separately defined for IPoE and PPPoE host types.
- Implicit default – in case that the sub-id string is not returned via RADIUS or LUDB and there is no def-sub-id configured, the sub-id name will be generated as a random 10 character encoded string based on the auto-sub-id-keys. This 10 characters encoded string will be unique per chassis as well as in dual-homed environment. It is generated based on auto-sub-id-keys. If auto-sub-id-keys are not explicitly configured, the default ones are:
  - → <mac, sap-id, session-id> for PPP type hosts
  - → <mac, sap-id> for IPoE type hosts.

This command does not apply to static subscribers.

Parameters **use-sap-id** — Specifies the sub-id name -id on which the original request for host creation arrived (DHCP Discover, or PADI or ARP Request).

string *sub-id* — Explicitly configured sub-id name.

use-auto-id — The sub-id name is the concatenated string of auto-sub-id-keys separated by a "|" character.

Default no def-sub-id

Implicit default – If the sub-id string is not supplied through RADIUS, LUDB orby configuration (def-sub-id), then a random 10 character encoded sub-id name will be generated. This random sub-id name will be based on the subscriber identification keys defined under the subscr-mgmt>auto-sub-id-key node. In case that the auto-sub-id-keys are not defined explicitly, the default ones are:

- <mac, sap-id, session-id>for PPPoE type hosts
- <mac, sap-id>for IPoE type hosts

def-sla-profile

Syntax **def-sla-profile** *default-sla-profile-name*
no def-sla-profile

Context config>subscr-mgmt>msap-policy>sub-sla-mgmt

Description This command specifies a default SLA profile for an MSAP.

An SLA profile is a named group of QoS parameters used to define per service QoS for all subscriber hosts common to the same subscriber within a provider service offering. A single SLA profile may define the QoS parameters for multiple subscriber hosts.

The **no** form of the command removes the default SLA profile from the MSAP configuration.

Default no def-sla-profile

Parameters *default-sla-profile-name* — Specifies a default SLA profile for an MSAP.

def-sub-profile

Syntax **def-sub-profile** *default-subscriber-profile-name*

Context config>subscr-mgmt>msap-policy>sub-sla-mgmt

Description This command specifies a default subscriber profile for an MSAP.

A subscriber profile defines the aggregate QoS for all hosts within a subscriber context. This is done through the definition of the egress and ingress scheduler policies that govern the aggregate SLA for subscriber using the subscriber profile.

The **no** form of the command removes the default SLA profile from the SAP configuration.

Parameters *default-sub-profile* — Specifies a default subscriber profile for this SAP.

multi-sub-sap

Syntax **multi-sub-sap** [**limit** *limit*]
no multi-sub-sap

Context config>subscr-mgmt>msap-policy>sub-sla-mgmt

Description This command defines the maximum number of subscribers (dynamic + static) that can be simultaneously active on an MSAP.

If the limit is reached, a new host will be denied access and the corresponding DHCP ACK will be dropped.

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

Default 1

Parameters **limit** *limit* — Specifies the maximum allowed.

Values 2 — 8000

single-sub-parameters

Syntax **single-sub-parameters**

Context config>subscr-mgmt>msap-policy>sub-sla-mgmt

Description This command enables the context to configure single subscriber MSAP parameters.

non-sub-traffic

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

Context config>subscr-mgmt>msap-policy>sub-sla-mgmt>single-sub

Description This command configures traffic profiles for non-IP traffic such as PPPoE.It is used in conjunction with the profiled-traffic-only on single subscriber SAPs and creates a subscriber host which is used to forward non-IP traffic through the single subscriber SAP without the need for SAP queues.

The **no** form of the command removes any configured profile.

Default no non-sub-traffic

Parameters *sub-profile-name* — Identifies the subscriber profile name.

Values 32 characters maximum

sla-profile-name — Identifies the SLA profile name.

Values 32 characters maximum

profiled-traffic-only

Syntax [**no**] **profiled-traffic-only**

Context config>subscr-mgmt>msap-policy>sub-sla-mgmt>single-sub

Description This command specifies whether only profiled traffic is applicable for an MSAP. When enabled, all queues will be deleted.

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

Default no profiled-traffic-only

sub-ident-policy

Syntax [**no**] **sub-ident-policy** *sub-ident-policy-name*

Context config>subscr-mgmt>msap-policy>sub-sla-mgmt

Description This command specifies an existing subscriber identification policy. Each subscriber identification policy can have a default subscriber profile defined. The subscriber identification policy default subscriber profile overrides the system default and the subscriber SAP default subscriber profiles. Defining a subscriber identification policy default subscriber profile is optional.

Defining a subscriber profile as a subscriber identification policy default subscriber profile will cause all active subscribers currently associated with a subscriber SAP using the policy and associated with a subscriber policy through the system default or subscriber SAP default subscriber profiles to be reassigned to the subscriber policy defined as default on the subscriber identification policy.

Attempting to delete a subscriber profile that is currently defined as a default for a subscriber identification policy will fail.

When attempting to remove a subscriber identification policy default subscriber profile definition, the system will evaluate each active subscriber on all subscriber SAPs the subscriber identification policy is currently associated with that are using the default definition to determine whether the active subscriber can be either reassigned to a subscriber SAP default or the system default subscriber profile. If all active subscribers cannot be reassigned, the removal attempt will fail.

Parameters *sub-ident-policy-name* — Specifies the name of the subscriber identification policy.

vpls-only-sap-parameters

Syntax **vpls-only-sap-parameters**

Context config>subscr-mgmt>msap-policy

Description This command enables the context to configure MSAP VPLS properties.

arp-host

Syntax **arp-host**

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

Description This command enables the context to configure ARP host parameters.

host-limit

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

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

Description This command configures the maximum number of ARP hosts.

Parameters *max-num-hosts* — Specifies the maximum number of ARP hosts.

Values 1 — 32767

min-auth-interval

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

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

Description This command configures the minimum authentication interval.

Parameters *min-auth-interval* — Specifies the minimum authentication interval.

Values 1 — 6000

sap-host-limit

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

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

Description This command configures the maximum number of ARP hosts per SAP.

Parameters *max-num-hosts-sap* — Specifies the maximum number of ARP hosts per SAP allowed on this IES interface.

Values 1 — 32767

arp-reply-agent

Syntax **arp-reply-agent** [**sub-ident**]
no arp-reply-agent

Context config>subscr-mgmt>msap-policy>vpls-only

Description This command enables a special ARP response mechanism in the system for ARP requests destined to static or dynamic hosts associated with the SAP. The system responds to each ARP request using the hosts MAC address as the both the source MAC address in the Ethernet header and the target hardware address in the ARP header.

ARP replies and requests received on an MSAP with **arp-reply-agent** enabled will be evaluated by the system against the anti-spoof filter entries associated with the ingress SAP (if the SAP has anti-spoof filtering enabled). ARPs from unknown hosts on the SAP will be discarded when anti-spoof filtering is enabled.

The ARP reply agent only responds if the ARP request enters an interface (SAP, spoke-SDP or mesh-SDP) associated with the VPLS instance of the MSAP.

A received ARP request that is not in the ARP reply agent table is flooded to all forwarding interfaces of the VPLS capable of broadcast except the ingress interface while honoring split-horizon constraints.

Static hosts can be defined using the **host** command. Dynamic hosts are enabled on the system by enabling the **lease-populate** command in the **dhcp** context. In the event that both a static host and a dynamic host share the same IP and MAC address, the VPLS ARP reply agent will retain the host information until both the static and dynamic information are removed. In the event that both a static and dynamic host share the same IP address, but different MAC addresses, the VPLS ARP reply agent is populated with the static host information.

The **arp-reply-agent** command will fail if an existing static host does not have both MAC and IP addresses specified. Once the ARP reply agent is enabled, creating a static host on the MSAP without both an IP address and MAC address will fail.

The ARP-reply-agent may only be enabled on SAPs supporting Ethernet encapsulation.

The **no** form of the command disables ARP-reply-agent functions for static and dynamic hosts on the MSAP.

Default not enabled

Parameters **sub-ident** — Configures the arp-reply-agent to discard ARP requests received on the MSAP that are targeted for a known host on the same MSAP with the same subscriber identification.

Hosts are identified by their subscriber information. For DHCP subscriber hosts, the subscriber hosts, the subscriber information is configured using the optional subscriber parameter string.

When arp-reply-agent is enabled with **sub-ident**:

- If the subscriber information for the destination host exactly matches the subscriber information for the originating host and the destination host is known on the same MSAP as the source, the ARP request is silently discarded.
- If the subscriber information for the destination host or originating host is unknown or undefined, the source and destination hosts are not considered to be the same subscriber. The ARP request is forwarded outside the MSAP's Split Horizon Group.
- When **sub-ident** is not configured, the arp-reply-agent does not attempt to identify the subscriber information for the destination or originating host and will not discard an ARP request based on subscriber information.

dhcp

Syntax **dhcp**

Context config>subscr-mgmt>msap-policy>vpls-only

Description This command enables the context to configure DHCP parameters.

lease-populate

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

Context config>subscr-mgmt>msap-policy>vpls-only

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

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

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

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

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

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

Default no lease-populate

option

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

Context config>subscr-mgmt>msap-policy>vpls-only>dhcp
config>service>ies>sub-if>dhcp

Description This command enables DHCP Option 82 (Relay Agent Information Option) parameters processing and enters the context for configuring Option 82 sub-options.

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

Default no option

action

Syntax **action** {**replace** | **drop** | **keep**}
no action

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

Description This command configures the Relay Agent Information Option (Option 82) processing.

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

Default The default is to keep the existing information intact.

Parameters **replace —** In the upstream direction (from the user), the Option 82 field from the router is inserted in the packet (overwriting any existing Option 82 field). In the downstream direction (towards the user) the Option 82 field is stripped (in accordance with RFC 3046).

drop — The DHCP packet is dropped if an Option 82 field is present, and a counter is incremented.

keep — The existing information is kept in the packet and the router does not add any additional information. In the downstream direction the Option 82 field is not stripped and is forwarded towards the client.

circuit-id

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

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

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

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

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

Default circuit-id

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

vlan-ascii-tuple — Specifies that the format will include VLAN-id and dot1p bits in addition to what is included in ascii-tuple already. The format is supported on dot1q and qinq ports only. Thus, when the option 82 bits are stripped, dot1p bits will be copied to the Ethernet header of an outgoing packet.

vendor-specific-option

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

Context config>subscr-mgmt>msap-policy>vpls-only>dhcp>option
config>service>ies>sub-if>dhcp

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

client-mac-address

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

Context config>subscr-mgmt>msap-policy>vpls-only>dhcp>option>vendor
config>service>ies>sub-if>dhcp>option

Description This command enables the sending of the MAC address in the Alcatel-Lucent vendor specific sub-option of the DHCP relay packet.

The **no** form of the command disables the sending of the MAC address in the Alcatel-Lucent vendor specific sub-option of the DHCP relay packet.

sap-id

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

Context config>subscr-mgmt>msap-policy>vpls-only>dhcp>option>vendor
config>service>ies>sub-if>dhcp>option

Description This command enables the sending of the SAP ID in the Alcatel-Lucent vendor specific sub-option of the DHCP relay packet.

The **no** form of the command disables the sending of the SAP ID in the Alcatel-Lucent vendor specific sub-option of the DHCP relay packet.

service-id

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

Context config>subscr-mgmt>msap-policy>vpls-only>dhcp>option>vendor
config>service>ies>sub-if>dhcp>option

Description This command enables the sending of the service ID in the Alcatel-Lucent vendor specific sub-option of the DHCP relay packet.

The **no** form of the command disables the sending of the service ID in the Alcatel-Lucent vendor specific sub-option of the DHCP relay packet.

string

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

Context config>subscr-mgmt>msap-policy>vpls-only>dhcp>option>vendor
config>service>ies>sub-if>dhcp>option

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

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

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

system-id

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

Context config>subscr-mgmt>msap-policy>vpls-only>dhcp>option>vendor
config>service>ies>sub-if>dhcp>option

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

proxy-server

Syntax **proxy-server**

Context config>subscr-mgmt>msap-policy>vpls-only>dhcp
config>service>ies>sub-if>dhcp

Description This command configures the DHCP proxy server.

emulated-server

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

Context config>subscr-mgmt>msap-policy>vpls-only>dhcp>proxy
config>service>ies>sub-if>dhcp

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

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

Parameters *ip-address* — Specifies the emulated server’s IP address. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).

lease-time

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

Context config>subscr-mgmt>msap-policy>vpls-only>dhcp>proxy

config>service>ies>sub-if>dhcp

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

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

Default 7 days 0 hours 0 seconds

Parameters **override —** Specifies that the local-proxy-server will use the configured lease-time information to provide DHCP clients.

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

Values 0 — 3650

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

Values 0 — 23

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

Values 0 — 59

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

Values 0 — 59

egress

Syntax **egress**

Context config>subscr-mgmt>msap-policy>vpls-only

Description This command configures egress policies for MSAPs.

multicast-group

Syntax **multicast-group** *group-name*
no multicast-group

Context config>subscr-mgmt>msap-policy>vpls-only>egress

Description This command specifies an existing egress multicast group (EMG). An EMG is created as an object used to group VPLS SAPs that are allowed to participate in efficient multicast replication (EMR). EMR is a method to increase the performance of egress multipoint forwarding by sacrificing some destination-based features. Eliminating the requirement to perform unique features for each destination allows the egress forwarding plane to chain together multiple destinations into a batch replication process. In order to perform this batch replication function, similar characteristics are required on each SAP within the EMG.

Only SAPs defined on Ethernet access ports are allowed into an egress-multicast-group.

In order to understand the purpose of an egress-multicast-group, an understanding of the system's use of flooding lists is required. A flooding list is maintained at the egress forwarding plane to define a set of destinations to which a packet must be replicated. Multipoint services make use of flooding lists to enable forwarding a single packet to many destinations. Examples of multipoint services that use flooding lists are VPLS, IGMP snooping and IP multicast routing. Currently, the egress forwarding plane will only use efficient multicast replication for VPLS and IGMP snooping flooding lists.

In VPLS services, a unique flooding list is created for each VPLS context. The flooding list is used when a packet has a broadcast, multicast or unknown destination MAC address. From a system perspective, proper VPLS handling requires that a broadcast, multicast or unknown destined packet be sent to all destinations that are in the forwarding state. The ingress forwarding plane ensures the packet gets to all egress forwarding planes that include a destination in the VPLS context. It is the egress forwarding plane's job to replicate the packet to the subset of the destinations that are reached through its interfaces and each of these destinations are included in the VPLS context's flooding list.

For IGMP snooping, a unique flooding list is created for each IP multicast (s,g) record. This (s,g) record is associated with an ingress VPLS context and may be associated with VPLS destinations in the source VPLS instance or other VPLS instances (in the case of MVR). Again, the ingress forwarding plane ensures that an ingress IP multicast packet matching the (s,g) record gets to all egress forwarding planes that have a VPLS destination associated with the (s,g) record. The egress forwarding plane uses the flooding list owned by the (s,g) record to replicate the packet to all VPLS destinations in the flooding list. The IGMP Snooping function identifies which VPLS destinations should be associated with the (s,g) record.

With normal multicast replication, the egress forwarding plane examines which features are enabled for each destination. This includes ACL filtering, mirroring, encapsulation and queuing. The resources used to perform this per destination multicast processing are very expensive to the egress forwarding plane when high replication bandwidth is required. If destinations with similar egress functions can be grouped together, the egress forwarding plane can process them in a more efficient manner and maximize replication bandwidth.

The egress-multicast-group object is designed to allow the identification of SAPs with similar egress characteristics. When a SAP is successfully provisioned into an egress-multicast-group, the system is ensured that it may be batched together with other SAPs in the same group at the egress forwarding plane for efficient multicast replication. A SAP that does not meet the common requirements is not allowed into the egress-multicast-group.

At the forwarding plane level, a VPLS flooding list is categorized into chainable and non-chainable destinations. Currently, the only chainable destinations are SAPs within an egress-multicast-group. The chainable destinations are further separated by egress-multicast-group association. Chains are then created following the rules below:

- A replication batch chain may only contain SAPs from the same egress-multicast-group
- A replication batch chain length may not exceed the dest-chain-limit of the egress-multicast-group to which the SAPs are members

Further subcategories are created for an IGMP (s,g) flooding list. A Layer 2 (s,g) record is created in a specific VPLS instance (the instance the (s,g) flow ingresses). SAPs within that VPLS context that join the (s,g) record are considered native SAPs within the flooding list. SAPs that join the (s,g) flooding list through the multicast VPLS registration process (MVR) from another VPLS context using the **from-vpls** command are considered alien SAPs. The distinction between native and alien in the list is maintained to allow the forwarding plane to enforce or suspend split-horizon-group (SHG) squelching. When the source of the (s,g) matching packet is in the same SHG as a native SAP, the packet must not be replicated to that SAP. For a SAP in another VPLS context, the source SHG of the

packet has no meaning and the forwarding plane must disregard SHG matching between the native source of the packet and the alien destination. Because the SHG squelch decision is done for the whole chain based on the first SAP in the chain, all SAPs in the chain must be all native or all alien SAPs. Chains for IGMP (s,g) flooding lists are created using the following rules:

1. A replication batch chain may only contain SAPs from the same egress-multicast-group.
2. A replication batch chain may only contain all alien or all native SAPs.
3. A replication batch chain length may not exceed the dest-chain-limit of the egress-multicast-group to which the SAPs are members

When a packet associated with a flooding list is received by the egress forwarding plane, it processes the packet by evaluating each destination on the list sequentially in a replication context. If the current entry being processed in the list is a non-chained destination, the forwarding plane processes the packet for that destination and then moves on to process other packets currently in the forwarding plane before returning to process the next destination in the list. If the current entry being processed is a chained destination, the forwarding plane remains in the replication context until it has forwarded to each entry in that chain. Once the replication context finishes with the last entry in the chain, it moves on to process other packets waiting for egress processing before returning to the replication context. Processing continues in this manner until the packet has been forwarded to all destinations in the list.

Batch chain processing of a chain of SAPs improves replication efficiency by bypassing the functions that perform egress mirroring decisions on SAPs within the chain and making a single ACL filtering decision for the whole chain. Each destination in the chain may have a unique egress QoS policy and per destination queuing is still performed for each destination in the chain. Also, while each SAP in the chain must be on access ports with the same encap-type, if the encap-type is dot1q, each SAP may have a unique dot1q tag.

One caveat to each SAP having a unique egress QoS policy in the chain is that only the Dot1P marking decisions for the first SAP in the list is enforced. If the first SAP's QoS policy forwarding class action states that the packet should not be remarked, none of the replicated packets in the chain will have the dot1P bits remarked. If the first SAP's QoS policy forwarding class action states that the packet should be remarked with a specific dot1P value, all the replicated packets for the remaining SAPs in the chain will have the same dot1P marking.

While the system supports 32 egress multicast groups, a single group would usually suffice. An instance where multiple groups would be needed is when all the SAPs requiring efficient multicast replication cannot share the same common requirements. In this case, an egress multicast group would be created for each set of common requirements. An egress multicast group may contain SAPs from many different VPLS instances. It should be understood that an egress multicast group is not equivalent to an egress forwarding plane flooding list. An egress multicast group only identifies which SAPs may participate in efficient multicast replication. As stated above, entries in a flooding list are populated due to VPLS destination creation or IGMP snooping events.

The **no** form of the command removes a specific egress multicast group. Deleting an egress multicast group will only succeed when the group has no SAP members. To remove SAP members, use the **no multicast-group** *group-name* command under each SAP's egress context.

Note: Efficient multicast replication will only be performed on IOMs that support chassis mode b If an IOM does not support mode b operation, egress-multicast-group membership is ignored on that IOM's egress forwarding planes. The chassis need not be placed into mode b for efficient multicast replication to be performed on the capable IOMs.

Parameters *group-name* — Multiple egress multicast groups may be created on the system. Each must have a unique name. The egress-multicast-group-name is an ASCII string up to 16 characters in length

and follows all the naming rules as other named policies in the system. The group's name is used throughout the system to uniquely identify the Egress Multicast Group and is used to provision a SAP into the group.

Default None, each egress multicast group must be explicitly configured.

Values Up to 32 egress multicast groups may be created on the system.

igmp-snooping

Syntax **igmp-snooping**

Context config>subscr-mgmt>msap-policy>vpls-only

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

Default none

fast-leave

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

Context config>subscr-mgmt>msap-policy>vpls-only>igmp-snp

Description This command enables fast leave.

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

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

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

Default no fast-leave

import

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

Context config>subscr-mgmt>msap-policy>vpls-only>igmp-snp

Description This command specifies the import routing policy to be used for IGMP packets to be used on this SAP or SDP. Only a single policy can be imported on a single SAP at any time.

The **no** form of the command removes the policy association from the SAP or SDP.

Default no import (No import policy is specified)

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

last-member-query-interval

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

Context config>subscr-mgmt>msap-policy>vpls-only>igmp-snp

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

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

Default 10

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

Values 1 — 50

max-num-groups

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

Context config>subscr-mgmt>msap-policy>vpls-only>igmp-snp

Description This command defines the maximum number of multicast groups that can be joined on an MSAP or SDP. If the router receives an IGMP join message that would exceed the configured number of groups, the request is ignored.

Default no max-num-groups

Parameters *max-num-groups* — Specifies the maximum number of groups that can be joined on an MSAP or SDP.

Values 1 — 1000

mcac

Syntax **mcac**

Context config>subscr-mgmt>msap-policy>vpls-only>igmp-snp>

Description This command enables the context to configure multicast CAC parameters.

Default none

mc-constraints

Syntax **mc-constraints**

Context config>subscr-mgmt>msap-policy>vpls-only>igmp-snp>mcac

Description This command enables the context to configure the level and its associated bandwidth for a bundle or a logical interface.

Default none

level

Syntax **level** *level-id* **bw** *bandwidth*
no level *level-id*

Context config>subscr-mgmt>msap-policy>vpls-only>igmp-snp>mcac

Description This command configures levels and their associated bandwidth for multicast CAC policy on an interface.

Parameters *level-id* — Specifies has an entry for each multicast CAC policy constraint level configured on a system.

Values 1 — 8

bandwidth — Specifies the bandwidth in kilobits per second (kbps) for the level.

Values 1 — 2147483647

number-down

Syntax **number-down** *number-lag-port-down* **level** *level-id*
no number-down *number-lag-port-down*

Context config>subscr-mgmt>msap-policy>vpls-only>igmp-snp>mcac

Description This command configures the number of ports down along with level for multicast CAC policy on an MSAP

Parameters *number-lag-port-down* — If the number of ports available in the LAG is reduced by the number of ports configured in this command here then bandwidth allowed for bundle and/or interface will be as per the levels configured in this context.

level *level-id* **—** Specifies the amount of bandwidth available within a given bundle for MC traffic for a specified level.

policy

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

Context config>subscr-mgmt>msap-policy>vpls-only>igmp-snp>mcac

Description This command configures the multicast CAC policy name.

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

unconstrained-bw

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

Context config>subscr-mgmt>msap-policy>vpls-only>igmp-snp>mcac

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

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

Values 0 — 2147483647

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

If the *bandwidth* value is 0, no mandatory channels are allowed. If the value of *bandwidth* is '-1', then all mandatory and optional channels are allowed.

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

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

Values 0 — 2147483647

sub-mcac-policy

Syntax **sub-mcac-policy** *sub-mcac-policy-name* [**create**]
no sub-mcac-policy *sub-mcac-policy-name*

Context config>subscr-mgmt

Description This command will create a policy template with mcac bandwidth limits that will be applied to the subscriber.

Per interface mcac bandwidth limits will be set directly under the interface (regular interface or group-interface) and no such policy templates are needed.

The need for a separate policy template for subscribers is due to the fact that groups of subscribers under the same group-interface can share certain settings that can be configured via this template.

To summarize, the mcac bandwidth constraints for subscribers are defined in the sub-mcac-policy while the mcac bandwidth constraints for the interface are configured directly under the **igmp>interface>mcac** or **igmp>group-interface>mcac** context without the need for policy templates.

Note that the sub-mcac-policy only deals with the mcac bandwidth limits and not the channel bandwidth definitions. Channels bandwidth is defined in a different policy (under the configure>router>mcac hierarchy) and that policy is applied on the interface level as follows:

- For group-interface: under the **configure>service>vprn>igmp>group-interface>mcac** context
- For regular interface: under the **configure>service/router>igmp>interface>mcac** context.

In case of HQoS Adjustment, it is mandatory that the sub-mcac-policy be created and applied to the subscriber. The sub-mac-policy does not have to contain any bandwidth constrains, but it has to be in a no shutdown state in order for HQoS Adjustment to work.

Parameters *policy-name* — Specifies the name of the policy.

mvr

Syntax **mvr**

Context config>subscr-mgmt>msap-policy>vpls-only>igmp-snp

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

from-vpls

Syntax **from-vpls** *service-id*
no from-vpls

Context config>subscr-mgmt>msap-policy>vpls-only>igmp-snp>mvr

Description This command configures the VPLS from which multicast traffic is copied upon receipt of an IGMP join request.

IGMP snooping must be enabled on the MVR VPLS.

Default no from-vpls

Parameters *service-id* — Specifies the MVR VPLS from which multicast channels should be copied into an MSAP.

Values *service-id*: 1 — 2147483647
svc-name: 64 characters maximum

query-interval

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

Context config>subscr-mgmt>msap-policy>vpls-only>igmp-snp

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

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

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

Default 125

Parameters *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>subscr-mgmt>msap-policy>vpls-only>igmp-snp

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

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

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

Default 10

Parameters *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>subscr-mgmt>msap-policy>vpls-only>igmp-snp

Description This command configures the IGMP robustness variable. If the **send-queries** command is enabled, this parameter allows tuning for the expected packet loss on a SAP or SDP. The robust-count variable allows tuning for the expected packet loss on a subnet and is comparable to a retry count. If an MSAP or SDP is expected to be "lossy", this parameter may be increased. IGMP snooping on an MSAP or SDP is robust to (robust-count-1) packet losses.

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

Default 2

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

Values 2 — 7

send-queries

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

Context config>subscr-mgmt>msap-policy>vpls-only>igmp-snp

Description This command specifies whether to send IGMP general query messages on the managed SAP. When send-queries is configured, all type of queries generate ourselves are of the configured version. If a report of a version higher than the configured version is received, the report will get dropped and a new wrong version counter will get incremented.

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 there is a v2 querier, a v3 group or group-source specific query is never sent when a host wants to leave a certain group.

Default no send-queries

version

Syntax **version** *version*
no version

Context config>subscr-mgmt>msap-policy>vpls-only>igmp-snp

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

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

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

Parameters *version* — Specify the IGMP version.

Values 1, 2, 3

mac-da-hashing

[no] mac-da-hashing

Context config>subscr-mgmt>msap-policy>vpls-only>igmp-snp

Description This command specifies whether subscriber traffic egressing a LAG SAP has its egress LAG link selected by a function of the MAC destination address instead of the subscriber ID.

This command is only meaningful if subscriber management is enabled and can be configured for a VPLS service.

split-horizon-group

Syntax **split-horizon-group** *group-name*

Context config>subscr-mgmt>msap-policy>vpls-only>igmp-snp

Description This command specifies the name of the split horizon group to which the MSAP belongs.

default-msap-policy

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

Context config>service>vpls>sap

Description This command specifies the default managed SAP policy to use to create MSAPs when the response from the RADIUS server does not specify a managed SAP policy.

The *policy-name* parameter is only valid for a SAP with the keywords **capture-sap** specified in the SAP's configuration. The **capture-sap** keyword in the SAP configuration captures the SAP where triggering packets will be sent to the CPM. Non-triggering packets captured by the capture SAP will be dropped.

The managed SAP policy must already be defined in the **config>subscr-mgmt>msap-policy** context

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

Default no default-msap-policy

Parameters *policy-name* — /Specifies an existing default managed SAP policy.

trigger-packet

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

Context config>service>vpls>sap

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

Default none

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

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

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

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

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

eval-msap

Syntax **eval-msap** {**policy** *msap-policy-name* | **msap** *sap-id*}

Context tools>perform>subscr-mgmt

Description This command evaluates managed SAP policies.

Parameters **policy** *msap-policy-name* **—** Specifies an existing MSAP policy.

msap *sap-id* **—** Specifies an MSAP sap-id.

Values [*port-id*|lag-*id*]:*qtag1*
[*port-id*|lag-*id*]:*qtag1.qtag2*

Subscriber Management Service Commands

This section contains:

Layer 3 Subscriber Interfaces Commands

subscriber-interface

Syntax **subscriber-interface** *ip-int-name* [**fwd-service** *service-id* **fwd-subscriber-interface** *ip-int-name*] [**create**]

Context config>service>ies
config>service>vprn

Description This command allows the operator to create special subscriber-based interfaces. It is used to contain multiple group interfaces. Multiple subnets associated with the subscriber interface can be applied to any of the contained group interfaces in any combination. The subscriber interface allows subnet sharing between group interfaces.

Use the **no** form of the command to remove the subscriber interface.

Default no subscriber interfaces configured

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

fwd-service *service-id* —

Values

fwd-subscriber-interface *ip-int-name* —

address

Syntax [**no**] **address** {*ip-address/mask* | *ip-address netmask*} [**gw-ip-address** *ip-address*] [**populate-host-routes**]

Context config>service>ies>subscriber-interface
config>service>vprn>subscriber-interface

Description This command creates or removes an IP address, IP subnet or broadcast address format for the interface. Multiple IP addresses can be associated with a subscriber-interface

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

In the IES subscriber interface context, this command is used to assign one or more 256maximum) host IP addresses and subnets. This differs from a normal IES interfaces where **secondary** command creates and additional subnet after the primary address is assigned. A user can then add or remove addresses without having to keep a primary address.

Use the **no** form of this command to remove the IP address assignment from the IP interface.

Default no IP address or subnet associations configured

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

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

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

netmask — The subnet mask in dotted decimal notation.

Values 0.0.0.0 - 255.255.255.255

gw-ip-address *ip-address* **—** Specifies a separate IP address within the subnet for SRRP routing purposes. This parameter must be followed by a valid IP interface that exists within the subscriber subnet created by the address command. The defined gateway IP address cannot currently exist as a subscriber host (static or dynamic). If the defined ip-address already exists as a subscriber host address, the address command will fail. The specified ip-address must be unique within the system.

The gw-address parameter may be specified at anytime. If the subscriber subnet was created previously, executing the address command with a gw-address parameter will simply add the SRRP gateway IP address to the existing subnet.

If the address command is executed without the gw-address parameter when the subscriber subnet is associated with an active SRRP instance, the address will fail. If the SRRP instance is inactive or removed, executing the address command without the gw-address parameter will remove the SRRP gateway IP address from the specified subscriber subnet.

If the address command is executed with a new gw-address, all SRRP instances currently associated with the specified subscriber subnet will be updated with the new SRRP gateway IP address.

populate-host-routes — Specifies to populate subscriber-host routes in local FIB. Storing them in FIB benefits topologies only where the external router advertises more specific routes than the one corresponding to locally configured subscriber-interface subnets.

allow-unmatching-subnets

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

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

Description This command allows address assignment for IPoEv4 and PPPoEv4 subscriber hosts in cases where the subscriber assigned IPv4 address falls outside of the subscriber-interface subnet configured under the same CLI hierarchy. Such subscriber host will be installed in the FIB as /32 hosts because the aggregated subscriber-interface route is not available for them (not configured under the subscriber-interface). Without the **allow-unmatching-subnets** command, such host are instantiated in the system but forwarding for them is disabled.

This command can be only configured in case where the subscriber-interface has an IP address (and therefore subnet) configured. In case where the subscriber interface does not have explicitly configured and IP address, execution of this command will fail.

IPv6 hosts are not affected by this command.

Default no allow-unmatching-subnets

allow-unmatching-subnets

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

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

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

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

IPv4 subscriber hosts are unaffected by this command.

Default no allow-unmatching-subnets

authentication-policy

Syntax **authentication-policy** *name*
no authentication-policy

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

Description This command assigns an authentication policy to the interface.

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

Default no authentication-policy

Parameters *name* — Specifies the authentication policy name. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

arp-populate

Syntax [**no**] **arp-populate**

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

Description This command enables populating static and dynamic hosts into the system ARP cache. When enabled, the host's IP address and MAC address are placed in the system ARP cache as a managed entry. Static hosts must be defined on the interface using the **host** command. Dynamic hosts are enabled on the system through enabling lease-populate in the IP interface DHCP context. In the event that both a static host and a dynamic host share the same IP and MAC address, the system's ARP cache retains the host information until both the static and dynamic information are removed. Both static and dynamic hosts override static ARP entries. Static ARP entries are marked as inactive when they conflict with static or dynamic hosts and will be repopulated once all static and dynamic host information for the IP address are removed. Since static ARP entries are not possible when static subscriber hosts are defined or when DHCP lease state table population is enabled, conflict between static ARP entries and the arp-populate function is not an issue.

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

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

arp-populate can only be enabled on VPRN interfaces supporting Ethernet encapsulation.

Use the **no** form of the command to disable ARP cache population functions for static and dynamic hosts on the interface. All static and dynamic host information in the systems ARP cache will be removed. Any existing static ARP entries previously inactive due to static or dynamic hosts will be populated in the system ARP cache.

When **arp-populate** is enabled, the system will not send out ARP Requests for hosts that are not in the ARP cache. Only statically configured and DHCP learned hosts are reachable through an IP interface with arp-populate enabled.

Default not enabled

arp-timeout

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

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

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

The **no** form of this command restores **arp-timeout** to the default value.

Default 14400 seconds

Parameters *seconds* — The minimum number of seconds a learned ARP entry will be stored in the ARP table, expressed as a decimal integer. A value of zero specifies that the timer is inoperative and learned ARP entries will not be aged.

Values 0 — 65535

client-applications

Syntax **client-applications** {[**dhcp**] [**ppp**]}
no client-applications

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

Description This command enables the clients that will try to contact the DHCP server(s).

The **no** form of the command removes the server client type from the configuration.

Parameters **dhcp —** Specifies that the DHCP relay will forward requests to the DHCP server(s).

pppoe — Specifies that PPPoE will attempt to request an IP address for a PPPoE client from the DHCP server(s)ly assigned to PPPoE node.

lease-populate

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

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

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

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

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

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

- Populate an MSAP based anti-spoof filter table to provide dynamic anti-spoof filtering. If the system is unable to populate the dynamic host information in the anti-spoof filter table on the

SAP, the DHCP ACK message must be discarded without adding new lease state entry or updating an existing lease state entry.

- Generate dynamic ARP replies if **arp-reply-agent** is enabled.

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

Default no lease-populate

delayed-enable

Syntax **delayed-enable** *seconds* [**init-only**]
no delayed-enable

Context config>service>ies>sub-if

Description This command delays making interface operational by the specified number of seconds.

In environments with many subscribers, it can take time to synchronize the subscriber state between peers when the subscriber-interface is enabled (perhaps, after a reboot). To ensure that the state has time to be synchronized, the **delayed-enable** timer can be specified. The optional parameter **init-only** can be added to use this timer only after a reboot.

Default no delayed-enable

Parameters *seconds* — Specifies the number of seconds to delay before the interface is operational.

Values 1 — 1200

init-only — Delays the initialization of the subscriber-interface to give the rest of the system time to complete necessary tasks such as allowing routing protocols to converge and/or to allow MCS to sync the subscriber information. The delay only occurs immediately after a reboot.

group-interface

Syntax **group-interface** *ip-int-name* [**create**]
group-interface *ip-int-name* [**create**] **lns**
group-interface *ip-int-name* [**create**] **softgre**
no group-interface *ip-int-name* [**create**]

Context config>service>ies>subscriber-interface
config>service>vprn>subscriber-interface

Description This command creates a group interface. This interface is designed for triple-play services where multiple SAPs are part of the same subnet. A group interface may contain one or more SAPs.

Use the **no** form of the command to remove the group interface from the subscriber interface.

Default no group interfaces configured

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

lns — Specifies to use LNS.

softgre — Specifies to use dynamic GRE encapsulation.

ip-mtu

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

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

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

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

Default none

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

Values 512 — 9000

ipoe-linking

Syntax **ipoe-linking**

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

Description This command enables the context to configure IPoE host linking.

gratuitous-rtr-adv

Syntax [**no**] **gratuitous-rtr-adv**

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

Description If enabled, this command controls generation of unsolicited Router-advertisement on creation of v4 host.

The **no** form of the command disables **gratuitous-rtr-adv**

Default gratuitous-rtr-adv

shared-circuit-id

Syntax [**no**] **shared-circuit-id**

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

Description If configured, circuit-id in DHCPv4 option-82 is used to authenticate DHCPv6. If DHCPv6 is received before DHCPv4, it is dropped. Also, a SLAAC host is created based on DHCPv4 authentication if RADIUS returns IPv6 framed-prefix. IPv6oE host is deleted if the linked IPv4oE host is deleted due to DHCP release or lease time-out. The linkage between IPv4 and IPv6 is based on SAP and MAC address. The sharing of circuit-id from DHCPv4 for authentication of DHCPv6 (or SLAAC) allows 7750 to work around lack of support for LDRA on Access-nodes.

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

Default no shared-circuit-id

ipv6

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

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

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

urpf-check

Syntax [**no**] **urpf-check**

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

Description This command enables unicast RPF (uRPF) Check on this interface.

The **no** form of the command disables unicast RPF (uRPF) Check on this interface.

Default disabled

mode

Syntax **mode** {**strict** | **loose** | **strict-no-ecmp**}
no mode

Context config>service>ies>if>urfp-check
config>service>ies>sub-if>group-if>ipv6>urfp-check

Description This command specifies the mode of unicast RPF check.

The **no** form of the command reverts to the default (strict) mode.

Default strict

Parameters **strict —** When specified, uRPF checks whether incoming packet has a source address that matches a prefix in the routing table, and whether the interface expects to receive a packet with this source address prefix.

loose — In **loose** mode, uRPF checks whether incoming packet has source address with a corresponding prefix in the routing table. However, the loose mode does not check whether the interface expects to receive a packet with a specific source address prefix. This object is valid only when **urpf-check** is enabled.

strict-no-ecmp — When a packet is received on an interface in this mode and the SA matches an ECMP route the packet is dropped by uRPF.

match-circuit-id

Syntax [**no**] **match-circuit-id**

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

Description This command enables Option 82 circuit ID on relayed DHCP packet matching. For routed CO, the group interface DHCP relay process is stateful. When packets are relayed to the server the virtual router ID, transaction ID, SAP ID, and client hardware MAC address of the relayed packet are tracked.

When a response is received from the server the virtual router ID, transaction ID, and client hardware MAC address must be matched to determine the SAP on which to send the packet out. In some cases, the virtual router ID, transaction ID, and client hardware MAC address are not guaranteed to be unique.

When the **match-circuit-id** command is enabled this as part of the key is used to guarantee correctness in our lookup. This is really only needed when dealing with an IP aware DSLAM that proxies the client hardware MAC address.

Default no match-circuit-id

mac

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

Context config>service>ies>subscriber-interface>group-interface

Description This command assigns a specific MAC address to a subscriber group interface.

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

Default The physical MAC address associated with the Ethernet interface that the SAP is configured on (the default MAC address assigned to the interface, assigned by the system).

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

oper-up-while-empty

Syntax [**no**] **oper-up-while-empty**

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

Description This command allows the subscriber interface to treat this group interface to be operationally enabled without any active SAPs.

This command is typically used with MSAPs where advertising the subnet prior to having a MSAP dynamically created is needed.

mode

Syntax **mode** *mode*

Context configure>card>mda>atm

Description This command configures the ATM MDA into a mode with the increased VC scale (16k VCs, as opposed to 8K VCs). ESM is supported only in 16K VCs mode. In 16K VCs mode, there is only one queue allocated to each VC in the ATM MDA. In 8K VCs mode, there are two queues allocated per VC.

The 16K VC mode is supported only on the 4 port oc-3/12c/STM-1/4c and the 16 port ATM oc-3/STM-1 ATM MDA.

Changing the ATM MDA mode requires a reset of the MDA. A warning is issued asking for the confirmation before the command is executed.

Default max8k-vc.

Parameters *mode* — Specifies VC scale.

Values max8k-vc | max16k-vc

agg-rate-limit

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

Context configure>port>sonet-sdh>path>access>egress>vport

Description This command configures the aggregate rate limit in the vport container. This command is mutually exclusive with the **port-scheduler-policy** command under the vport. The **agg-rate-limit** is only effective when used with the port scheduler that is defined under the **configure>port>sonet-**

sdh>path>egress-scheduler-policy *port-scheduler-policy-name* hierarchy. Without the **port-scheduler**, the aggregate rate limit will have no effect.

The **agg-limit-rate** is on-the-wire rate of the directly connected ATM port. The method of calculating on-the-wire rate depends on whether the encap-offset (sub-profile) command is configured or not.

Default None.

Parameters *agg-rate* — Specifies the on-the-wire rate in Kbps.

Values [1..100000000 | max] Kbps

vpi

Syntax **vpi** *vpi* **egress-traffic-desc** ***a****tm-td-profile-id*
no vpi *vpi*

Context configure>port>sonet-sdh>path>atm

Description This command enables the ATM VP shaper under the ATM port. The type of ATM shaper are CBR or rt/nrt-VBR as defined by the traffic descriptor. It cannot be a UBR service-type.

All VCs within the shaper will degrade into a UBR type service class. For example, when a CBR type VC is associated with the shaper, it will degrade into a UBR type VC. Scheduling traffic amongst VCs within the shaper is based on WRR using the weight parameter.

If the VP shaper is deleted, the VCs that were under it is restored to their original service category.

The VP shaper is statically configured and instantiated upon configuration.

A VP shaper can be seamlessly added to or removed from the active VCs in the system.

Default none

Parameters *atm-td-profile-id* — Specifies ATM traffic description id.

Values [1..1000]

vpi —

Values [0..4095]

egress-traffic-desc — References an atm traffic descriptor profile.

traffic-desc

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

Context configure>service>vprn>sub-if>grp-if>sap>atm>egress
configure>service>vprn>sub-if>grp-if>sap>atm>ingress
configure>service>ies>sub-if>grp-if>sap>atm>egress
configure>service>ies>sub-if>grp-if>sap>atm>ingress
configure>subscr-mgmt>msap-policy>atm>egress
configure>subscr-mgmt>msap-policy>atm>ingress

Description This command references traffic-descriptor id for VPs and VCs.

The VP Shaper cannot be of service-type UBR.

Default Default traffic descriptor (id=1) of UBR type.

Parameters *atm-td-profile-id* — Specifies traffic-descriptor id.

Values [1..1000]

weight

Syntax **weight** *weight*
no weight

Context configure>qos>atm-td-profile

Description VCs within the VP tunnel are serviced by a single scheduler assigned to each VP tunnel. VCs within the shaped VP tunnel will be degraded from the originally assigned service category to a common UBR service category (default traffic descriptor). Scheduling between VCs will be based on WRR with a weight parameter that can be explicitly configured in the ATM traffic descriptor profile. If weight is not specifically configured, the defaults are taken.

The explicitly configured weight parameter is honored only on the ATM MDA in the max16k-vc mode. On all other ATM capable MDAs (ASAP or ATM MDA in max8k-vc mode), the weight parameter is ignored.

Default VC degraded from CBR = weight 10

VC degraded from rt-VBR = weight 7

VC degraded from nrt-VBR = weight 5

VC degraded from UBR+ = weight 2

VC degraded from UBR = weight 1

Parameters *weight* —

Values [1-255]

encapsulation

Syntax **encapsulation** [**aal5auto | aal5nlpid-ppp | aal5mux-ppp | aal5snap-bridged | aal5mux-bridged-eth-nofcs**]
no encapsulation

Context configure>service>ies>sub-if>grp-if>sap>atm
configure>service>vprn>sub-if>grp-if>sap>atm
configure>service>vpls>sap>atm

Description This command is a SAP level command and it will either statically set or enable dynamic detection of the encapsulation.

Default snap-bridged

Parameters **aal5auto —** This option is available only in max16k-vc mode on dynamic or static SAPs. It will enable automatic detection of one of the four supported encapsulation types.

aal5mux-bridged-eth-nofcs — This option already exist outside of the ESM context on regular interfaces. Within the ESM context (group-interfaces and capture SAPs), this option is available only in max16K-vc mode. The encapsulation is statically set to VC-MUX bridged Ethernet with no FCS. This is a valid encapsulation only for PPPoEoA.

aal5mux-ppp — This option is available only in max16k-vc mode on dynamic or static SAPs. The encapsulation is statically set VC-MUX PPP encapsulation. This is a valid encapsulation only for PPPoA.

aal5nlpid-ppp — dynamic or static SAPs. The encapsulation is statically set to NLPID (LLC) PPP encapsulation. This is a valid encapsulation only for PPPoA.

aal5snap-bridged — This option already exist outside of the ESM context on regular interfaces. Within the ESM context (group-interfaces and capture SAPs), this option is available only in max16k-vc mode. The encapsulation is statically set to bridged Ethernet with or without FCS. Both PIDs (0x 00-01 and 0x 00-07) are accepted on ingress and use this to determine whether to strip four bytes from the end of the encapsulated Ethernet frame. The inner FCS is not checked. This is a valid encapsulation only for PPPoEoA.

Note that on ATM frames with Ethernet FCS or without FCS are accepted but only frames with no Ethernet FCS are sent.

def-inter-dest-id

Syntax **def-inter-dest-id string** *interest-string*
def-inter-dest-id {**use-top-q** | **use-vpi**}
no def-inter-dest-id

Context configure>service>ies>sub-if>grp-if>sap>sub-sla-mgmt
configure>service>vprn>sub-if>grp-if>sap>sub-sla-mgmt
configure>subscr-mgmt>msap-policy>sub-sla-mgmt

Description This command is used to associate the vport with the subscriber. The association method will depend on the configured option.

Default Disabled

Parameters *string* — A RADIUS VSA (Alc-Int-Dest-Id-Str, type 28) obtained during the subscriber authentication phase will contain the destination string name that will be matched against the string defined under the vport. In this fashion the subscriber host will be associated with the corresponding vport.

Alternatively, the destination string can be defined in LUDB.

use-top-q — This is applicable only to Ethernet ports.

use-vpi — VP Identifier (VPI) will be used to make the association between the subscriber and the vport automatically.

Control Plane will be aware of the VPI during the session initiation phase. This VPI will be used to make the association between the host and the vport with the same name (VPI number). Note

that in this case the vport name under the **configure>port>sonet-sdh>path>access>egress** context must be the VPI number.

pppoe-user-db

Syntax **pppoe-user-db** *ludb-name*
no pppoe-user-db

Context configure>services>vpls>sap

Description This command will enable LUDB authentication on capture SAPs for PPPoE(oA) clients. In case that this command is configured along with the authentication-policy command (RADIUS authentication), then the authentication-policy command will take precedence.

Optionally, with a separate command (ppp-user-db) PPPoA clients can be authenticated under a separate LUDB.

Default Disabled

Parameters *ludb-name* — Name of local user database.

ppp-user-db

Syntax **pppp-user-db** *ludb-name*
no pppp-user-db

Context configure>services>vpls>sap

Description This command will enable LUDB authentication on capture SAPs for PPPoA clients. In case that this command is configured along with the authentication-policy command (RADIUS authentication), then the authentication-policy command will take precedence.

Optionally, with a separate command (pppoe-user-db) PPPoE(oA) clients can be authenticated under a separate LUDB.

Default Disabled

Parameters *ludb-name* — Name of local user database.

ppp-policy

Syntax **ppp-policy** *ppp-pol-name*
no ppp-policy

Context configure>services>vpls>sap

Description This command will reference a ppp-policy that will define session parameters (ppp-mtu, authentication options, etc.) during the session initiation phase. Normally, ppp-policy is referenced under the group-interface hierarchy. But with capture SAP is it not known at the session initiation phase to which group-interface the session belongs. This is why, with the capture SAP, the ppp-policy must be

referenced directly under the capture SAP. The ppp-policy referenced under the group-interface must be the same as the ppp-policy referenced under the capture SAP. Otherwise the session will not come up.

Default Disabled

Parameters *ppp-pol-name* — Name of the ppp-policy.

pppoe-policy

Syntax **pppoe-policy** *ppoep-pol-name*
no pppoe-policy

Context configure>services>vpls>sap

Description This command will reference a pppoe-policy that will define session parameters (ppp-mtu, authentication options, etc.) during the session initiation phase. Normally, pppoe-policy is referenced under the group-interface hierarchy. But with capture SAP is it not known at the session initiation phase to which group-interface the session belongs. This is why, with the capture SAP, the ppp-policy must be referenced directly under the capture SAP. The pppoe-policy referenced under the group-interface must be the same as the pppoe-policy referenced under the capture SAP. Otherwise the session will not come up.

Default Disabled

Parameters *pppoe-pol-name* — Name of the pppoe-policy

vc-range

Syntax **vc-range** *num* **vpi-range** *vpi-range* **vci-range** *vci-range*
no vc-range num

Context configure>services>vpls>sap>atm

Description This command is supported only in max16k-vc ATM MDA mode. An ATM MDA supports a number (see scaling guides for more info) of passive (or listening) VCs, of which a subset can be simultaneously active.

Default Disabled

Parameters *num* — Specifies the VC range.

Values 1 — 5 (Five ranges are supported to accommodate non-contiguous ranges of VPI/VCI pairs.)

vci-range *vci-range* **—** Specifies the VCI range.

Values 1, 2, 5 — 65535 (Contiguous VCI ranges in the form of ‘x’-‘y’.)

vpi-range *vpi-range.* — Specifies the VPI range.

Values 0 — 255 for UNI
0 — 4095 for NNI
(Contiguous VPI range in the form of 'x'-'y'.)

Layer 3 Subscriber Interfaces SAP Commands

accounting-policy

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

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

Description This command specifies the policy to use to collect accounting statistics on a subscriber profile.

A maximum of one accounting policy can be associated with a profile at one time.

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

Default no accounting policy

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

Values 1 — 99

anti-spoof

Syntax **anti-spoof** {**ip** | **ip-mac** | **nh-mac**}
no anti-spoof

Context config>service>vprn>if>sap
config>service>ies>sub-if>grp-if>sap
config>service>vprn>sub-if>grp-if>sap
config>subscr-mgmt>msap-policy

Description This command configures the anti-spoof type of the MSAP.

The type of anti-spoof filtering defines what information in the incoming packet is used to generate the criteria to lookup an entry in the anti-spoof filter table. The type parameter (**ip**, **ip-mac**) defines the anti-spoof filter type enforced by the SAP when anti-spoof filtering is enabled.

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

Note that for IES and VPRN subscriber group interfaces, setting no anti-spoof will set the default anti-spoofing type which is **ip-mac**.

Default no anti-spoof

Parameters **ip —** Configures SAP anti-spoof filtering to use only the source IP address in its lookup. If a static host exists on the SAP without an IP address specified, the anti-spoof type **ip** command will fail. Note that this parameter is not applicable in the **config>subscr-mgmt>msap-policy** context.

ip-mac — Configures SAP anti-spoof filtering to use both the source IP address and the source MAC address in its lookup. If a static host exists on the SAP without both the IP address and MAC address specified, the anti-spoof type **ip-mac** command will fail. This is also true if the default anti-spoof filter type of the SAP is **ip-mac** and the default is not overridden. The anti-spoof type **ip-mac** command will also fail if the SAP does not support Ethernet encapsulation.

nh-mac — Indicates that the ingress anti-spoof is based on the source MAC and egress anti-spoof is based on the nh-ip-address.

app-profile

Syntax **app-profile** *app-profile-name*
no app-profile

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

Description This command configures the application profile name.

Parameters *app-profile-name* — Specifies an existing application profile name configured in the **config>app-assure>group>policy** context.

collect-stats

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

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

Description When enabled, the agent collects non-RADIUS accounting statistics on a subscriber profile.

When the **no collect-stats** command is issued the statistics are still accumulated by the IOM cards. However, the CPU will not obtain the results and write them to the billing file. If a subsequent **collect-stats** command is issued then the counters written to the billing file include all the traffic while the **no collect-stats** command was in effect.

Default collect-stats

cpu-protection

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

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

Description This command assigns an existing CPU protection policy to the associated group interface. The CPU protection policies are configured in the **config>sys>security>cpu-protection>policy** *cpu-protection-policy-id* context.

If no CPU-Protection policy is assigned to a group interface SAP, then the default policy is used to limit the overall-rate. The default policy is policy number 254 for access interfaces and 255 for network interfaces.

The **no** form of the command removes the association of the CPU protection policy from the associated interface and reverts to the default policy values.

Default cpu-protection 254 (for access interfaces)

cpu-protection 255 (for network interfaces)

The configuration of no cpu-protection returns the interface/SAP to the default policies as shown above.

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

Values 1 — 255

mac-monitoring — This keyword enables MAC monitoring.

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

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

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

egress

Syntax **egress**

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

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

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

filter

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

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

Description This command associates an IP filter policy with an ingress or egress Service Access Point (SAP). Filter policies control the forwarding and dropping of packets based on the matching criteria.

MAC filters are only allowed on Epipe and Virtual Private LAN Service (VPLS) SAPs.

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

In general, filters applied to SAPs (ingress or egress) apply to all packets on the SAP. One exception is non-IP packets are not applied to the match criteria, so the default action in the filter policy applies to these packets.

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

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

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

ip-filter-id — Specifies the ID for the IP filter policy. Allowed values are an integer in the range of 1 and 65535 that corresponds to a previously created IP filter policy in the **configure>filter** context.

qos

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

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

Description Associates a Quality of Service (QoS) policy with an egress Service Access Point (SAP) or IP interface.

QoS egress policies are important for the enforcement of SLA agreements. The policy ID must be defined prior to associating the policy with a SAP or IP interface. If the *policy-id* does not exist, an error will be returned.

The **qos** command is used to associate egress QoS policies. The **qos** command only allows egress policies on SAP or IP interface egress. Attempts to associate a QoS policy of the wrong type returns an error.

Only one ingress QoS policy can be associated with a SAP or IP interface at one time. Attempts to associate a second QoS policy of a given type will return an error.

By default, no specific QoS policy is associated with the SAP or IP interface for egress, so the default QoS policy is used.

The normal behavior is for queues to be created per destination.

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

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

Values 1 — 65535

qos

Syntax **qos** *policy-id* [**shared-queuing**]
no qos

Context config>service>ies>sub-if>grp-if>sap>ingress

Description Associates a Quality of Service (QoS) policy with an ingress Service Access Point (SAP) or IP interface.

QoS ingress policies are important for the enforcement of SLA agreements. The policy ID must be defined prior to associating the policy with a SAP or IP interface. If the *policy-id* does not exist, an error will be returned.

This **qos** command is used to associate ingress QoS policies. The **qos** command only allows ingress policies to be associated on SAP or IP interface ingress.

Only one ingress and one egress QoS policy can be associated with a SAP or IP interface at one time. Attempts to associate a second QoS policy of a given type will return an error.

By default, no specific QoS policy is associated with the SAP or IP interface for ingress so the default QoS policy is used.

The normal behavior is for queues to be created per destination. Shared and multipoint shared change this behavior creating either unicast or unicast and mcast shared queues.

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

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

Values 1 — 65535

shared-queuing — This keyword can only be specified on SAP ingress. Specify the ingress shared queue policy used by a SAP. When the value of this object is null it means that the SAP will use individual ingress QoS queues, instead of the shared ones.

scheduler-policy

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

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

Description This command applies an existing scheduler policy to an ingress or egress scheduler used by SAP queues associated with this multi-service customer site. The schedulers defined in the scheduler policy can only be created once the customer site has been appropriately assigned to a chassis port, channel or slot. Scheduler policies are defined in the **config>qos>scheduler-policy** *scheduler-policy-name* context.

The **no** form of this command removes the configured ingress or egress scheduler policy from the multi-service customer site. When the policy is removed, the schedulers created due to the policy are removed also making them unavailable for the ingress SAP queues associated with the customer site. Queues that lose their parent scheduler association are deemed to be orphaned and are no longer subject to a virtual scheduler. The SAPs that have ingress queues reliant on the removed schedulers enter into an operational state depicting the orphaned status of one or more queues. When the **no scheduler-policy** command is executed, the customer site ingress or egress node will not contain an applied scheduler policy.

scheduler-policy-name: — The *scheduler-policy-name* parameter applies an existing scheduler policy that was created in the **config>qos>scheduler-policy** *scheduler-policy-name* context to create the hierarchy of ingress or egress virtual schedulers. The scheduler names defined within the policy are created and made available to any ingress or egress queues created on associated SAPs.

Values Any existing valid scheduler policy name.

host

Syntax [**no**] **host ip** *ip-address* [**mac** *ieee-address*]] [**subscriber** *sub-ident-string*] [**sub-profile** *sub-profile-name*] [**sla-profile** *sla-profile-name*]
no host {[**ip** *ip-address*] [**mac** *ieee-address*]}
no host all

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

Description This command creates a static subscriber host for the SAP. Static subscriber hosts may be used by the system for various purposes. Applications within the system that make use of static host entries include anti-spoof filters and ARP cache population.

Multiple static hosts may be defined on the SAP. Each host is identified by either a source IP address, a source MAC address or both a source IP and source MAC address. Every static host definition must have at least one address defined, IP or MAC.

Static hosts can exist on the SAP even with anti-spoof and ARP populate features disabled. When enabled, each feature has different requirements for static hosts.

anti-spoof — When enabled, this feature uses static and dynamic host information to populate entries into an anti-spoof filter table. The anti-spoof filter entries generated will be of the same type as specified in the anti-spoof type parameter. If the SAP anti-spoof filter is defined as **ip**, each static host definition must specify an IP address. If the SAP anti-spoof filter is defined as **ip-mac**, each static host definition must specify both an IP address and MAC address. If definition of a static host is attempted without the appropriate addresses specified for the enabled anti-spoof filter, the static host definition will fail.

arp-populate — When enabled, this feature uses static and dynamic host information to populate entries in the system ARP cache.

Attempting to define a static subscriber host that conflicts with an existing DHCP lease state table entry will fail.

Use the **no** form of the command to remove a static entry from the system. The specified *ip-address* and *mac-address* must match the host's exact IP and MAC addresses as defined when it was created. When a static host is removed from the SAP, the corresponding anti-spoof entry and/or ARP cache entry is also removed.

Default none

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

mac *mac-address* — Specify this optional parameter when defining a static host. The MAC address must be specified for **anti-spoof ip-mac** and **arp-populate**. Multiple static hosts may be configured with the same MAC address given that each definition is distinguished by a unique IP address.

subscriber *sub-ident-string* — Specify this optional parameter to specify an existing subscriber identification profile to be associated with the static subscriber host. The subscriber identification profile is configured in the **config>subscr-mgmt>sub-ident-policy** context. The subscriber information is used by the VPRN SAP arp-reply-agent to determine the proper handling of received ARP requests from subscribers.

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

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

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

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

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

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

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

ingress

Syntax **ingress**

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

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

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

multi-service-site

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

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

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

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

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

Default None — Each customer site must be explicitly created.

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

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

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

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

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

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

ATM Commands

atm

Syntax **atm**

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

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

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

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

egress

Syntax **egress**

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

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

encapsulation

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

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

Description This command configures RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*, encapsulation for an ATM PVCC delimited SAP.

This command specifies the data encapsulation for an ATM PVCC delimited SAP. The definition references RFC 2684 and to the ATM Forum LAN Emulation specification.

Ingress traffic that does not match the configured encapsulation will be dropped.

Default The encapsulation is driven by the services for which the SAP is configured.
For IES service SAPs, the default is **aal5snap-routed**.

Parameters *atm-encap-type* — Specify the encapsulation type.

Values **aal5snap-routed** — Routed encapsulation for LLC encapsulated circuit (LLC/SNAP precedes protocol datagram) as defined in RFC 2684.
aal5mux-ip — Routed IP encapsulation for VC multiplexed circuit as defined in RFC 2684

ingress

Syntax **ingress**

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

Description This command configures ingress ATM attributes for the SAP.

traffic-desc

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

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

Description This command assigns an ATM traffic descriptor profile to a given context (for example, a SAP). When configured under the ingress context, the specified traffic descriptor profile defines the traffic contract in the forward direction. When configured under the egress context, the specified traffic descriptor profile defines the traffic contract in the backward direction.

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

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

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

oam

Syntax **oam**

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

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

The ATM-capable MDAs support F5 end-to-end OAM functionality (AIS, RDI, Loopback):

- ITU-T Recommendation I.610 - B-ISDN Operation and Maintenance Principles and Functions version 11/95
- GR-1248-CORE - Generic Requirements for Operations of ATM Network Elements (NEs). Issue 3 June 1996
- GR-1113-CORE - Bellcore, Asynchronous Transfer Mode (ATM) and ATM Adaptation Layer (AAL) Protocols Generic Requirements, Issue 1, July 1994

alarm-cells

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

Context config>service>ies>sub-if>grp-if>sap>atm>oam
config>service>vprn>if>sap>atm>oam
config>service>vprn>sub-if>grp-if>sap>atm>oam

Description This command configures AIS/RDI fault management on a PVCC. Fault management allows PVCC termination to monitor and report the status of their connection by propagating fault information through the network and by driving PVCCs operational status.

When alarm-cells functionality is enabled, PVCCs operational status is affected when a PVCC goes into AIS or RDI state because of an AIS/RDI processing (i.e. assuming nothing else affects PVCCs operational status, PVCC goes DOWN, when it enters a fault state and comes back UP, when it exits that fault state) and RDI cell are generated when PVCC is operationally DOWN. No OAM-specific SNMP trap is raised whenever an endpoint enters/exits an AIS or RDI states, however, if as result of an OAM state change, the PVCC changes operational status, then a trap is expected from an entity the PVCC is associated with (for example a SAP).

The no command disables alarm-cells functionality for a PVCC. When alarm-cells functionality is disabled, PVCCs operational status is no longer affected by PVCCs OAM state changes due to AIS/RDI processing (Note that when alarm-cells is disabled, a PVCC will change operational status to UP, if it was DOWN because of the alarm-cell processing) and RDI cells are not generated as result of PVCC going into AIS or RDI state, however, PVCCs OAM status will record OAM faults as described above.

Default Enabled for PVCCs delimiting IES SAPs

periodic-loopback

Syntax [**no**] **periodic-loopback**

Context config>service>ies>sub-if>grp-if>sap>atm>oam
config>service>vprn>if >sap>atm>oam
config>service>vprn>sub-if>grp-if>sap>atm

Description This command enables periodic OAM loopbacks on this SAP. This command is only configurable on IES and VPRN SAPs. When enabled, an ATM OAM loopback cell is transmitted every period as configured in the `config>system>atm>oam>loopback-period` *`period`* context.

If a response is not received and consecutive retry-down retries also result in failure, the endpoint will transition to an alarm indication signal/loss of clock state. Then, an ATM OAM loopback cell will be transmitted every period as configured in the `loopback-period` *`period`*. If a response is received for the periodic loopback and consecutive retry-up retries also each receive a response, the endpoint will transition back to the up state.

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

Default no periodic-loopback

Redundant Interface Commands

redundant-interface

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

Context config>service>ies
config>service>vprn

Description This command configures a redundant interface.

Parameters *ip-int-name* — Specifies the name of the IP interface. Interface names can be from 1 to 32 alphanumeric characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

address

Syntax **address** {*ip-address/mask* | *ip-address netmask*} [**remote-ip** *ip-address*]
no address

Context config>service>vprn>redundant-interface

Description This command assigns an IP address mask or netmask and a remote IP address to the interface.

Parameters *ip-address/mask* — Assigns an IP address/IP subnet format to the interface.

ip-address netmask — Specifies a string of 0s and 1s that mask or screen out the network part of an IP address so that only the host computer part of the address remains.

Assigns an IP address netmask to the interface.

remote-ip *ip-address* — Assigns a remote IP to the interface.

spoke-sdp

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

Context config>service>vprn

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

A spoke SDP is treated like the equivalent of a traditional bridge "port" where flooded traffic received on the spoke SDP is replicated on all other "ports" (other spoke and mesh SDPs or SAPs) and not transmitted on the port it was received.

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

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

SDPs must be explicitly associated and bound to a service. If an SDP is not bound to a service, no far-end 7750 SRdevices can participate in the service.

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

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

Special Cases **VPRN —** Several SDPs can be bound to a VPRN service. Each SDP must be destined to a different 7750 SR router. If two *sdp-id* bindings terminate on the same 7750 SR, an error occurs and the second SDP binding is rejected.

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

vc-id — The virtual circuit identifier.

Values 1 — 4294967295

egress

Syntax **egress**

Context config>service>vprn>red-if>spoke-sdp

Description This command configures egress SDP parameters.

ingress

Syntax **ingress**

Context config>service>vprn>red-if>spoke-sdp

Description This command configures ingress SDP parameters.

vc-label

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

Context config>service>vprn>red-if>spoke-sdp>egress

Description This command configures the egress VC label.

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

Values 16 — 1048575

vc-label

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

Context config>service>vprn>red-if>spoke-sdp>ingress

Description This command configures the ingress VC label.

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

Values 2048 — 18431

filter

Syntax **filter** {**ip** *ip-filter-i*d}
no filter

Context config>service>vprn>red-if>spoke-sdp>ingress
config>service>vprn>red-if>spoke-sdp>egress

Description This command associates an IP filter policy with an ingress or egress Service Access Point (SAP) or IP interface. An IP filter policy can be associated with spoke SDPs.

Filter policies control the forwarding and dropping of packets based on IP or MAC matching criteria.

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

In general, filters applied to SAPs (ingress or egress) apply to all packets on the SAP. One exception is non-IP packets are not applied to IP match criteria, so the default action in the filter policy applies to these packets.

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

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

Values 1 — 65535

SDP Binding Commands

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 PW 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 PW 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 PW-port.

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

Default none

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

Values 1 — 10239

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

Values 1 — 4294967295

egress

Syntax **egress**

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

Description This command enables the context to configure PW-port egress side parameters.

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 PW-port as dot1q or qinq.

Default dot1q

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

qinq — Specifies **qinq** encapsulation type.

shaper

Syntax **shaper**
no shaper

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

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

Default no shaper.

int-dest-id

Syntax **int-dest-id** *int-dest-id*
no int-dest-id *int-dest-id*

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

Description This command specifies the intermediate destination string configured for dynamic vport selection.

The **no** form of the command removes the configured intermediate destination string.

This command is only valid for PW ports used for enhanced subscriber management (ESM on PW).

Default no .int-dest-id

Parameters *int-dest-id* — A text string that describes the intermediate destination ID.

pw-sap-secondary-shaper

Syntax [**no**] **pw-sap-secondary-shaper** *secondary-shaper-name*

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

Description This command configures the use of secondary shaper name as a reference to a shaper to use for a PW port on the HSMDA.

The **no** form of the command removes the configured shaper.

This command is valid for PW ports used PW SAPs on the HSMDA.

Default no pw-sap-secondary-shaper

Parameters *secondary-shaper-name* — specifies a text string representing the name of the secondary shaper.

vport

Syntax [**no**] **vport** *vport-name*

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

Description This command configures the name of the vport to be used for the PW port.

The **no** form of the command removes the configured vport name.

This command is valid for PW ports used for enhanced subscriber management (ESM on pseudowire) and pseudowire SAPs on Ethernet ports. It is not valid for pseudowire ports on the HSMDA.

Default no vport

Parameters *vport-name* — Specifies a text string representing the name of the vport.

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 PW-port. The vc-type is signaled to the peer, and must be configured consistently on both ends of the PW. vc-type VLAN is only configurable with dot1q encapsulation on the PW-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 PW, and a configured vlan-tag (for vc-type vlan) is inserted when forwarding into the PW. With vc-type ether, the tags if present (max 2), are transparently preserved when forwarding in our out of the PW.

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 PW.

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

Default 0

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

Values 0 — 4094

Category Map Commands

category-map

Syntax **category-map** *category-map-name* [**create**]
no category-map *category-map-name*

Context config>subscr-mgmt
config>subscr-mgmt>sla-prof

Description This command specifies the category map name.

Default none

Parameters *category-map-name* — Specifies the category map name up to 32 characters in length.

create — Mandatory keyword when creating a new category map.

credit-control-policy

Syntax **credit-control-policy** *policy-name* [**create**]
no credit-control-policy *policy-name*

Context config>subscr-mgmt

Description This command creates, configures or deletes a credit control policy.

Parameters *policy-name* — Specifies the policy name, 32 characters max.

credit-control-server

Syntax **credit-control-server radius**
credit-control-server diameter *policy-name*
no credit-control-server

Context config>subscr-mgmt>credit-control-policy

Description This command configures the credit control server to use. In case of RADIUS, the servers defined in the authentication policy are used. For Diameter, the peers defined in the specified Diameter policy are used.

Default no credit-control-server

Parameters **radius —** Use the RADIUS authentication servers defined in the RADIUS authentication policy in the group-interface to report credit usage and obtain new credit.

diameter *policy-name* **—** Use the diameter peers specified in the diameter **policy** *policy-name* to report credit usage and obtain new credit.

default-category-map

Syntax **default-category-map** *category-map-name*
no default-category-map

Context config>subscr-mgmt>credit-control-policy

Description This command configures the default category map.

Parameters *category-map-name* — Specifies the category map name, 32 chars max.

error-handling-action

Syntax **error-handling-action** {**continue | block**}
no error-handling-action

Context config>subscr-mgmt>credit-control-policy

Description This command configures the error handling action for the policy.

out-of-credit-action

Syntax **out-of-credit-action** {**continue | disconnect-host | block-category | change-service-level**}
no out-of-credit-action

Context config>subscr-mgmt>credit-control-policy

Description This command configures the action to be performed when out of credit is reached.

Parameters {**continue** | **disconnect-host** | **block-category** |**change-service-level**} — Specifies the action to be taken when out of credit is reached.

activity-threshold

Syntax **activity-threshold** *kilobits-per-second*
no activity-threshold

Context config>subscr-mgmt>cat-map

Description This command configures the threshold that is applied to determine whether or not there is activity. This is only valid for credit-type = time (not volume).

Default 0

Parameters *kilobits-per-second* — Specifies the activity threshold value in kilobits per second.

Values 1 — 100000000

category

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

Context config>subscr-mgmt>cat-map

Description This command specifies the category name.

Default none

Parameters *category-name* — Specifies the category name up to 32 characters in length.

create — Mandatory keyword when creating a new category.

category-map

Syntax **category-map** *category-map-name*
no category-map

Context config>subscr-mgmt>sla-prof

Description This command references the category-map to be used for the idle-timeout monitoring of subscriber hosts associated with this sla-profile. The **category-map** must already exist in the **config>subscr-mgmt** context.

Parameters *category-map-name* — Specifies the name of the category map (up to 32 characters in length) where the activity-threshold and the category is defined for idle-timeout monitoring of subscriber hosts.

category

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

Context config>subscr-mgmt>sla-prof>cat-map

Description This command defines the category in the category-map to be used for the idle-timeout monitoring of subscriber hosts.

Parameters *category-name* — Specifies the name (up to 32 characters in length) of the category where the queues and policers are defined for idle-timeout monitoring of subscriber hosts.

create — Mandatory keyword when creating a new category

idle-timeout

Syntax **idle-timeout** *timeout*
no idle-timeout

Context config>subscr-mgmt>sla-prof>cat-map>category

Description This command defines the idle-timeout value.

Default no idle-timeout – corresponds with an infinite idle-timeout

Parameters *timeout* — Specifies the idle-timeout in seconds.

Values 60 — 15552000

idle-timeout-action

Syntax **idle-timeout-action** {**shcv-check** | **terminate**}
no idle-timeout-action

Context config>subscr-mgmt>sla-prof>cat-map>category

Description This command defines the action to be executed when the idle-timeout is reached. The action is performed for all hosts associated with the sla-profile instance.

Default terminate

Parameters **shcv-check** — performs a subscriber host connectivity verification check (IPoE hosts only). Note that host connectivity verification must be enabled on the group-interface where the host is connected.

If the check is successful, the hosts are not disconnected and the idle-timeout timer is reset.

If the check fails, the hosts are deleted, similar as for "idle-timeout-action=terminate".

terminate — Deletes the subscriber host from the system: for PPP hosts, a terminate request is send; for IPoE hosts a DHCP release is send to the DHCP server.

credit-type-override

Syntax **credit-type-override** {**volume** | **time**}
no credit-type-override

Context config>subscr-mgmt>cat-map>category

Description This command overrides the **credit-type** configured in the **config>subscr-mgmt>cat-map** context for the given category.

Default no credit-type-override

Parameters **volume —** If different than the value specified in the **credit-type** command, the value overrides the credit-type.

time — If different than the value specified in the **credit-type** command, the value overrides the credit-type.

default-credit

Syntax **default-credit volume** *credits* **bytes | kilobytes | megabytes | gigabytes**
default-credit time *seconds*
no default-credit

Context config>subscr-mgmt>cat-map>category

Description This command configures the default credit used by this category.

Default no default-credit

Parameters **volume** *credits* **bytes|kilobytes|megabytes|gigabytes —** Specifies the default value for the volume credit and the unit in which the default value is expressed.

Values 1 — 4294967295 (minimum 100 megabytes)

time *seconds* **—** Specifies the default value for the time credit, in seconds.

Values 900 — 4294967295 (minimum 900 seconds)

exhausted-credit-service-level

Syntax [**no**] **exhausted-credit-service-level**

Context config>subscr-mgmt>cat-map>category

Description This command enables the context to configure the exhausted credit service level

Default exhausted-credit-service-level

egress-ip-filter-entries

Syntax [**no**] **egress-ip-filter-entries**

Context config>subscr-mgmt>cat-map>category>exh-lvl

Description This command configures the egress IP filter entries.

egress-ipv6-filter-entries

Syntax [**no**] **egress-ipv6-filter-entries**

Context config>subscr-mgmt>cat-map>category>exh-lvl

Description This command configures the egress IPv6 filter entries.

ingress-ip-filter-entries

Syntax [**no**] **ingress-ip-filter-entries**

Context config>subscr-mgmt>cat-map>category>exh-lvl

Description This command configures the ingress IP filter entries.

ingress-ipv6-filter-entries

Syntax [**no**] **ingress-ipv6-filter-entries**

Context config>subscr-mgmt>cat-map>category>exh-lvl

Description This command configures the ingress IPv6 filter entries.

pir

Syntax [**no**] **pir**

Context config>subscr-mgmt>cat-map>category>exh-lvl

Description This command configures the PIR.

entry

Syntax **entry** *entry-id* [**create**]

Context config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ip
config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ipv6
config>subscr-mgmt>cat-map>category>exh-lvl>egr-ip
config>subscr-mgmt>cat-map>category>exh-lvl>egr-ipv6

Description This command configures the IP filter entry.

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

Values 1..65535

action

Syntax **action drop**
action forward
action http-redirect *url*
no action

Context config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ip>entry
config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ipv6>entry

config>subscr-mgmt>cat-map>category>exh-lvl>egr-ip>entry
config>subscr-mgmt>cat-map>category>exh-lvl>egr-ipv6>entry

Description This command configures the action for the filter entry.

Parameters **drop —** Specifies to drop the IP filter entry.

forward — Specifies to forward the IP filter entry.

http-redirect *url* **—** Specifies the HTTP web address that will be sent to the user's browser. Note that http-redirect is not supported on 7750 SR-1 or 7450 ESS-1 models.

The following displays information that can optionally be added as variables in the portal URL (http-redirect url):

- $IP – Customer's IP address
- $MAC – Customer's MAC address
- $URL – Original requested URL
- $SAP – Customer's SAP
- $SUB – Customer's subscriber identification string"

Values 255 characters maximum

match

Syntax **match** [**protocol** *protocol-id*]
no match

Context config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ip>entry
config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ipv6>entry
config>subscr-mgmt>cat-map>category>exh-lvl>egr-ip>entry
config>subscr-mgmt>cat-map>category>exh-lvl>egr-ipv6>entry

Description This command configures the match criteria for this IP filter entry.

Parameters *protocol-id* — Specifies the protocol number accepted in DHB.

Values 0..255

dscp

Syntax **dscp** *dscp-name*
no dscp

Context config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ip>entry>match
config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ipv6>entry>match
config>subscr-mgmt>cat-map>category>exh-lvl>egr-ip>entry>match
config>subscr-mgmt>cat-map>category>exh-lvl>egr-ipv6>entry>match

Description This command configures DSCP match conditions.

Parameters *dscp-name* — Specifies the DSCP name.

Values 32 chars max

dst-ip

Syntax **dst-ip** {*ip-address/mask* | *ip-address netmask*}
no dst-ip

Context config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ip>entry>match
config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ipv6>entry>match

Description This command configures the destination IP match condition.

Parameters *ip-address/mask* — Specifies the IPv4 address and mask.

Values ip-address a.b.c.d
mask 0..32

ipv6-address/prefix-length — Specifies the IPv6 address and length.

Values ipv6-address x:x:x:x:x:x:x:x (where x is [0..FFFFH])
x:x:x:x:x:x:d.d.d.d (where d is [0..255]D)

prefix-length — Specifies the prefix length.

Values 1..128

netmask — Specifies the mask, expressed as a dotted quad.

Values a.b.c.d

dst-port

Syntax **dst-port** {**lt | gt | eq**} *dst-port-number*
dst-port range *start end*
no dst-port

Context config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ip>entry>match
config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ipv6>entry>match
config>subscr-mgmt>cat-map>category>exh-lvl>egr-ip>entry>match
config>subscr-mgmt>cat-map>category>exh-lvl>egr-ipv6>entry>match

Description This command configures the destination port match condition.

Parameters *lt*|*gt*|*eq* — Specifies the operator.

dst-port-number — Specifies the destination port number as a decimal hex or binary.

Values 0..65535

fragment

Syntax **fragment** {**true | false**}

Context config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ip>entry>match
config>subscr-mgmt>cat-map>category>exh-lvl>egr-ip>entry>match

Description This command configures the fragmentation match condition.

Parameters **true**|**false** — Sets/resets fragmentation check.

icmp-code

Syntax **icmp-code** *icmp-code*
no icmp-code

Context config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ip>entry>match
config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ipv6>entry>match
config>subscr-mgmt>cat-map>category>exh-lvl>egr-ip>entry>match
config>subscr-mgmt>cat-map>category>exh-lvl>egr-ipv6>entry>match

Description This command configures the ICMP code match condition.

Parameters *icmp-code* — Specifies the ICMP code numbers accepted in DHB.

Values 0..255

icmp-type

Syntax **icmp-type** *icmp-type*
no icmp-type

Context config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ip>entry>match
config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ipv6>entry>match
config>subscr-mgmt>cat-map>category>exh-lvl>egr-ip>entry>match
config>subscr-mgmt>cat-map>category>exh-lvl>egr-ipv6>entry>match

Description This command configures the ICMP type match condition.

Parameters *icmp-type* — Specifies the ICMP type numbers accepted in DHB.

Values 0..255

ip-option

Syntax **ip-option** *ip-option-value* [*ip-option-mask*]
no ip-option

Context config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ip>entry>match
config>subscr-mgmt>cat-map>category>exh-lvl>egr-ip>entry>match

Description This command configures the IP option match condition.

Parameters *ip-option-value* — Specifies the IP option value as a decimal hex or binary.

Values 0..255

ip-option-mask — Specifies the IP opition mask as a decimal hex or binary.

Values 0..255

multiple-option

Syntax **multiple-option** {**true | false**}

Context config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ip>entry>match
config>subscr-mgmt>cat-map>category>exh-lvl>egr-ip>entry>match

Description This command configures the multiple-option match condition.

Parameters **true**|**false —** Sets or resets the multiple option check.

option-present

Syntax **option-present** {**true | false**}

Context config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ip>entry>match
config>subscr-mgmt>cat-map>category>exh-lvl>egr-ip>entry>match

Description This command configures the option-present match condition.

Parameters **true | false —** Sets or resets the option present check.

src-ip

Syntax **src-ip** {*ip-address/mask* | *ip-address netmask*}
no src-ip

Context config>subscr-mgmt>cat-map>category>exh-lvl>egr-ip>entry>match
config>subscr-mgmt>cat-map>category>exh-lvl>egr-ipv6>entry>match

Description This command configures the source IP match condition.

Parameters *ip-address/mask* — Specifies the IPv4 address and mask.

Values

ip-address	a.b.c.d
mask	0 — 32

netmask — Specifies the mask, expressed as a dotted quad.

Values a.b.c.d

ipv6-address/prefix-length — Specifies the IPv6 address and length.

Values ipv6-address x:x:x:x:x:x:x:x (where x is [0..FFFFH])
x:x:x:x:x:x:d.d.d.d (where d is [0..255]D)

prefix-length — Specifies the prefix length.

Values 1..128

src-port

Syntax **src-port** {**lt | gt | eq**} *src-port-number*
src-port range *start end*
no src-port

Context config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ip>entry>match
config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ipv6>entry>match
config>subscr-mgmt>cat-map>category>exh-lvl>egr-ip>entry>match
config>subscr-mgmt>cat-map>category>exh-lvl>egr-ipv6>entry>match

Description This command configures the source port match condition.

Parameters **lt|gt|eq —** Specifies the operators.

src-port-number — Specifies the source port number as a decimal hex or binary.

Values 0..65535

dst-port-number — Specifies the destination port number as a decimal hex or binary.

Values 0..65535

tcp-ack

Syntax **tcp-ack** {**true | false**}
no tcp-ack

Context config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ip>entry>match
config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ipv6>entry>match
config>subscr-mgmt>cat-map>category>exh-lvl>egr-ip>entry>match
config>subscr-mgmt>cat-map>category>exh-lvl>egr-ipv6>entry>match

Description This command configures the TCP ACK match condition. The **no** tcp-ack command disables the checking on the presence or absence of the tcp-ack flag.

Parameters **true|false —** True|false indicates that the entry will match on the presence resp. absence of the tcp-ack flag in the received packet. .

tcp-syn

Syntax **tcp-syn** {**true | false**}
no tcp-syn

Context config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ip>entry>match
config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ipv6>entry>match
config>subscr-mgmt>cat-map>category>exh-lvl>egr-ip>entry>match
config>subscr-mgmt>cat-map>category>exh-lvl>egr-ipv6>entry>match

Description This command configures the TCP SYN match condition. The **no** tcp-syn command disables the checking on the presence or absence of the tcp-syn flag.

Parameters **true|false —** True|false indicates that the entry will match on the presence resp. absence of the tcp-syn flag in the received packet.

pir

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

Context config>subscr-mgmt>cat-map>category>svc-lvl

Description This command configures the PIR which will be enforced for all queues pertaining to this category.

Default no pir

Parameters *pir-rate* — Specifies the amount of bandwidth in kilobits per second (thousand bits per second).

Values 1 — 40000000

max — Specifies to use the maximum amount of bandwidth.

out-of-credit-action-override

Syntax **out-of-credit-action-override** {**continue** | **block-category** | **change-service-level**}
no out-of-credit-action-override

Context config>subscr-mgmt>cat-map>category

Description This command specifies the action to be taken if the credit is exhausted.

Default no out-of-credit-action-override

Parameters **continue —** Specifies to continue when running out of credit.

block-category — Specifies to block the category when running out of credit.

change-service-level — Specifies to change the service level when running out of credit.

queue

Syntax **queue** *queue-id* {**ingress-only** | **egress-only** | **ingress-egress**}
no queue *queue-id*

Context config>subscr-mgmt>cat-map>category

Description This command configures a queue in this category.

Default none

Parameters *queue-id* — Specifies the queue ID for this instances. Each queue nominated in the category map is monitored for activity (over a period of approximately 60 seconds), should the activity fall below the threshold value then a time is started. Whenever this timer exceeds the configured timeout under the idle-timeout the action (currently disconnect) is executed for that subscriber and all hosts under that given SLA-profile-instance.

Values 1 — 32

ingress-only — Specifies that ingress queues are defined in this category.

egress-only — Specifies that egress queues are defined in this category.

ingress-egress — Specifies that ingress and egress queues are defined in this category.

rating-group

Syntax **rating-group** *rating-group-id*
no rating-group

Context config>subscr-mgmt>cat-map>category

Description This command configures the rating group applicable for this category.

Default no rating group

Parameters *rating-group-id* — Specifies the rating group applicable for this category.

credit-exhaust-threshold

credit-exhaust-threshold *threshold-percentage*
no credit-exhaust-threshold

Context config>subscr-mgmt>cat-map

Description This command specifies the credit exhaust threshold taken into account to take action as depicted in Figure 129.

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

Default 100

Parameters *threshold-percentage* — Specifies the percent to use for the credit exhaust threshold.

Values 50 — 100

credit-type

Syntax **credit-type** {**volume** | **time**}
no credit-type

Context config>subscr-mgmt>cat-map

Description This command specifies whether volume or time based accounting is performed.

Default volume

Parameters **volume —** specifies volume-based accounting.

time — Specifies time-based accounting.

Vport Commands

ethernet

Syntax **ethernet**

Context config>port

Description This command enables access to the context to configure Ethernet port attributes.

This context can only be used when configuring Fast Ethernet, gigabit or 10Gig Fast Ethernet or Ethernet LAN ports on an appropriate MDA.

egress-scheduler-override

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

Context config>port>ethernet

Description This command applies egress scheduler overrides. When a port scheduler is associated with an egress port, it is possible to override the following parameters:

- The **max-rate** allowed for the scheduler.
- The maximum **rate** for each priority level 8 through 1.
- The CIR associated with each priority level 8 through 1.

See the SR OS Quality of Service Guide for command syntax and usage for the **port-scheduler-policy** command.

The **no** form of this command removes all override parameters from the egress port or channel scheduler context. Once removed, the port scheduler reverts all rate parameters back to the parameters defined on the port-scheduler-policy associated with the port.

level

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

Context config>port>ethernet>egress-scheduler-override

Description This command overrides the maximum and CIR rate parameters for a specific priority level on the port or channel's port scheduler instance. When the **level** command is executed for a priority level, the corresponding priority level command in the port-scheduler-policy associated with the port is ignored.
The override level command supports the keyword **max** for the **rate** and **cir** parameter.
When executing the level override command, at least the **rate** or **cir** keywords and associated parameters must be specified for the command to succeed.

The **no** form of this command removes the local port priority level rate overrides. Once removed, the port priority level will use the port scheduler policies level command for that priority level.

Parameters *priority-level* — Identifies which of the eight port priority levels are being overridden.

Values 1 — 8

rate *pir-rate* **—** Overrides the port scheduler policy's maximum level rate and requires either the **max** keyword or a rate defined in kilobits-per-second to follow.

Values 1 — 40000000, max

cir *cir-rate* **—** Overrides the port scheduler policy's within-cir level rate and requires either the max keyword or a rate defined in kilobits-per-second to follow.

Values 0— 40000000, max

max — removes any existing rate limit imposed by the port scheduler policy for the priority level allowing it to use as much total bandwidth as possible.

access

Syntax **access**

Context config>port>ethernet

Description This command configures Ethernet access port parameters.

egress

Syntax **egress**

Context config>port>ethernet>access

Description This command configures Ethernet access egress port parameters.

vport

Syntax **vport** *name* [**create**]
no vport *name*

Context config>port>ethernet>access>egress

Description This command configures a scheduling node, referred to as virtual port, within the context of an egress Ethernet port. The vport scheduler operates either like a port scheduler with the difference that multiple vport objects can be configured on the egress context of an Ethernet port, or it can be an aggregate rate when an egress port-scheduler policy is applied to the port.

The vport is always configured at the port level even when a port is a member of a LAG.

When a port scheduler policy is applied to a vport the following command is used:

configure>port>ethernet>acess>egress>vport>port-scheduler-policy *port-scheduler-policy-name*

The CLI will not allow the user to apply a port scheduler policy to a vport if one has been applied to the port. Conversely, the CLI will not allow the user to apply a port scheduler policy to the egress of an Ethernet port if one has been applied to any vport defined on the access egress context of this port. The agg-rate-limit, along with an egress port-scheduler, can be used to ensure that a given vport does not oversubscribe the port's rate.

SAP and subscriber host queues can be port-parented to a vport scheduler in a similar way they port-parent to a port scheduler or can be port-parented directly to the egress port-scheduler if the agg-rate-limit is used.

When the vport uses an aggregate rate, the following command is used:

configure>port>ethernet>acess>egress>vport>agg-rate-limit

Parameters *name* — Specifies the name of the vport scheduling node and can be up to 32 ASCII characters in length. This does not need to be unique within the system but is unique within the port or a LAG.

agg-rate-limit

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

Context configure>port>ethernet>access>egress>vport

Description This command configures an aggregate rate for the vport. This command is mutually exclusive with the port-scheduler-policy command.

Parameters *agg-rate* — Specifies the rate limit for the vport.

Values **max**, 1— 10000000

egress-rate-modify

Syntax [**no**] **egress-rate-modify**

Context configure>port>ethernet>access>egress>vport

Description This command is used to apply HQoS Adjustment to a vport. HQoS Adjustment refers to the dynamic adjustment of the rate limit at an QoS enforcement point within 7x50 when the multicast traffic stream is disjointed from the unicast traffic stream. This QoS enforcement point within 7x50 represents the physical point further down in the access part of the network where the two streams join each other and potentially can cause congestion.

An example would be a PON port which is shared amongst subscriber's multicast traffic (single copy of each channel) and subscriber's unicast traffic. The bandwidth control point for this PON port resides in the upstream 7x50 BNG node in the form of a vport. In case that the multicast delivery method in the 7x50 BNG utilizes redirection, the multicast traffic in the 7x50 BNG will flow outside of the subscriber or the vport context and thus will bypass any bandwidth enforcement in 7x50. To correct this, a vport bandwidth adjustment is necessary in 7x50 that will account for the multicast bandwidth consumption that is bypassing vport in 7x50 but is present in the PON port whose bandwidth is controlled by vport.

An estimate of the multicast bandwidth consumption on the PON port can be made at the vport level based on the IGMP messages sourced from the subscribers behind the PON port. This process is called HQoS Adjustment.

A multicast channel bandwidth is subtracted from or added to the vport rate limit according to the received IGMP Join/Leave messages and the channel bandwidth definition policy associated with the vport (indirectly through a group-interface). Since the multicast traffic on the PON port is shared amongst subscribers behind this PON port, only the first IGMP Join or the last IGMP Leave per multicast channel is tracked for the purpose of the vport bandwidth modification.

The vport rate that will be affected by this functionality depends on the configuration:

- In case the agg-rate-limit within the vport is configured, its value will be modified based on the IGMP activity associated with the subscriber under this vport.
- In case the port-scheduler-policy within the vport is referenced, the max-rate defined in the corresponding port-scheduler-policy will be modified based on the IGMP activity associated with the subscriber under this vport.

The channel bandwidth definition policy is defined in the mcac policy in the **configure>router>mcac>policy** context. The policy is applied under the group-interface or in case of redirection under the redirected-interface.

The rates in effect can be displayed with the following two commands:

show port 1/1/5 vport *name*

qos scheduler-hierarchy port *port-id* vport *vport-name*

Context HQoS Adjustment for vport is disabled.

host-match

Syntax **host-match dest** *destination-string* [**create**]
no host-match dest *destination-string*

Context config>port>ethernet>access>egr>qgrp

Description This command configures host matching for the Ethernet port egress queue-group.

The no form of the command removes

Parameters **dest** *destination-string* — Specify a host match destination string up to 32 characters in length.

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

port-scheduler-policy

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

Context config>port>ethernet>access>egress>vport

Description This command specifies the destination and organization strings to be used for matching subscriber hosts with this vport.

The parent vport of a subscriber host queue, which has the port-parent option enabled, is determined by matching the destination string dest string associated with the subscriber and the organization string org string associated with the subscriber host with the strings defined under a vport on the port associated with the subscriber.

If a given subscriber host queue does not have the port-parent option enabled, it will be foster-parented to the vport used by this subscriber and which is based on matching the dest string and org string. If the subscriber could not be matched with a vport on the egress port, the host queue will not be bandwidth controlled and will compete for bandwidth directly based on its own PIR and CIR parameters.

By default, a subscriber host queue with the port-parent option enabled is scheduled within the context of the port's port scheduler policy.

The no form of the command removes the port-scheduler-policy-name from the configuration.

Parameters *port-scheduler-policy-name* — Specifies an existing port-scheduler-policy configured in the **config>qos** context.

scheduler-policy

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

Context config>port>ethernet>access>egress>vport

Description This command associates a virtual scheduler policy with a port queue group. 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 queue-group.

Parameters *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.

Show Commands

radius-configuration

Syntax **radius-configuration**

Context show>aaa

Description This command displays RADIUS configuration information.

Sample Output

```
# show aaa radius-configuration
===============================================================================
RADIUS configuration
===============================================================================
CoA Port                    : 3799
===============================================================================
```

ancp-policy

Syntax **ancp-policy** [*policy-name*]
ancp-policy *policy-name* **association**

Context show>subscr-mgmt

Description This command displays subscriber Access Node Control Protocol (ANCP) policy information.

Parameters *policy-name* — Displays information for the specified ANCP policy.

association — Displays the information configured with the ANCP policy.

Sample Output

```
A:cses-E11>config>subscr-mgmt>ancp# show subscriber-mgmt ancp-policy "test"
===============================================================================
ANCP Policy "test"
===============================================================================
I. Rate Reduction      : 0 kbps
I. Rate Adjustment     : 100 percent
I. Rate Monitor        : 63360 kbps
I. Rate Monitor Alarm  : Yes
I. Rate Modify         : N/A

E. Rate Reduction      : 0 kbps
E. Rate Adjustment     : 100 percent
E. Rate Monitor        : 0 kbps
E. Rate Monitor Alarm  : no
E. Rate Modify         : N/A
```

```
Port Down : N/A

Last Mgmt Change: 02/13/2013 19:15:28
===============================================================================
*A:cses-E11>config>subscr-mgmt>ancp#
```

ancp-string

Syntax **ancp-string**
ancp-string *ancp-string*
ancp-string customer *customer-id* **site** *customer-site-name*
ancp-string sap *sap-id*

Context show>subscr-mgmt

Description This command displays subscriber Access Node Control Protocol (ANCP) string information.

Parameters *ancp-string* — Specifies an Access Node Control Protocol (ANCP) string up to 63 characters in length.

customer *customer-id* — Specifies an existing customer ID.

Values 1..2147483647

site *customer-site-name* — Specifies an existing customer site name up to 32 characters in length.

sap *sap-id* — Displays ANCP string information for the specified SAP ID.

Values

<sap-id>	null	<port-id|bundle-id|bpgrp-id|lag-id|aps-id>
	dot1q	<port-id|bundle-id|bpgrp-id|lag-id|aps-id|pw-id>:qtag1
	qinq	<port-id|bundle-id|bpgrp-id|lag-id| pw-id>:qtag1.qtag2
	atm	<port-id|aps-id>[:vpi/vci|vpi|vpi1.vpi2|cp.conn-prof-id]
		cp - keyword
		conn-prof-id - [1..8000]
	frame	<port-id|aps-id>:dlci
	cisco-hdlc	slot/mda/port.channel
	cem	slot/mda/port.channel
	ima-grp	<bundle-id>[:vpi/vci|vpi|vpi1.vpi2|cp.conn-prof-id]
		cp keyword
		conn-prof-id [1..8000]
	port-id	slot/mda/port[.channel]
	bundle-id	bundle-<type>-slot/mda.<bundle-num>
		bundle keyword
		type ima|fr|ppp
		bundle-num [1..336]
	bpgrp-id	bpgrp-<type>-<bpgrp-num>
		bpgrp keyword
		type ima|ppp
		bpgrp-num [1..2000]
	aps-id	aps-<group-id>[.channel]
		aps keyword
		group-id [1..64]
	ccag-id	ccag-<id>.<path-id>[cc-type]:<cc-id>

	ccag keyword
	id [1..8]
	path-id [a|b]
	cc-type [.sap-net|.net-sap]
	cc-id [0..4094]
eth-tunnel	eth-tunnel-<id>[:<eth-tun-sap-id>]
	id [1..1024]
	eth-tun-sap-id [0..4094]
lag-id	lag-<id>
	lag keyword
	id [1..200]
pw-id	pw-<id>
	pw keyword
	id [1..10239]
qtag1	[0..4094]
qtag2	[*|0..4094]
vpi	[0..4095] (NNI)
	[0..255] (UNI)
vci	[1|2|5..65535]
dlci	[16..1022]
tunnel-id	tunnel-<id>.<private|public>:<tag>
	tunnel keyword
id	[1..16]
tag	[0..4094]

Sample Output

```
show subscriber-mgmt ancp-string "ANCP-0000003-0000001"
===============================================================================
ANCP-String "ANCP-0000003-0000001"
===============================================================================
Type       : SUB - "4AACAHCU74"
State      : Up                Ancp Policy: N/A
I. Rate    : 129 kbps          E. Rate    : 130 kbps
Adj I. Rate: N/A               Adj E. Rate: N/A
Act I. Rate: N/A               Act E. Rate: N/A
Service Id : 50 (VPRN)
Group      : linux
Neighbor   : 10.0.0.2:34885
Persist Key: N/A
-------------------------------------------------------------------------------
Actual-Net-Data-Rate-Upstream              : 129 kbits/s
Actual-Net-Data-Rate-Downstream            : 130 kbits/s
Minimum-Net-Data-Rate-Upstream             : 131 kbits/s
Minimum-Net-Data-Rate-Downstream           : 132 kbits/s
Attainable-Net-Data-Rate-Upstream          : 133 kbits/s
Attainable-Net-Data-Rate-Downstream        : 134 kbits/s
Maximum-Net-Data-Rate-Upstream             : 135 kbits/s
Maximum-Net-Data-Rate-Downstream           : 136 kbits/s
Minimum-Net-Low-Power-Data-Rate-Upstream   : 137 kbits/s
Minimum-Net-Low-Power-Data-Rate-Downstream : 138 kbits/s
Maximum-Interleaving-Delay-Upstream        : 139 ms
Actual-Interleaving-Delay-Upstream         : 140 ms
Maximum-Interleaving-Delay-Downstream      : 141 ms
Actual-Interleaving-Delay-Downstream       : 142 ms
```

```
DSL-Line-State                              : 2 (IDLE)
Access-Loop-Encapsulation                   : 16909056 (0x01020300)
===============================================================================
```

authentication

Syntax **authentication** *policy-name* **association**
authentication [*policy-name*]
authentication [*policy-name*] **statistics**
authentication coa-statistics

Context show>subscr-mgmt

Description This command displays subscriber management RADIUS authentication policy information and statistics.

Parameters *policy-name* — Specifies the subscriber management RADIUS authentication policy name, up to 32 characters, for which information is requested.

association — Displays SAP, interface, local user database host, AA and L2TP associations of this policy.

coa-statistics — Displays the overall statistics for incoming RADIUS Change of Authorization (CoA) messages and Disconnect Messages. For dropped requests, a counter for different drop reasons is available.

statistics — Displays a list of policies with basic statistics (without specifying a policy name) or detailed statistics, including per-server statistics for the specified policy-name. These statistics apply only to the legacy RADIUS server configuration method where the servers are directly configured in the authentication policy.

Sample Output

```
# show subscriber-mgmt authentication
===============================================================================
Authentication Policies
===============================================================================
Name                            Description
-------------------------------------------------------------------------------
auth-policy-1                   Radius auth policy - servers
auth-policy-2                   Radius auth policy - radius-server-policy
-------------------------------------------------------------------------------
Number of Authentication Policies : 2
===============================================================================


# show subscriber-mgmt authentication "auth-policy-2"
===============================================================================
Authentication Policy auth-policy-2
===============================================================================
Description          : Radius auth policy - radius-server-policy
Re-authentication    : Yes                 Username Format      : MAC Address
PPPoE Access Method  : PAP/CHAP            Username Mac-Format  : "aa:"
PPP-Username Oper    : None
```

```
PPP-Domain-Name      : N/A
Username Oper        : None
Domain-Name          : N/A
Acct-Stop-On-Fail    :
RADIUS Server Policy : "aaa-server-policy-1"
Fallback Action      : deny
Last Mgmt Change     : 06/24/2013 21:16:50
-------------------------------------------------------------------------------
Include Radius Attributes
-------------------------------------------------------------------------------
Remote Id            : Yes                 Circuit Id           : Yes
NAS Port Id          : Yes                 NAS Identifier       : Yes
PPPoE Service Name   : Yes                 DHCP Vendor Class Id : Yes
Access Loop Options  : Yes                 MAC Address          : Yes
NAS Port Prefix      : None                NAS Port Suffix      : None
NAS-Port-Type        : Yes (standard)      Acct Session Id      : Host
Calling Station Id   : Yes (sap-string)    Called Station Id    : Yes
Tunnel Server Attr   : Yes                 DHCP Options         : Yes
NAS Port             : Yes
NAS Port Bits Spec   : *3s*1m*4p*12o*12i
-------------------------------------------------------------------------------
Radius Servers
-------------------------------------------------------------------------------
Router               : management + Base   Source Address       : N/A
Access Algorithm     : Direct              Retry                : 3
Timeout (s)          : 5                   Hold down time (s)   : 30
-------------------------------------------------------------------------------
Index IP Address      Port  Pend-Req-Limit Out/Overload time (s) Oper State
-------------------------------------------------------------------------------
No Radius Servers configured.
-------------------------------------------------------------------------------
Accept Radius Attributes
-------------------------------------------------------------------------------
No Matching Entries
-------------------------------------------------------------------------------
Radius Script Policies
-------------------------------------------------------------------------------
Access-Request          : "N/A"
Access-Accept           : "N/A"
Change-of-Authorization : "N/A"
===============================================================================


# show subscriber-mgmt authentication "auth-policy-2" association
===============================================================================
Authentication Policy auth-policy-2
===============================================================================
-------------------------------------------------------------------------------
SAP Associations
-------------------------------------------------------------------------------
No associations found.
-------------------------------------------------------------------------------
Interface Associations
-------------------------------------------------------------------------------
Service-Id : 3000 (VPRN)
 - If Name : group-int-ws-1-1
-------------------------------------------------------------------------------
Local-User-Db PPPoE Host Associations
-------------------------------------------------------------------------------
Local-User-Db : ludb-1
 - Host : host-1
```

```
-------------------------------------------------------------------------------
Local-User-Db DHCP Host Associations
-------------------------------------------------------------------------------
Local-User-Db : ludb-1
 - Host : default
-------------------------------------------------------------------------------
Application Assurance Associations
-------------------------------------------------------------------------------
No associations found.
===============================================================================
No associated L2TP groups found.
No associated L2TP tunnels found.


# show subscriber-mgmt authentication statistics
===============================================================================
Authentication Policy Statistics
===============================================================================
Policy Name                       Subscr. Pkts   Subscr. Pkts   Subscr. Pkts
                                  Authenticated  Rejected       Rejected
                                                                Send Failed
-------------------------------------------------------------------------------
auth-policy-1                     0              0              0
auth-policy-2                     0              0              0
-------------------------------------------------------------------------------
Number of Authentication Policies : 2
===============================================================================


# show subscriber-mgmt authentication "auth-policy-1" statistics
===============================================================================
Authentication Policy Statistics
===============================================================================
-------------------------------------------------------------------------------
Policy name                              : auth-policy-1
subscriber packets authenticated         : 0
subscriber packets rejected              : 0
subscriber packets rejected send failed  : 0
-------------------------------------------------------------------------------
radius server     requests  requests  requests  requests   requests requests
idx IP-address    accepted  rejected  no reply  md5 failed pending  send failed
-------------------------------------------------------------------------------
1 172.16.1.1      0         0         0         0          0        0
-------------------------------------------------------------------------------
===============================================================================


# show subscriber-mgmt authentication coa-statistics
===============================================================================
Radius Notify Statistics    Change-Of-Authorization   Disconnect-Messages
===============================================================================
Requests Received             7                         10
Requests Accepted             5                         6
Requests Rejected             2                         4
Requests Dropped              0                         0
    No Auth Policy found      0                         0
    Invalid message           0                         0
    Out of resources          0                         0
    Authentication failure    0                         0
===============================================================================
```

I

explicit-subscriber-map

Syntax **explicit-subscriber-map**

Context show>subscriber-mgmt

Description This command displays explicit subscriber mappings.

Sample Output

```
B:Dut-A>show>subscr-mgmt# explicit-subscriber-map
===============================================================================
Explicit Subscriber Map
===============================================================================
Key                              Sub profile
                                 SLA profile
-------------------------------------------------------------------------------
sub_ident_A_1                    sub_prof80
                                 sla_prof80
-------------------------------------------------------------------------------
Number of Explicit Subscriber Mappings : 1
===============================================================================
B:Dut-A>show>subscr-mgmt#
```

host-lockout-policy

Syntax **host-lockout-policy**
host-lockout-policy *policy-name* **association**
host-lockout-policy *policy-name*
host-lockout-policy *policy-name* **all**
host-lockout-policy *policy-name* **sap** *sap-id* [**circuit-id** | **mac** | **remote-id**]

Context show>subscriber-mgmt

Description This command displays host lockout policy information.

Parameters *policy-name* —

association —

all —

sap *sap-id* —

circuit-id —

mac —

remote-id —

Sample Output

```
*A:cses-E11# show subscriber-mgmt host-lockout-policy
===============================================================================
Host Lockout Policies
===============================================================================
Lockout Policy                   Last Mgmt Change
  Lockout Time Min                 Lockout Time Max
Description
  Lockout Reset Time               Max Lockout Hosts
-------------------------------------------------------------------------------
test                             04/20/2012 19:51:02
  10                               3600
test
  60                               100
===============================================================================
*A:cses-E11#


*A:cses-E11# show subscriber-mgmt host-lockout-policy "test"
===============================================================================
Host Lockout Policy "test"
===============================================================================
Description                        : test
Last Mgmt Change                   : 04/20/2012 19:51:02
Lockout time min                   : 10
Lockout time max                   : 3600
Lockout reset time                 : 60
Max lockout hosts                  : 100
===============================================================================
*A:cses-E11#
```

igmp-policy

Syntax **igmp-policy**
igmp-policy *policy-name* **association**
igmp-policy *policy-name*

Context show>subscriber-mgmt

Description This command displays IGMP policy information.

Parameters *policy-name* — Specifies an existing IGMP policy.

association — Displays the information configured with the IGMP policy.

Sample Output

```
*B:Dut-C# show subscriber-mgmt igmp-policy
===============================================================================
IGMP Policies
===============================================================================
IGMP Policy
  Import Policy                    Admin Version
Description
  Num Subscribers                  Host Max Groups
```

```
  Fast Leave
-------------------------------------------------------------------------------
pol1
                                  3
  2                               0
  fast-leave
pol2
                                  3
  0                               0
  fast-leave
===============================================================================
*B:Dut-C#


*B:Dut-C# show subscriber-mgmt igmp-policy "pol1"
===============================================================================
IGMP Policy pol1
===============================================================================
Import Policy                     :
Admin Version                     : 3
Num Subscribers                   : 2
Host Max Group                    : 0
Fast Leave                        : yes
===============================================================================
*B:Dut-C#
```

```
*B:Dut-C# show subscriber-mgmt igmp-policy "pol1" association
===============================================================================
IGMP Policy pol1 Associations
===============================================================================
sub_1
sub_2
-------------------------------------------------------------------------------
No. of subscriber(s): 2
===============================================================================
*B:Dut-C#
```

local-user-db

Syntax **local-user-db** *local-user-db-name* **association** [**dhcp**] [**ppp**]
local-user-db *local-user-db-name* **dhcp-all-hosts**
local-user-db *local-user-db-name* **dhcp-host** *dhcp-host-name*
local-user-db *local-user-db-name* **dhcp-unmatched-hosts**
local-user-db [*local-user-db-name*]
local-user-db *local-user-db-name* **pppoe-all-hosts**
local-user-db *local-user-db-name* **pppoe-host** *pppoe-host-name*
local-user-db *local-user-db-name* **pppoe-unmatched-hosts**

Context show>subscriber-mgmt

Description This command displays local user database information.

Sample Output

```
*A:ALA-48>show>subscr-mgmt# local-user-db
===============================================================================
Local User Databases
===============================================================================
Name                              Admin Host  Description
                                  State Count
-------------------------------------------------------------------------------
database01                        Down  1
database02 Provider001/Class0002 Down  0     This is a long testdescription wi*
test                              Down  2
-------------------------------------------------------------------------------
Number of Local User Databases : 3     Number of Hosts : 3
===============================================================================
* indicates that the corresponding row element may have been truncated.
*A:ALA-48>show>subscr-mgmt#
```

```
*A:ALA-48>show>subscr-mgmt# local-user-db database01
===============================================================================
Local User Database "database01"
===============================================================================
Admin State          : Down
Last Mgmt Change     : 11/08/2007 12:27:36
Host Count           : 1
DHCP Match Types     : circ-id
DHCP CircId Mask Pfx : test
DHCP CircId Mask Sfx : N/A
PPPoE Match Types    : N/A
```

```
PPPoE CircId Mask Pfx: N/A
PPPoE CircId Mask Sfx: N/A
===============================================================================
*A:ALA-48>show>subscr-mgmt#


*A:ALA-48>show>subscr-mgmt# local-user-db database01 dhcp-all-hosts
===============================================================================
Local User Database "database01" DHCP hosts
===============================================================================
Name                            Admin     Matched objects
                                State
-------------------------------------------------------------------------------
host001                         Down      -
-------------------------------------------------------------------------------
Number of DHCP Hosts : 1
===============================================================================
*A:ALA-48>show>subscr-mgmt#


*A:ALA-48>show>subscr-mgmt# local-user-db "database01" dhcp-host host001
===============================================================================
DHCP Host "host001"
===============================================================================
Admin State          : Down
Last Mgmt Change     : 11/08/2007 12:13:42

Host Indentification
 Circuit Id          : N/A
 Mac Address         : N/A
 Remote Id           : N/A
 Sap Id              : N/A
 Service Id          : N/A
 String              : N/A
 Option 60           : N/A
 System Id           : N/A

Matched Objects      : N/A

Address              : N/A

Identification Strings
 Subscriber Id       : N/A
 SLA Profile String  : N/A
 Sub Profile String  : N/A
 App Profile String  : N/A
 ANCP String         : N/A
 Inter Destination Id: N/A
===============================================================================
*A:ALA-48>show>subscr-mgmt#
```

```
*A:ALA-48>show>subscr-mgmt# local-user-db "database01" dhcp-unmatched-hosts
===============================================================================
Local User Database "database01" DHCP unmatched hosts
===============================================================================
Name                              Reason      Duplicate Host
-------------------------------------------------------------------------------
host002                           No match    N/A
host003                           Duplicate   host001
host004                           No match    N/A
host005                           Duplicate   host001
-------------------------------------------------------------------------------
Number of DHCP Unmatched Hosts : 4
===============================================================================
*A:ALA-48>show>subscr-mgmt#


*A:ALA-48>show>subscr-mgmt# local-user-db "database01" association
===============================================================================
DHCP Servers where database01 is used
===============================================================================
Server-Name                       Router-Name
-------------------------------------------------------------------------------
dhcpS1                            vprn1000
-------------------------------------------------------------------------------
No. of Server(s): 1
===============================================================================
Interfaces where database01 is used for authentication
===============================================================================
Interface-Name                    Service-Id Type
-------------------------------------------------------------------------------
No. of Interface(s): 0
===============================================================================
*A:ALA-48>show>subscr-mgmt#


*A:ALA-48>show>subscr-mgmt# local-user-db "database01" association dhcp
===============================================================================
DHCP Servers where database01 is used
===============================================================================
Server-Name                       Router-Name
-------------------------------------------------------------------------------
dhcpS1                            vprn1000
-------------------------------------------------------------------------------
No. of Server(s): 1
===============================================================================
*A:ALA-48>show>subscr-mgmt#
```

msap-policy

Syntax **msap-policy** [*msap-policy-name* [**association**]]

Context show>subscr-mgmt

Description This command displays Managed SAP policy information.

Sample Output

```
*A:ALA-48>show>subscr-mgmt# msap-policy
===============================================================================
Managed SAP Policies
===============================================================================
Name                              Num    Description
                                  MSAPs
-------------------------------------------------------------------------------
test                              0      (Not Specified)
test 1                            0      (Not Specified)
-------------------------------------------------------------------------------
Number of MSAP Policies : 2
Number of MSAPs         : 0
===============================================================================
*A:ALA-48>show>subscr-mgmt#
```

sla-profile

Syntax **sla-profile** [*sla-profile-name* [**association**]]

Context show>subscriber-mgmt

Description This command displays SLA profile information.

Parameters *sla-profile-name* — Specifies an existing SLA profile name.

association — Displays the information configured with the specified *sla-profile-name*.

Sample Output

```
A:Dut-A# show subscriber-mgmt sla-profile
===============================================================================
SLA Profiles
===============================================================================
Name                              Description
-------------------------------------------------------------------------------
sla_default
sla_prof100_VOIP
sla_prof110_VOIP
sla_prof120_VOIP
sla_prof130_VOIP
sla_prof140_VOIP
sla_prof230_VOIP
sla_prof80
sla_prof80_VOIP
sla_prof81_VOIP
sla_prof90_VOIP
sla_profPC1
sla_profPC2
sla_profPC3
-------------------------------------------------------------------------------
Number of SLA Profiles : 14
===============================================================================
A:Dut-A#
```

```
A:Dut-A# show subscriber-mgmt sla-profile sla_prof100_VOIP
===============================================================================
SLA Profile sla_prof100_VOIP
===============================================================================
Host Limit           : 3 (Remove Oldest)
Ingress Qos-Policy   : 100                    Egress Qos-Policy : 100
Ingress Queuing Type : Service-queuing
Ingress Filter-Id    : N/A                    Egress Filter-Id  : N/A
Last Mgmt Change     : 07/10/2006 12:55:33
-------------------------------------------------------------------------------
Ingress Queue Overrides
-------------------------------------------------------------------------------
Queue Rate      CIR       HiPrio  CBS     MBS
-------------------------------------------------------------------------------
2     4000      -         -       -       -
3     2500      -         -       -       -

-------------------------------------------------------------------------------
Egress Queue Overrides
-------------------------------------------------------------------------------
Queue Rate      CIR       HiPrio  CBS     MBS
-------------------------------------------------------------------------------
2     4000      -         -       -       -
3     2500      -         -       -       -
===============================================================================
A:Dut-A#


A:Dut-A# show subscriber-mgmt sla-profile sla_prof100_VOIP association
===============================================================================
SLA Profile sla_prof100_VOIP
-------------------------------------------------------------------------------
SAP Default-Profile Associations
-------------------------------------------------------------------------------
No associations found.
-------------------------------------------------------------------------------
SAP Static Host Associations
-------------------------------------------------------------------------------
No associations found.
-------------------------------------------------------------------------------
SAP Non-Sub-Traffic-Profile Associations
-------------------------------------------------------------------------------
No associations found.
-------------------------------------------------------------------------------
Sub-Ident-Policy Profile Map Associations
-------------------------------------------------------------------------------
Policy-name : sub_ident_all
 - Key : sla_prof100_VOIP
-------------------------------------------------------------------------------
Sub-Profile Map Associations
-------------------------------------------------------------------------------
No associations found.
-------------------------------------------------------------------------------
Explicit Subscriber Map Associations
-------------------------------------------------------------------------------
No associations found.
===============================================================================
A:Dut-A#
```

statistics

Syntax **statistics iom** (*slot* | **all**) [**host**|**session**|**subscriber**|**summary**] [**non-zero-value-only**]
statistics mda (*mda* | **all**) [**host**|**session**|**subscriber**|**summary**] [**non-zero-value-only**]
statistics port (*port-id* | **all**) [**host**|**session**|**subscriber**|**summary**] [**non-zero-value-only**]
statistics pw-port (*pw-port* | **all**) [**host**|**session**|**subscriber**|**summary**] [**non-zero-value-only**]
statistics system [**host**|**session**|**subscriber**|**summary**] [**non-zero-value-only**]

Context show>subscr-mgmt

Description This command displays enhanced subscriber management statistics per port/pw-port/MDA/IOM/system.

For each statistic, there is current value and peak value, peak value is the highest value since last reset via system boot or command **clear subscriber-mgmt peakvalue-stats**.

Note that the peak values can be reset via the **clear subscriber-mgmt peakvalue-stats** command.

Parameters. **iom** *slot* — Displays specified IOM slot information.

mda *mda* — Displays specified slot/mda information.

port *port-id* — Specifies to display information for both the physical port ID and LAG.

pw-port *pw-port* — Specifies to display information for a pseudowire port ID.

Values 1 — 10239

all — displays statistics of all IOM or MDA or port or pseudowire port in the system.

host — Displays v4/v6 host statistics only.

session — Displays PPPoX/LAC/LNS session statistics only.

subscriber — Displays subscriber statistics only.

summary — Displays summary statistics only.

non-zero-value-only — Displays only non-zero value counters.

Sample Output

```
show subscriber-mgmt statistics port 1/1/1
===============================================================================
Subscriber Management Statistics for Port 1/1/1
===============================================================================
          Type                              Current      Peak      Peak Timestamp
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
Host Statistics
-------------------------------------------------------------------------------
IPv4      PPP Hosts  - IPCP                     0         0
          PPP Hosts  - L2TP                     0         0
          IPOE Hosts - DHCP                     0         0
          IPOE Hosts - ARP                      0         0
          IPOE Hosts - Static                   0         0
          Non-Sub-Traffic Hosts                 0         0
-------------------------------------------------------------------------------
```

```
IPv6      PPP Hosts  - SLAAC                    0         0
          PPP Hosts  - DHCP6 (PD)               0         0
          PPP Hosts  - DHCP6 (NA)               0         0
          IPOE Hosts - SLAAC                    0         0
          IPOE Hosts - DHCP6 (PD)               0         0
          IPOE Hosts - DHCP6 (NA)               0         0
-------------------------------------------------------------------------------
Total     PPP Hosts                             0         0
          IPOE Hosts                            0         0
          IPv4 Hosts                            0         0
          IPv6 Hosts                            0         0
          IPv4 + IPv6 Hosts                     0         0
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Session Statistics
-------------------------------------------------------------------------------
Local     PPP Sessions - PPPoE                  0         0
          PPP Sessions - PPPoEoA                0         0
          PPP Sessions - PPPoA                  0         0
          PPP Sessions - L2TP (LNS)             0         0
-------------------------------------------------------------------------------
LAC       PPP Sessions - PPPoE                  0         0
          PPP Sessions - PPPoEoA                0         0
          PPP Sessions - PPPoA                  0         0
          PPP Sessions - L2TP (LTS)             0         0
-------------------------------------------------------------------------------
Total     PPP Sessions - established            0         0
          PPP Sessions - in setup               0         0
          PPP Sessions - local                  0         0
          PPP Sessions - LAC                    0         0
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Subscriber Statistics
-------------------------------------------------------------------------------
Total     Subscribers                           0         0
-------------------------------------------------------------------------------
===============================================================================
Peak values last reset at : n/a Default:N/A
```

sub-ident-policy

Syntax **sub-ident-policy** [*sub-ident-policy-name* [**association**]]
sub-ident-policy *sub-ident-policy-name* **script** {**primary** | **secondary** | **tertiary**}

Context show>subscriber-mgmt

Description This command displays subscriber identification policy information.

Parameters *sub-ident-policy-name* — Specifies an existing subscriber identification policy name.

association — Displays information configured with the specified *sub-ident-policy-name*.

script {**primary** | **secondary** | **tertiary**} **—** Displays information for the specified identification script.

Sample Output

```
B:Dut-A>show>subscr-mgmt# sub-ident-policy
===============================================================================
Subscriber Identification Policies
===============================================================================
Name                            Description
-------------------------------------------------------------------------------
sub_ident_all
sub_ident_pc
-------------------------------------------------------------------------------
Number of Subscriber Identification Policies : 2
===============================================================================
B:Dut-A>show>subscr-mgmt#


B:Dut-A>show>subscr-mgmt# sub-ident-policy sub_ident_all
===============================================================================
Subscriber Identification Policy sub_ident_all
===============================================================================
Sub Profile Map
-------------------------------------------------------------------------------
Key                             Sub profile
-------------------------------------------------------------------------------
sub_prof100                     sub_prof100
sub_prof110                     sub_prof110
sub_prof120                     sub_prof120
sub_prof130                     sub_prof130
sub_prof140                     sub_prof140
sub_prof230                     sub_prof230
sub_prof80                      sub_prof80
sub_prof81                      sub_prof81
sub_prof90                      sub_prof90
-------------------------------------------------------------------------------
SLA Profile Map
-------------------------------------------------------------------------------
Key                             SLA profile
-------------------------------------------------------------------------------
sla_prof100_VOIP                sla_prof100_VOIP
sla_prof110_VOIP                sla_prof110_VOIP
sla_prof120_VOIP                sla_prof120_VOIP
sla_prof130_VOIP                sla_prof130_VOIP
sla_prof140_VOIP                sla_prof140_VOIP
sla_prof230_VOIP                sla_prof230_VOIP
sla_prof80_VOIP                 sla_prof80_VOIP
sla_prof81_VOIP                 sla_prof81_VOIP
sla_prof90_VOIP                 sla_prof90_VOIP
-------------------------------------------------------------------------------
Python Scripts
-------------------------------------------------------------------------------
#         Admin Oper  Script
          State State Name
-------------------------------------------------------------------------------
Primary   Down  Down  pyTom.py
Secondary Up    Up    pyTomDebug.py
Tertiary  Up    Up    hardcoded.py
===============================================================================
B:Dut-A>show>subscr-mgmt#
B:Dut-A>show>subscr-mgmt# sub-ident-policy sub_ident_all association
===============================================================================
Subscriber Identification Policy sub_ident_all
```

```
===============================================================================
SAP Associations
-------------------------------------------------------------------------------
Service-Id : 80 (VPLS)
 - SAP : 1/2/1:80
Service-Id : 90 (VPLS)
 - SAP : 1/2/1:90
Service-Id : 100 (VPLS)
 - SAP : 1/2/1:100
 - SAP : 1/2/1:101
 - SAP : 1/2/1:102
Service-Id : 110 (VPLS)
 - SAP : 1/2/1:110
 - SAP : 1/2/1:111
 - SAP : 1/2/1:112
Service-Id : 120 (VPLS)
 - SAP : 1/2/1:120
 - SAP : 1/2/1:121
 - SAP : 1/2/1:122
Service-Id : 130 (VPLS)
 - SAP : 1/2/1:130
Service-Id : 140 (VPLS)
 - SAP : 1/2/1:140
===============================================================================
B:Dut-A>show>subscr-mgmt#


B:Dut-A>show>subscr-mgmt# sub-ident-policy sub_ident_all script primary
===============================================================================
Subscriber Identification Policy sub_ident_all
===============================================================================
Primary Script
-------------------------------------------------------------------------------
URL         : ftp://xxx:yyy@a.b.c.d/pyTom.py
Admin State : Down                      Oper State : Down
-------------------------------------------------------------------------------
Source (dumped from memory)
-------------------------------------------------------------------------------
Script is not active.
-------------------------------------------------------------------------------
===============================================================================
B:Dut-A>show>subscr-mgmt#


B:Dut-A>show>subscr-mgmt# sub-ident-policy sub_ident_all script secondary
===============================================================================
Subscriber Identification Policy sub_ident_all
===============================================================================
Secondary Script
-------------------------------------------------------------------------------
URL         : ftp://xxx:yyy@a.b.c.d/pyTomDebug.py
Admin State : Up                        Oper State : Up
-------------------------------------------------------------------------------
Source (dumped from memory)
-------------------------------------------------------------------------------
      1 import alc
      2 yiaddr = alc.dhcp.yiaddr
      3 # Subscriber ID equals full client IP address.
      4 # Note: IP address 10.10.10.10 yields 'sub-168430090'
      5 # and not 'sub-10.10.10.10'
      6 alc.dhcp.sub_ident = 'sub-' + str(yiaddr)
```

```
      7 # DHCP server is configured such that the third byte (field) of the IP
      8 # address indicates the session Profile ID.
      9 alc.dhcp.sla_profile = 'sp-' + str((yiaddr & 0x0000FF00) >> 8)
===============================================================================
B:Dut-A>show>subscr-mgmt#


B:Dut-A>show>subscr-mgmt# sub-ident-policy sub_ident_all script tertiary
===============================================================================
Subscriber Identification Policy sub_ident_all
===============================================================================
Tertiary Script
-------------------------------------------------------------------------------
URL         : ftp://xxx:yyy@a.b.c.d/hardcoded.py
Admin State : Up                        Oper State : Up
-------------------------------------------------------------------------------
Source (dumped from memory)
-------------------------------------------------------------------------------
      1 from alc import dhcp
      2
      3 dhcp.sub_ident = 'sub_ident_A_1'
      4 dhcp.sub_profile_string = 'sub_prof_B_2'
      5 dhcp.sla_profile_string = 'sla_prof_C_3'
      6
===============================================================================
B:Dut-A>show>subscr-mgmt#
```

sub-profile

Syntax **sub-profile** [*sub-profile-name* [**association**]]

Context show>subscriber-mgmt

Description This command displays subscriber profile information.

Parameters *sub-profile-name* — Specifies an existing subscriber profile name.

association — Displays the information configured with the specified *sub-profile-name*.

Sample Output

```
A:Dut-A# show subscriber-mgmt sub-profile
===============================================================================
Subscriber Profiles
===============================================================================
Name                              Description
-------------------------------------------------------------------------------
sub_default
sub_prof100
sub_prof110
sub_prof120
sub_prof130
sub_prof140
sub_prof230
sub_prof80
sub_prof81
sub_prof90
```

```
sub_profPC1
sub_profPC2
sub_profPC3
-------------------------------------------------------------------------------
Number of Subscriber Profiles : 13
===============================================================================
A:Dut-A#


A:Dut-A# show subscriber-mgmt sub-profile sub_prof100
===============================================================================
Subscriber Profile sub_prof100
===============================================================================
I. Sched. Policy : service100
E. Sched. Policy : service100
Acct. Policy     : 1                                Collect Stats : Enabled
Last Mgmt Change : 07/10/2006 12:55:33
-------------------------------------------------------------------------------
Ingress Scheduler Overrides
-------------------------------------------------------------------------------
Scheduler                       Rate      CIR
-------------------------------------------------------------------------------
serv100                         8000      sum
-------------------------------------------------------------------------------
Egress Scheduler Overrides
-------------------------------------------------------------------------------
Scheduler                       Rate      CIR
-------------------------------------------------------------------------------
serv100                         8000      sum
-------------------------------------------------------------------------------
SLA Profile Map
-------------------------------------------------------------------------------
Key                             SLA Profile
-------------------------------------------------------------------------------
No mappings configured.
===============================================================================
A:Dut-A#


A:Dut-A# show subscriber-mgmt sub-profile sub_prof100 association
===============================================================================
Subscriber Profile sub_prof100
-------------------------------------------------------------------------------
SAP Default-Profile Associations
-------------------------------------------------------------------------------
No associations found.
-------------------------------------------------------------------------------
SAP Static Host Associations
-------------------------------------------------------------------------------
No associations found.
-------------------------------------------------------------------------------
SAP Non-Sub-Traffic-Profile Associations
-------------------------------------------------------------------------------
No associations found.
-------------------------------------------------------------------------------
Sub-Ident-Policy Profile Map Associations
-------------------------------------------------------------------------------
Policy-name : sub_ident_all
 - Key : sub_prof100
-------------------------------------------------------------------------------
Explicit Subscriber Map Associations
```

```
-------------------------------------------------------------------------------
No associations found.
===============================================================================
A:Dut-A#
```

port-scheduler-policy

Syntax **port-scheduler-policy** [*port-scheduler-policy-name*] [**association**]
port-scheduler-policy *port-scheduler-policy-name* **network-policy** *network-queue-policy-name*
port-scheduler-policy *port-scheduler-policy-name* **sap-egress** *policy-id*
port-scheduler-policy *port-scheduler-policy-name* **scheduler-policy** *scheduler-policy-name*
port-scheduler-policy *port-scheduler-policy-name* **scheduler-policy** *scheduler-policy-name* **sap-egress** *policy-id*

Context show>qos

Description This command displays scheduler policy information.

Sample Output

```
A:NS072860910>config>qos>port-sched-plcy# info
----------------------------------------------
            max-rate 10000
            group "group1" create
                rate 3000 cir 1000
            exit
            group "group2" create
                rate 2000 cir 500
            exit
            level 7 rate 7000 cir 700 group "group1" weight 3
            level 6 rate 6000 cir 600 group "group1" weight 2
            level 5 rate 5000 cir 500 group "group1" weight 1
            level 2 rate 2000 cir 200 group "group2" weight 2
            level 1 rate 1000 cir 100 group "group2" weight 1
----------------------------------------------

A:NS072860910# show qos scheduler-hierarchy port 5/1/2 vport "fred"
===============================================================================
Scheduler Hierarchy - Port 5/1/2, Vport "fred"
===============================================================================
Port-scheduler-policy psp1
    Port Bandwidth : 1000000    Max Rate : 10000
    Consumed : 0          Offered : 0

[Within CIR Level 8]
    Rate     : max
    Consumed : 0          Offered : 0

    (Q) : 1->5/1/2:1->1
    (Q) : 1->5/1/2:2->1

[Within CIR Group "group1"]
    Rate     : 1000
    Consumed : 0          Offered : 0
```

```
        [Within CIR Level 7]
            Weight   : 3
            Rate     : 700
            Consumed : 0          Offered : 0

            (Q) : 1->5/1/2:1->2
            (Q) : 1->5/1/2:2->2

        [Within CIR Level 6]
            Weight   : 2
            Rate     : 600
            Consumed : 0          Offered : 0

            (Q) : 1->5/1/2:1->3
            (Q) : 1->5/1/2:2->3

        [Within CIR Level 5]
            Weight   : 1
            Rate     : 500
            Consumed : 0          Offered : 0

            (Q) : 1->5/1/2:1->4
            (Q) : 1->5/1/2:2->4

    [Within CIR Level 4]
        Rate     : max
        Consumed : 0          Offered : 0

    [Within CIR Level 3]
        Rate     : max
        Consumed : 0          Offered : 0

        (Q) : 1->5/1/2:1->5
        (Q) : 1->5/1/2:2->5

    [Within CIR Group "group2"]
        Rate : 500
        Consumed : 0          Offered : 0

        [Within CIR Level 2]
            Weight   : 2
            Rate     : 200
            Consumed : 0          Offered : 0

            (Q) : 1->5/1/2:1->6
            (Q) : 1->5/1/2:2->6

        [Within CIR Level 1]
            Weight   : 1
            Rate     : 200
            Consumed : 0          Offered : 0

            (Q) : 1->5/1/2:1->7
            (Q) : 1->5/1/2:2->7

    [Within CIR Level 0]
        Rate     : 0
        Consumed : 0          Offered : 0

        (Q) : 1->5/1/2:1->8
```

```
    (Q) : 1->5/1/2:2->8

[Above CIR Level 8]
    Rate     : max
    Consumed : 0          Offered : 0

[Above CIR Group "group1"]
    Rate     : 3000
    Consumed : 0          Offered : 0

    [Above CIR Level 7]
        Weight   : 3
        Rate     : 7000
        Consumed : 0          Offered : 0

    [Above CIR Level 6]
        Weight   : 2
        Rate     : 6000
        Consumed : 0          Offered : 0

    [Above CIR Level 5]
        Weight   : 1
        Rate     : 5000
        Consumed : 0          Offered : 0

[Above CIR Level 4]
    Rate     : max
    Consumed : 0          Offered : 0

[Above CIR Level 3]
    Rate     : max
    Consumed : 0          Offered : 0

[Above CIR Group "group2"]
    Rate     : 2000
    Consumed : 0          Offered : 0

    [Above CIR Level 2]
        Weight   : 2
        Rate     : 2000
        Consumed : 0          Offered : 0

    [Above CIR Level 1]
        Weight   : 1
        Rate     : 1000
        Consumed : 0          Offered : 0

        (Q) : 1->5/1/2:1->1
        (Q) : 1->5/1/2:1->2
        (Q) : 1->5/1/2:1->3
        (Q) : 1->5/1/2:1->4
        (Q) : 1->5/1/2:1->5
        (Q) : 1->5/1/2:1->6
        (Q) : 1->5/1/2:1->7
        (Q) : 1->5/1/2:1->8
        (Q) : 1->5/1/2:2->1
        (Q) : 1->5/1/2:2->2
        (Q) : 1->5/1/2:2->3
        (Q) : 1->5/1/2:2->4
        (Q) : 1->5/1/2:2->5
        (Q) : 1->5/1/2:2->6
```

```
        (Q) : 1->5/1/2:2->7
        (Q) : 1->5/1/2:2->8
===============================================================================
A:NS072860910#

*A:Bennet-Dut-A>config>qos>port-sched-plcy# show qos port-scheduler-policy "psp"
===============================================================================
QoS Port Scheduler Policy
===============================================================================
Policy-Name       : psp
Description       : (Not Specified)
Max Rate          : max                 Last changed      : 04/15/2010 00:37:02
Group             : 1
Group PIR         : 80000               Group CIR         : max

Group             : 2
Group PIR         : 80000               Group CIR         : max

Group             : 3
Group PIR         : 80000               Group CIR         : max

Group             : 4
Group PIR         : 80000               Group CIR         : max

Lvl1 PIR          : max                 Lvl1 CIR          : max
Lvl1 Group        : 1                   Lvl1 Grp Weight   : 10

Lvl2 PIR          : max                 Lvl2 CIR          : max
Lvl2 Group        : 1                   Lvl2 Grp Weight   : 20

Lvl3 PIR          : max                 Lvl3 CIR          : max
Lvl3 Group        : 2                   Lvl3 Grp Weight   : 30

Lvl4 PIR          : max                 Lvl4 CIR          : max
Lvl4 Group        : 2                   Lvl4 Grp Weight   : 40

Lvl5 PIR          : max                 Lvl5 CIR          : max
Lvl5 Group        : 3                   Lvl5 Grp Weight   : 50

Lvl6 PIR          : max                 Lvl6 CIR          : max
Lvl6 Group        : 3                   Lvl6 Grp Weight   : 60

Lvl7 PIR          : max                 Lvl7 CIR          : max
Lvl7 Group        : 4                   Lvl7 Grp Weight   : 70

Lvl8 PIR          : max                 Lvl8 CIR          : max
Lvl8 Group        : 4                   Lvl8 Grp Weight   : 80

Orphan Lvl        : default             Orphan Weight     : default
Orphan CIR-Lvl    : default             Orphan CIR-Weight : default
===============================================================================
*A:Bennet-Dut-A>config>qos>port-sched-plcy#


*A:Bennet-Dut-A# show qos port-scheduler-policy "psp" association
===============================================================================
QoS Port Scheduler Policy
===============================================================================
Policy-Name       : psp
Description       : (Not Specified)
```

```
-------------------------------------------------------------------------------
Associations
-------------------------------------------------------------------------------
 - Port : 1/1/2 VPort : vp1
===============================================================================
*A:Bennet-Dut-A#


*A:Bennet-Dut-A# show qos port-scheduler-policy "psp" sap-egress 1000
===============================================================================
Compatibility : Port-scheduler Policy psp & Sap Egress Queue 1000
===============================================================================
Orphan Queues :

None Found

Hierarchy     :

Root
|
|---(Q) : 1
|
|---(Q) : 2
|
|---(Q) : 3
|
|---(Q) : 4
|
|---(Q) : 5
|
|---(Q) : 6
|
|---(Q) : 7
|
|---(Q) : 8
===============================================================================
*A:Bennet-Dut-A#
```

sap-egress

Syntax **sap-egress** [*policy-id*] [**association** | **detail**]

Context show>qos

Description This command displays SAP egress policy information.

Parameters *policy-id* — Displays information for the specified SAP egress policy.

association — Displays the information configured with the specified *sap-egress* policy.

detail — Displays detailed information.

Sample Output

```
A:Dut-A# show qos sap-egress
===============================================================================
```

```
Sap Egress Policies
===============================================================================
Policy-Id            Scope     Description
-------------------------------------------------------------------------------
1                    Template  Default SAP egress QoS policy.
30                   Template  1 video channel, 1 BE
31                   Template  1 video EF, 2xvideo AF, 1 BE
80                   Template  Limit outgoing 80
100                  Template  User100
110                  Template  User110
120                  Template  User120
130                  Template  User130
140                  Template  User140
901                  Template  User90_1
902                  Template  User90_2
903                  Template  User90_3
904                  Template  User90_4
905                  Template  User90_5
1000                 Template  Service all
===============================================================================
A:Dut-A#


A:Dut-A# show qos sap-egress 31 detail
===============================================================================
QoS Sap Egress
===============================================================================
Sap Scheduler Policy (31)
-------------------------------------------------------------------------------
Policy-id      : 31                            Scope        : Template
Description    : 1 video EF, 2xvideo AF, 1 BE

-------------------------------------------------------------------------------
Queue         CIR Admin PIR Admin CBS     HiPrio  PIR Lvl/Wt    Parent
              CIR Rule  PIR Rule  MBS             CIR Lvl/Wt
-------------------------------------------------------------------------------
1             0         max       def     def     1/1           limit_8000
              closest   closest   def             0/1
2             0         max       def     def     2/1           limit_8000
              closest   closest   def             0/1
3             0         max       def     def     3/1           limit_8000
              closest   closest   def             0/1
-------------------------------------------------------------------------------
FC Name               Queue-id    Explicit/Default
-------------------------------------------------------------------------------
be                    1           Explicit (0)
af                    2           Explicit (0)
ef                    3           Explicit (0)
-------------------------------------------------------------------------------
Associations
-------------------------------------------------------------------------------
Service-Id     : 23 (VPLS)                    Customer-Id  : 1
 - SAP : 1/2/2:4000
Service-Id     : 30 (VPLS)                    Customer-Id  : 2
 - SAP : lag-1
 - SAP : lag-2:5
Service-Id     : 31 (VPLS)                    Customer-Id  : 2
 - SAP : 1/2/1:31
SLA Profiles :
 - sla_profPC1                      override
-------------------------------------------------------------------------------
```

```
Mirror SAPs
-------------------------------------------------------------------------------
No Mirror SAPs Found.
===============================================================================
A:Dut-A#
```

sap-ingress

Syntax **sap-ingress** [*policy-id*] [**association** | **match-criteria** | **detail**]

Context show>qos

Description This command displays SAP ingress policy information.

Parameters *policy-id* — Displays information for the specified SAP ingress policy.

association — Displays the information configured with the specified *sap-ingress* policy.

match-criteria — Displays information about the matching criteria.

detail — Displays detailed information.

Sample Output

```
A:Dut-A# show qos sap-ingress
===============================================================================
Sap Ingress Policies
===============================================================================
Policy-Id            Scope     Description
-------------------------------------------------------------------------------
1                    Template  Default SAP ingress QoS policy.
80                   Template   Dot1p mappings/service for service 80
90                   Template   Dot1p mappings/service for service 90
100                  Template   Dot1p mappings/service for service 100
110                  Template   Dot1p mappings/service for service 110
120                  Template   Dot1p mappings/service for service 120
130                  Template   Dot1p mappings/service for service 130
140                  Template   Dot1p mappings/service for service 140
901                  Template  User90_1
902                  Template  User90_2
903                  Template  User90_3
904                  Template  User90_4
905                  Template  User90_5
1000                 Template   Dot1p mappings/service for all services
===============================================================================
A:Dut-A#
```

```
A:Dut-A# show qos sap-ingress 80 detail
===============================================================================
QoS Sap Ingress
===============================================================================
Sap Ingress Policy (80)
-------------------------------------------------------------------------------
Policy-id      : 80                          Scope        : Template
Default FC     : be                          Priority     : Low
Criteria-type  : IP
Description    :  Dot1p mappings/service for service 80
```

```
-------------------------------------------------------------------------------
Queue Mode    CIR Admin PIR Admin CBS     HiPrio  PIR Lvl/Wt    Parent
              CIR Rule  PIR Rule  MBS             CIR Lvl/Wt
-------------------------------------------------------------------------------
1     Prio    0         7000      def     def     1/1           serv80
              closest   closest   def             0/1
2     Prio    0         3500      def     def     2/1           serv80
              closest   closest   def             0/1
3     Prio    0         2000      def     def     3/1           serv80
              closest   closest   def             0/1
11    Prio    0         max       def     def     1/1           None
              closest   closest   def             0/1
-------------------------------------------------------------------------------
FC                  UCastQ         MCastQ         BCastQ         UnknownQ
-------------------------------------------------------------------------------
be                  1              def            def            def
af                  2              def            def            def
ef                  3              def            def            def
-------------------------------------------------------------------------------
SubFC                              Profile        In-Remark      Out-Remark
-------------------------------------------------------------------------------
af                                 None           None           None
be                                 None           None           None
ef                                 None           None           None
-------------------------------------------------------------------------------
Dot1p           FC                                 Priority
-------------------------------------------------------------------------------
0               be                                 Default
2               af                                 Default
5               ef                                 Default
-------------------------------------------------------------------------------
DSCP            FC                                 Priority
-------------------------------------------------------------------------------
No DSCP-Map Entries Found.
-------------------------------------------------------------------------------
Prec Value      FC                                 Priority
-------------------------------------------------------------------------------
No Prec-Map Entries Found.
-------------------------------------------------------------------------------
Match Criteria
-------------------------------------------------------------------------------
IP Match Criteria
-------------------------------------------------------------------------------
Entry          : 1
Source IP      : Undefined                    Source Port  : None
Dest. IP       : Undefined                    Dest. Port   : None
Protocol       : None                         DSCP         : None
Fragment       : Off
FC             : Default                      Priority     : Default
-------------------------------------------------------------------------------
Associations
-------------------------------------------------------------------------------
Service-Id     : 80 (VPLS)                    Customer-Id  : 80
 - SAP : 1/2/1:80

SLA Profiles :
 - sla_prof80                       override
 - sla_prof80_VOIP                  override
 - sla_prof81_VOIP                  override
===============================================================================
A:Dut-A#
```

scheduler-hierarchy

Syntax **scheduler-hierarchy**

Context show>qos

Description This command enables the context to display information about policies that use this scheduler.

customer

Syntax **customer** *customer-id* **site** *customer-site-name* [**scheduler** *scheduler-name*] [**ingress**|**egress**] [**detail**]

Context show>qos>scheduler-hierarchy
show>qos>scheduler-stats

Description This command displays the scheduler hierarchy per customer multi-service-site.

Parameters **customer** *customer-id* — Displays information for the specified customer ID.

site ***customer-site-name*** — Displays information for the specified multi-service *customer-site-name*.

scheduler ***scheduler-name*** — Displays information for the specified scheduler-name.

ingress — Displays information for the ingress policy.

egress — Displays information for the egress policy.

detail — Displays detailed information.

sap

Syntax **sap** *sap-id* [**scheduler** *scheduler-name*] [**ingress**|**egress**] [**detail**]

Context show>qos>scheduler-hierarchy
show>qos>scheduler-stats

Description This command displays the scheduler stats per SAP.

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition. See Common Service Commands on page 1740 for *sap-id* command syntax.

scheduler ***scheduler-name*** — Displays information for the specified scheduler-name.

ingress — Displays information for the ingress policy.

egress — Displays information for the egress policy.

detail — Displays detailed information.

subscriber

Syntax **subscriber** *sub-ident-string* [**scheduler** *scheduler-name*] [**ingress**|**egress**][**detail**]

Context show>qos>scheduler-hierarchy
show>qos>scheduler-stats

Description This command displays the scheduler stats per subscriber.

Parameters **subscriber** *sub-ident-string* **—** Displays the subscriber profile name.

scheduler *scheduler-name* — Displays information for the specified scheduler-name.

ingress — Displays information for the ingress policy.

egress — Displays information for the egress policy.

detail — Displays detailed information.

Sample Output

```
B:Dut-A# show qos scheduler-hierarchy subscriber alcatel_100
===============================================================================
Scheduler Hierarchy - Subscriber alcatel_100
===============================================================================
Root (Ing)
| slot(1)
|--(S) : root
|   |
|   |--(S) : serv_all
|   |   |
|   |   |--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:101->11 MCast
|   |   |
|   |   |--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:100->11 MCast
|   |   |
|   |   |--(S) : AccessIngress:Sub=6:1 100->1/2/1:100->3
|   |   |   |
|   |   |   |--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:100->3 1/1
|   |   |   |
|   |   |   |--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:100->3 3/2
|   |   |   |
|   |   |   |--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:100->3 1/2
|   |   |   |
|   |   |
|   |   |--(S) : AccessIngress:Sub=6:1 100->1/2/1:100->2
|   |   |   |
|   |   |   |--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:100->2 1/1
|   |   |   |
|   |   |   |--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:100->2 3/2
|   |   |   |
|   |   |   |--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:100->2 1/2
|   |   |   |
|   |   |
|   |   |--(S) : AccessIngress:Sub=6:1 100->1/2/1:100->1
|   |   |   |
|   |   |   |--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:100->1 1/1
|   |   |   |
|   |   |   |--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:100->1 3/2
|   |   |   |
```

```
|   |   |   |--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:100->1 1/2
|   |   |   |
|   |   |
|   |   |--(S) : AccessIngress:Sub=6:1 100->1/2/1:101->1
|   |   |   |
|   |   |   |--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:101->1 1/1
|   |   |   |
|   |   |   |--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:101->1 3/2
|   |   |   |
|   |   |   |--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:101->1 1/2
|   |   |   |
|   |   |
|   |   |--(S) : AccessIngress:Sub=6:1 100->1/2/1:101->2
|   |   |   |
|   |   |   |--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:101->2 1/1
|   |   |   |
|   |   |   |--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:101->2 3/2
|   |   |   |
|   |   |   |--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:101->2 1/2
|   |   |   |
|   |   |
|   |   |--(S) : AccessIngress:Sub=6:1 100->1/2/1:101->3
|   |   |   |
|   |   |   |--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:101->3 1/1
|   |   |   |
|   |   |   |--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:101->3 3/2
|   |   |   |
|   |   |   |--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:101->3 1/2
|   |   |   |


Root (Egr)
| slot(1)
|--(S) : root
|   |
|   |--(S) : serv_all
|   |   |
|   |   |--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:101->1
|   |   |
|   |   |--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:101->2
|   |   |
|   |   |--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:101->3
|   |   |
|   |   |--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:100->1
|   |   |
|   |   |--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:100->2
|   |   |
|   |   |--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:100->3
|   |   |
===============================================================================
B:Dut-A#


B:Dut-A# show qos scheduler-hierarchy subscriber alcatel_100 scheduler serv_all
===============================================================================
Scheduler Hierarchy - Subscriber alcatel_100
===============================================================================
serv_all (Ing)
| slot(1)
|--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:101->11 MCast
|
```

```
|--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:100->11 MCast
|
|--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:102->11 MCast
|
|--(S) : AccessIngress:Sub=6:1 100->1/2/1:100->2
|   |
|   |--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:100->2 1/1
|   |
|   |--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:100->2 3/2
|   |
|   |--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:100->2 1/2
|   |
|
|--(S) : AccessIngress:Sub=6:1 100->1/2/1:100->1
|   |
|   |--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:100->1 1/1
|   |
|   |--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:100->1 3/2
|   |
|   |--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:100->1 1/2
|   |
|
|--(S) : AccessIngress:Sub=6:1 100->1/2/1:100->3
|   |
|   |--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:100->3 1/1
|   |
|   |--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:100->3 3/2
|   |
|   |--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:100->3 1/2
|   |
|
|--(S) : AccessIngress:Sub=6:1 100->1/2/1:102->1
|   |
|   |--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:102->1 1/1
|   |
|   |--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:102->1 3/2
|   |
|   |--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:102->1 1/2
|   |
|
|--(S) : AccessIngress:Sub=6:1 100->1/2/1:102->2
|   |
|   |--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:102->2 1/1
|   |
|   |--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:102->2 3/2
|   |
|   |--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:102->2 1/2
|   |
...
===============================================================================
B:Dut-A#


B:Dut-A# show qos scheduler-hierarchy subscriber alcatel_100 scheduler serv_all
detail
===============================================================================
Scheduler Hierarchy - Subscriber alcatel_100
===============================================================================
Legend :
(U) - Unrestricted     (P) - Provisioned
(A) - Administrative   (O) - Operational
```

```
MIR - Measured Info Rate
-------------------------------------------------------------------------------
serv_all (Ing)
| slot(1)
|--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:101->11 MCast
|   |
|   |    PIR Lvl:4            PIR Wt :1
|   |    CIR Lvl:0            CIR Wt :1
|   |
|   |    MIR    :0
|   |    PIR (P):0            PIR (U):7000
|   |    CIR (P):0            CIR (U):0
|   |
|   |    PIR (A):1000000      PIR (O):7000
|   |    CIR (A):0            CIR (O):0
|   |    CBS    :0            MBS    :1280
|   |    Depth  :0            Hi Prio:256
|
|--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:102->11 MCast
|   |
|   |    PIR Lvl:4            PIR Wt :1
|   |    CIR Lvl:0            CIR Wt :1
|   |
|   |    MIR    :0
|   |    PIR (P):0            PIR (U):7000
|   |    CIR (P):0            CIR (U):0
|   |
|   |    PIR (A):1000000      PIR (O):7000
|   |    CIR (A):0            CIR (O):0
|   |    CBS    :0            MBS    :1280
|   |    Depth  :0            Hi Prio:256
|
|--(S) : AccessIngress:Sub=6:1 100->1/2/1:102->1
|   |
|   |    PIR Lvl:1            PIR Wt :1
|   |    CIR Lvl:0            CIR Wt :1
|   |
|   |    MIR    :1687
|   |    PIR (P):1690         PIR (U):3510
|   |    CIR (P):0            CIR (U):0
|   |
|   |    PIR (A):7000
|   |    CIR (A):0
|   |
|   |--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:102->1 1/1
|   |   |
|   |   |    PIR Lvl:1            PIR Wt :1
|   |   |    CIR Lvl:1            CIR Wt :1
|   |   |
|   |   |    MIR    :0
|   |   |    PIR (P):0            PIR (U):1830
|   |   |    CIR (P):0            CIR (U):0
|   |   |
|   |   |    PIR (A):7000         PIR (O):1850
|   |   |    CIR (A):0            CIR (O):0
|   |   |    CBS    :0            MBS    :64
|   |   |    Depth  :0            Hi Prio:8
|   |
...
|--(Q) : Sub=alcatel_100:sla_default 100->1/2/1:102->3
|   |
```

```
|   |    PIR Lvl:3            PIR Wt :1
|   |    CIR Lvl:0            CIR Wt :1
|   |
|   |    MIR    :0
|   |    PIR (P):0            PIR (U):2000
|   |    CIR (P):0            CIR (U):0
|   |
|   |    PIR (A):2000         PIR (O):2000
|   |    CIR (A):0            CIR (O):0
|   |    CBS    :0            MBS    :64
|   |    Depth  :0            Hi Prio:8
===============================================================================
B:Dut-A#
```

scheduler-name

Syntax **scheduler-name** *scheduler-name*

Context show>qos

Description This command displays information about the specified scheduler name.

Parameters *scheduler-name* — Displays information about the specified scheduler.

scheduler-policy

Syntax **scheduler-policy** [*scheduler-policy-name*] [**association** | **sap-ingress** *policy-id* | **sap-egress** *policy-id*]

Context show>qos

Description This command displays information about the specified scheduler policy.

Parameters *scheduler-policy-name* — Displays information for the specified scheduler policy.

sap-ingress *policy-id* **—** Displays information for the ingress policy.

sap-egress *policy-id* **—** Displays information for the egress policy.

association — Displays the information currently configured with the specified *scheduler-policy-name*.

Sample Output

```
B:Dut-A# show qos scheduler-policy
===============================================================================
Sap Scheduler Policies
===============================================================================
Policy-Id                       Description
-------------------------------------------------------------------------------
maximum_4000_1xEF_1xBE
maximum_8000_1xEF_2xAF_1xBE
multiservice-site
root
```

```
scheduler-7Mbps
service100
service110
service120
service130
service140
service80
service90
service_all
===============================================================================
B:Dut-A#


B:Dut-A# show qos scheduler-policy root association
===============================================================================
QoS Scheduler Policy
===============================================================================
Policy-Name    : root
-------------------------------------------------------------------------------
Associations
-------------------------------------------------------------------------------
No Association Found.
===============================================================================
B:Dut-A#


B:Dut-A# show qos scheduler-policy association
===============================================================================
QoS Scheduler Policy
===============================================================================
Policy-Name    : maximum_4000_1xEF_1xBE
-------------------------------------------------------------------------------
Associations
-------------------------------------------------------------------------------
No Association Found.

Policy-Name    : maximum_8000_1xEF_2xAF_1xBE
-------------------------------------------------------------------------------
Associations
-------------------------------------------------------------------------------
Service-Id     : 23 (VPLS)                    Customer-Id  : 1
 - SAP : 1/3/2:4000 (Egr)
Service-Id     : 30 (VPLS)                    Customer-Id  : 2
 - SAP : lag-1 (Egr)
 - SAP : lag-2:5 (Egr)
Policy-Name    : multiservice-site
-------------------------------------------------------------------------------
Associations
-------------------------------------------------------------------------------
Service-Id     : 90 (VPLS)                    Customer-Id  : 90
 - SAP : 1/1/12:95 (Ing) (Egr) MSS : site1
 - SAP : 1/1/20:94 (Ing) (Egr) MSS : site1

 - Customer : 2         MSS : site1  (Ing) (Egr)
 - Customer : 90        MSS : site1  (Ing) (Egr)

Policy-Name    : root
-------------------------------------------------------------------------------
Associations
-------------------------------------------------------------------------------
No Association Found.
```

```
Policy-Name    : scheduler-7Mbps
-------------------------------------------------------------------------------
Associations
-------------------------------------------------------------------------------
No Association Found.

Policy-Name    : service100
-------------------------------------------------------------------------------
Associations
-------------------------------------------------------------------------------
Service-Id     : 100 (VPLS)                  Customer-Id  : 100
 - SAP : 1/2/1:100 (Ing) (Egr)
 - SAP : 1/2/1:101 (Ing) (Egr)
 - SAP : 1/2/1:102 (Ing) (Egr)

 - Customer : 100       MSS : site100  (Ing) (Egr)

Sub Profiles :
 - sub_prof100 (Ing) (Egr)

Policy-Name    : service110
-------------------------------------------------------------------------------
Associations
-------------------------------------------------------------------------------
Service-Id     : 110 (VPLS)                  Customer-Id  : 110
 - SAP : 1/2/1:110 (Ing) (Egr)
 - SAP : 1/2/1:111 (Ing) (Egr)
 - SAP : 1/2/1:112 (Ing) (Egr)

Sub Profiles :
 - sub_prof110 (Ing) (Egr)

Policy-Name    : service120
-------------------------------------------------------------------------------
Associations
-------------------------------------------------------------------------------
Service-Id     : 120 (VPLS)                  Customer-Id  : 120
 - SAP : 1/2/1:120 (Ing) (Egr)
 - SAP : 1/2/1:121 (Ing) (Egr)
 - SAP : 1/2/1:122 (Ing) (Egr)

Sub Profiles :
 - sub_prof120 (Ing) (Egr)

Policy-Name    : service130
-------------------------------------------------------------------------------
Associations
-------------------------------------------------------------------------------
Service-Id     : 130 (VPLS)                  Customer-Id  : 130
 - SAP : 1/2/1:130 (Ing) (Egr)

Sub Profiles :
 - sub_prof130 (Ing) (Egr)

Policy-Name    : service140
-------------------------------------------------------------------------------
Associations
-------------------------------------------------------------------------------
Service-Id     : 140 (VPLS)                  Customer-Id  : 140
 - SAP : 1/2/1:140 (Ing) (Egr)
```

```
Sub Profiles :
 - sub_prof140 (Ing) (Egr)

Policy-Name    : service80
-------------------------------------------------------------------------------
Associations
-------------------------------------------------------------------------------
Service-Id     : 80 (VPLS)                    Customer-Id  : 80
 - SAP : 1/2/1:80 (Ing) (Egr)

 - Customer : 80        MSS : site80  (Ing) (Egr)

Sub Profiles :
 - sub_prof80 (Ing) (Egr)
 - sub_prof81 (Ing) (Egr)

Policy-Name    : service90
-------------------------------------------------------------------------------
Associations
-------------------------------------------------------------------------------
Service-Id     : 90 (VPLS)                    Customer-Id  : 90
 - SAP : 1/2/1:90 (Ing) (Egr)

Sub Profiles :
 - sub_prof90 (Ing) (Egr)

Policy-Name    : service_all
-------------------------------------------------------------------------------
Associations
-------------------------------------------------------------------------------
Sub Profiles :
 - sub_default (Ing) (Egr)
===============================================================================
B:Dut-A#
```

scheduler-stats

Syntax **scheduler-stats**

Context show>qos

Description This command enables the context to display scheduler statistics information.

Sample Output

```
A:Dut-A# show qos scheduler-stats subscriber alcatel_100
===============================================================================
Scheduler Stats
===============================================================================
Scheduler                          Forwarded Packets      Forwarded Octets
-------------------------------------------------------------------------------
Ingress Schedulers
root                               112777                 25218126
serv_all                           112777                 25218126
Egress Schedulers
root                               113781                 26008462
```

```
serv_all                            113781                 26008462
===============================================================================
A:Dut-A#


A:Dut-A# show qos scheduler-stats subscriber alcatel_100 scheduler root
===============================================================================
Scheduler Stats
===============================================================================
Scheduler                           Forwarded Packets      Forwarded Octets
-------------------------------------------------------------------------------
Ingress Schedulers
root                                0                      0
Egress Schedulers
root                                0                      0
===============================================================================
A:Dut-A#
```

shared-queue

Syntax **shared-queue** [*shared-queue-policy-name*] [**detail**]

Context show>qos

Description This command displays shared policy information.

Sample Output

```
A:Dut-A# show qos shared-queue
===============================================================================
Shared Queue Policies
===============================================================================
Policy-Id                       Description
-------------------------------------------------------------------------------
default                         Default Shared Queue Policy
===============================================================================
A:Dut-A#


A:Dut-A# show qos shared-queue detail
===============================================================================
QoS Network Queue Policy
-------------------------------------------------------------------------------
Shared Queue Policy (default)
-------------------------------------------------------------------------------
Policy         : default
Description    : Default Shared Queue Policy
-------------------------------------------------------------------------------
Queue CIR       PIR       CBS     MBS     HiPrio  Multipoint
-------------------------------------------------------------------------------
1     0         100       1       50      10      FALSE
2     25        100       3       50      10      FALSE
3     25        100       10      50      10      FALSE
4     25        100       3       25      10      FALSE
5     100       100       10      50      10      FALSE
6     100       100       10      50      10      FALSE
7     10        100       3       25      10      FALSE
```

```
8     10        100       3       25      10      FALSE
9     0         100       1       50      10      TRUE
10    25        100       3       50      10      TRUE
11    25        100       10      50      10      TRUE
12    25        100       3       25      10      TRUE
13    100       100       10      50      10      TRUE
14    100       100       10      50      10      TRUE
15    10        100       3       25      10      TRUE
16    10        100       3       25      10      TRUE
17    0         100       1       50      10      TRUE
18    25        100       3       50      10      TRUE
19    25        100       10      50      10      TRUE
20    25        100       3       25      10      TRUE
21    100       100       10      50      10      TRUE
22    100       100       10      50      10      TRUE
23    10        100       3       25      10      TRUE
24    10        100       3       25      10      TRUE
25    0         100       1       50      10      TRUE
26    25        100       3       50      10      TRUE
27    25        100       10      50      10      TRUE
28    25        100       3       25      10      TRUE
29    100       100       10      50      10      TRUE
30    100       100       10      50      10      TRUE
31    10        100       3       25      10      TRUE
32    10        100       3       25      10      TRUE
-------------------------------------------------------------------------------
FC    UCastQ    MCastQ    BCastQ  UnknownQ
-------------------------------------------------------------------------------
be    1         9         17      25
l2    2         10        18      26
af    3         11        19      27
l1    4         12        20      28
h2    5         13        21      29
ef    6         14        22      30
h1    7         15        23      31
nc    8         16        24      32
-------------------------------------------------------------------------------
Associations
-------------------------------------------------------------------------------
Service : 10            SAP : 1/1/4:101 (shared Q)
Service : 10            SAP : 1/1/4:102 (shared Q)
Service : 10            SAP : 1/1/4:103 (shared Q)
Service : 10            SAP : 1/1/4:104 (shared Q)
Service : 10            SAP : 1/1/4:105 (shared Q)
Service : 10            SAP : 1/1/4:106 (shared Q)
Service : 10            SAP : 1/1/4:107 (shared Q)
Service : 10            SAP : 1/1/4:108 (shared Q)
Service : 10            SAP : 1/1/4:109 (shared Q)
Service : 10            SAP : 1/1/4:110 (shared Q)
Service : 10            SAP : 1/1/4:111 (shared Q)
Service : 10            SAP : 1/1/4:112 (shared Q)
Service : 10            SAP : 1/1/4:113 (shared Q)
Service : 10            SAP : 1/1/4:114 (shared Q)
Service : 10            SAP : 1/1/4:115 (shared Q)
Service : 10            SAP : 1/1/4:116 (shared Q)
Service : 10            SAP : 1/1/4:117 (shared Q)
Service : 10            SAP : 1/1/4:118 (shared Q)
Service : 10            SAP : 1/1/4:119 (shared Q)
Service : 10            SAP : 1/1/4:120 (shared Q)
Service : 10            SAP : 1/1/4:121 (shared Q)
Service : 10            SAP : 1/1/4:122 (shared Q)
```

```
Service : 10            SAP : 1/1/4:123 (shared Q)
Service : 10            SAP : 1/1/5:279 (shared Q)
===============================================================================
A:Dut-A#
```

ancp-policy

Syntax **ancp-policy** [*policy-name*]

Context show>subscriber-management

Description This command displays subscriber ANCP policy information.

Sample Output

```
A:active# show subscriber-mgmt ancp-policy
===============================================================================
ANCP Policies
===============================================================================
adsl-operator1
vdsl-operator1
-------------------------------------------------------------------------------
Number of ANCP policies : 2
===============================================================================
A:active#

A:active# show subscriber-mgmt ancp-policy adsl-operator1
===============================================================================
ANCP Policy "adsl-operator1"
===============================================================================
I. Rate Reduction      : 0 kbps
I. Rate Adjustment     : 100 percent
I. Rate Monitor        : 0 kbps
I. Rate Monitor Alarm  : no
I. Rate Modify         : scheduler "root"

E. Rate Reduction      : 10 kbps
E. Rate Adjustment     : 100 percent
E. Rate Monitor        : 0 kbps
E. Rate Monitor Alarm  : no
E. Rate Modify         : scheduler "root"
Port Down : N/A
Last Mgmt Change: 01/26/2007 17:10:51
===============================================================================
A:active#

A:active# show subscriber-mgmt ancp-policy adsl-operator1 association
===============================================================================
ANCP Policy "adsl-operator1" associations
===============================================================================
SAP Static Map Associations
-------------------------------------------------------------------------------
- SAP     : 1/1/3                            Svc-id  : 333 (VPLS)
     String : "ANCP-String-1"
     String : "ANCP-String-2"
-------------------------------------------------------------------------------
MSS Static Map Associations
```

```
-------------------------------------------------------------------------------
- Cust-id : 1                                    MSS-name: mss1
     String : "ANCP-String-3"
-------------------------------------------------------------------------------
Subscriber Associations
-------------------------------------------------------------------------------
No associations found.
Number of associations : 3
===============================================================================
A:active#
```

ancp-string

Syntax **ancp-string**
ancp-string *ancp-string*
ancp-string customer *customer-id* **site** *customer-site-name*
ancp-string sap *sap-id*

Context show>subscriber-management

Description This command displays subscriber ANCP string information.

Parameters *ancp-string* — Specify the ASCII representation of the DSLAM circuit-id name.

customer *customer-id* — Specify the associated existing customer name.

site *customer-site-name* — Specify the associated customer's configured MSS name.

sap *sap-id* — Specifies the physical port identifier portion of the SAP definition. See Common Service Commands on page 1740 for *sap-id* command syntax.

Sample Output

```
A:active# show subscriber-mgmt ancp-string
===============================================================================
ANCP-Strings
===============================================================================
ANCP-String                                                 Assoc State
-------------------------------------------------------------------------------
"ANCP-String-1"                                             SAP   Up
"ANCP-String-2"                                             SAP   Down
"ANCP-String-3"                                             MSS   Up
"ANCP-String-4"                                             MSS   Unknown
"ANCP-String-5"                                             ANCP  Up
"ANCP-String-6"                                             MSS   Unknown
-------------------------------------------------------------------------------
Number of ANCP-Strings : 6
===============================================================================
A:active#


*A:Dut-C# show subscriber-mgmt ancp-string hpolSub43
===============================================================================
ANCP-String "hpolSub43"
===============================================================================
Type       : SUB - "hpolSub43"
State      : Up                Ancp Policy: ancpPol
```

```
I. Rate    : 100 kbps           E. Rate    : 200 kbps
Adj I. Rate: N/A                Adj E. Rate: 200 kbps
Act I. Rate: N/A                Act E. Rate: 182 kbps
Service Id : 1 (VPRN)
Group      : Alu
Neighbor   : 100.100.100.1:49063
===============================================================================
*A:Dut-C#
```

Other applicable show command output:

```
A:active# show service id 333 sap 1/1/3 detail
===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 333
SAP                : 1/1/3                    Encap             : null
...
-------------------------------------------------------------------------------
ANCP Override
-------------------------------------------------------------------------------
Ing Sched Name: root
- PIR    : 100 kbps
- String : "ANCP-String-1"
Egr Sched Name: root
- PIR    : 100 kbps
- String : "ANCP-String-1"
-------------------------------------------------------------------------------
...
Dro. InProf           : 0                        0
Dro. OutProf          : 0                        0
===============================================================================
A:active#

A:active# show service customer 1 site mss1
===============================================================================
Customer  1
===============================================================================
Customer-ID        : 1
Description        : Default customer
...
-------------------------------------------------------------------------------
ANCP Override
-------------------------------------------------------------------------------
Egr Sched Name: root
- PIR    : 90 kbps
- String : "ANCP-String-3"

-------------------------------------------------------------------------------
Service Association
-------------------------------------------------------------------------------
No Service Association Found.

===============================================================================
A:active#
```

Clear Commands

ancp-sub-string

Syntax **ancp-sub-string** *string*

Context clear>subscr-mgmt>ancp>ancp

Description This command clears subscriber ANCP data.

Parameters *string* — Clears the ANCP string corresponding to this subscriber ID.

authentication

Syntax **authentication** [*policy-name*]
authentication coa-statistics

Context clear

Description This command clears subscriber authentication data.

Parameters *policy-name* — Clears the authentication policy name. The policy must be already configured.

coa-statistics — Clears statistics for incoming RADIUS Change of Authorization requests.

msap-policy

Syntax **msap-policy** *msap-policy-name*

Context clear> subscriber-mgmt

Description This command deletes managed SAPs created by the managed SAP policy.

Parameters *msap-policy-name* — Specifies an existing managed SAP policy name. Any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

peakvalue-stats

Syntax **peakvalue-stats iom** (*slot* | **all**) [**recursive**]
peakvalue-stats mda (*mda* | **all**) [**recursive**]
peakvalue-stats port (*port-id* | **all**)
peakvalue-stats pw-port (*pw-port* | **all**)
peakvalue-stats system [**recursive**]

Context clear> subscriber-mgmt

Description This command resets the most recent peak counter.

Note that clearing one counter will not impact other counters. For example, clearing one IOM's most recent peak value will not impact chassis peak value.

Parameters **iom** *slot* **—** Clears IOM host peak value statistics for the specified IOM.

mda *mda* **—** Clears MDA host peak value statistics for the specified MDA.

port *port-id* **—** Clears port host peak value statistics for the specified port ID.

pw-port *pw-port* **—** Clears pseudowire port host peak value statistics for the specified port.

Values 1 — 10239

system — Clears system host peak value statistics.

all — Clears all host peak value statistics.

recursive — Resets the sub-level counters. For example, clearing IOM counters with the **recursive** keyword will also clear counters of all ports counters on that IOM.

radius-accounting

Syntax **radius-accounting** [*policy-name*]

Context clear> subscriber-mgmt

Description This command clears RADIUS accounting data for the specified policy.

Parameters *policy-name* — The name of the policy. The string is case sensitive and limited to 32 ASCII 7-bit printable characters with no spaces

scheduler-stats

Syntax **scheduler-stats**

Context clear>qos

Description This command clears scheduler statistics.

subscriber

Syntax **subscriber** *sub-ident-string* [**scheduler** *scheduler-name*] [**ingress**|**egress**]

Context clear>qos>scheduler-stats

Description This command clears scheduler stats per subscriber.

Parameters *sub-ident-string* — Clears information for the subscriber profile name.

scheduler *scheduler-name* — Clears information for the specified scheduler-name.

egress — Clears egress information for the subscriber.

ingress — Clears ingress information for the subscriber.

srrp

Syntax **srrp**

Context clear>router

Description This command enables the context to clear and reset SRRP virtual router instances.

interface

Syntax **interface** *subscriber-interface* [**id** *srrp-id*]

Context clear>router>srrp

Description This command clears and resets SRRP interface instances.

Parameters *subscriber-interface* — Specifies an existing subscriber interface name.

Values 32 chars max

id *srrp-id* **—** Specifies an existing SRRP ID.

Values 1 — 4294967295

statistics

Syntax **statistics interface** *subscriber-interface* [**id** *srrp-id*]

Context clear>router>srrp

Description This command clears statistics for SRRP instances.

Parameters *subscriber-interface* — Specifies an existing subscriber interface name.

Values 32 chars max

id *srrp-id* **—** Specifies an existing SRRP ID.

Values 1 — 4294967295

route-downloader

Syntax **route-downloader** *name* [**vprn** *vprn*] [**family** *family*]

Context clear>aaa

Description This command clears all the radius-downloaded routes from the internal downloader cache (or protocol RIB/db) (and thus eventually from the RTM itself). The parameters **vprn** and/or **family** allow to restrict the deletion of those routes learned in a particular address family (IPv4 or IPv6) and/or a particular VPRN.

By default, all VPRNs and both IPv4 and IPv6 families are affected.

Note that A clear of the internal protocol DB means the corresponding prefix that were deleted should be removed from the RTM (and from any other exports) as well.

Parameters **vprn —** Specifies to limit the removal of prefixes to only the specific VPRN. The parameter can be either the service-id or service-name that identifies a VPRN.

family *family* **—** Specifies to limit he removal or prefixes only belonging to the address family IPv4 or IPv6. Only these two values will be accepted.

Values ipv4, ipv6

Tools Commands

tools

Syntax **tools**

Context <root>

Description The context to enable useful tools for debugging purposes.

Default none

Parameters **dump —** Enables dump tools for the various protocols.

perform — Enables tools to perform specific tasks.

perform

Syntax **perform**

Context tools

Description This command enables the context to enable tools to perform specific tasks.

Default none

persistence

Syntax **persistence**

Context tools>perform

Description This command enables the context to configure downgrade paramters.

downgrade

Syntax **downgrade target-version** *target* [**reboot**]

Context tools>perform>persistence

Description This command downgrades persistence files to a previous version.

Parameters **target-version** *target* — Specifies the downgrade version.

reboot — Specifies to reboot the system after a successful conversion.

subscriber-mgmt

Syntax **subscriber-mgmt**

Context tools>perform

Description This command enables tools to control subscriber management.

edit-lease-state

Syntax **edit-lease-state sap** *sap-id* **ip** *ip-address* [**subscriber** *sub-ident-string*] [**sub-profile-string** *sub-profile-string*] [**sla-profile-string** *sla-profile-string*]
edit-lease-state svc-id *service-id* ip *ip-address* [**subscriber** *sub-ident-string*] [**sub-profile-string** *sub-profile-string*] [**sla-profile-string** *sla-profile-string*]

Context tools>perform>subscr-mgmt

Parameters **sap *sap-id*** — Specifies the physical port identifier portion of the SAP definition. See Common Service Commands on page 1740 for *sap-id* command syntax.

ip *ip-address* — Modifies lease state information for the specified IP address.

subscriber *sub-ident-string* — Modifies lease state information for the specified subscriber identification.

sub-profile-string *sub-profile-string* — Modifies lease state information for the specified subscriber profile.

sla-profile-string *sla-profile-string* — Modifies lease state information for the specified SLA profile.

svc-id *service-id* — Modifies lease state information for the specified service ID.

Values *service-id*: 1 — 2147483647
svc-name: 64 characters maximum

credit-reset

Syntax **credit-reset sap** *sap-id* **subscriber** *sub-ident-string* **sla-profile** *sla-profile-name* {**category** *category-name*|**all-categories**}
credit-reset sap *sap-id* **ip** *ip-address* {**category** *category-name*| **all-categories**}
credit-reset svc *service-id* **ip** *ip-address* {**category** *category-name*| **all-categories**}

Context tools>perform>subscr-mgmt

Description This command resets the credit for an SLA-profile instance.

Parameters **sap *sap-id*** — Specifies the physical port identifier portion of the SAP definition. See Common Service Commands on page 1740 for *sap-id* command syntax.

ip *ip-address* — Modifies lease state information for the specified IP address.

subscriber *sub-ident-string* — Modifies lease state information for the specified subscriber identification.

sub-profile-string *sub-profile-string* — Modifies lease state information for the specified subscriber profile.

sla-profile-string *sla-profile-string* — Modifies lease state information for the specified SLA profile.

svc-id *service-id* — Modifies lease state information for the specified service ID.

Values *service-id*: 1 — 2147483647
svc-name: 64 characters maximum

eval-lease-state

Syntax **eval-lease-state** [**svc-id** *service-id*] [**sap** *sap-id*] [**subscriber** *sub-ident-string*] [**ip** *ip-address*]

Context tools>perform>subscr-mgmt

Description This command evaluates lease state information.

Parameters **svc-id *service-id* —** Evaluates lease state information for the specified service.

Values *service-id*: 1 — 2147483647
svc-name: 64 characters maximum

sap *sap-id* — Evaluates lease state information for the specified SAP.

sap-id — Specifies the physical port identifier portion of the SAP definition. See Common Service Commands on page 1740 for *sap-id* command syntax.

subscriber *sub-ident-string* — Evaluates lease state information for the specified subscriber identification string.

ip *ip-address* — Evaluates lease state information for the specified IP address.

re-ident-sub

Syntax **re-ident-sub** *old-sub-ident-string* **to** *new-sub-ident-string*

Context tools>perform>subscr-mgmt

Description This command renames a subscriber identification string.

Parameters *old-sub-ident-string* — Specifies the existing subscriber identification string to be renamed.

new-sub-ident-string — Specifies the new subscriber identification string name.

redundancy

Syntax **redundancy**

Context tools>dump

Description This command enables the context to dump redundancy parameters.

multi-chassis

Syntax **multi-chassis**

Context tools>dump>redundancy

Description This command enables the context to dump multi-chassis parameters.

mc-ipsec

Syntax **mc-ipsec**

Context tools>perform>redundancy>multi-chassis

Description This command enters the mc-ipsec context.

force-switchover

Syntax **force-switchover tunnel-group** *local-group-id*

Context tools>perform>redundancy>multi-chassis>mc-ipsec

Description This command manually switches over mc-ipsec mastership of the specified tunnel-group.

Parameters *local-group-id* — Specifies the local tunnel-group ID configured under config>redundancy.multi-chassis>peer>mc-ipsec.

mc-ring

Syntax **mc-ring**

Context tools>dump>redundancy>multi-chassis

Description This command dumps multi-chassis ring data.

sync-database

Syntax **sync-database** [**peer** *ip-address*] [**port** *port-id* | *lag-id*] [**sync-tag** *sync-tag*] [**application** *application*] [**detail**] [**type** *type*]

Context tools>dump>redundancy>multi-chassis

Description This command dumps multi-chassis sync database information.

Parameters **peer** *ip-address* — Dumps the specified address of the multi-chassis peer.

port *port-id* — Dumps the specified port ID of the multi-chassis peer.

port *lag-id* — Dumps the specified Link Aggregation Group (LAG) on this system.

sync-tag *sync-tag* — Dumps the synchronization tag used while synchronizing this port with the multi-chassis peer.

application — Dumps the specified application information that was synchronized with the multi-chassis peer.

Values dhcps, igmp, igmp-snooping, mc-ring, srrp, sub-mgmt, mld-snooping, all

detail — Displays detailed information.

type *type* — Filters by the specified entry type.

Values alarm-deleted, local-deleted

srrp-sync-data

Syntax **srrp-sync-database** [**instance** *instance-id*] [**peer** *ip-address*]

Context tools>dump>redundancy>multi-chassis

Description This command dumps multi-chassis SRRP sync database information.

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

Values 1 —4294967295

ip-address — Dumps the specified address (in the form of a.b.c.d).

route-downloader

Syntax **route-downloader start** [**force**]

Context tools>perform>aaa

Description This command causes the download process to start immediately. If an ongoing download is already in progress then no further action is needed, except if the **force** keyword is added. In case the **force** keyword is added, then the current download is aborted and a new one is immediately restarted. If aborting the current download, the internal route table should not be emptied or cleared.

Parameters **start —** Starts the download process immediately.

force — Causes the current download to be aborted and a new one is immediately restarted.

Debug Commands

arp

Syntax **arp** {**all** | *ip-address*}
arp interface [*ip-int-name* | *ip-address*]

Context clear>router

Description This command clears all or specific ARP entries.

The scope of ARP cache entries cleared depends on the command line option(s) specified.

Parameters **all —** Clears all ARP cache entries.

ip-addr — Clears the ARP cache entry for the specified IP address.

interface *ip-int-name* **—** Clears all ARP cache entries for the interface with the specified name.

interface *ip-addr* **—** Clears all ARP cache entries for the specified interface with the specified address.

arp-host

Syntax [**no**] arp-host

Context debug>service>id

Description This command enables and configures ARP host debugging.

The no form of the command disables ARP host debugging.

one-time-http-redirection

Syntax one-time-http-redirection

Context debug>service>id

Description This command produces one-time http redirection debug output.

ppp

Syntax [**no**] **ppp**

Context debug>service>id>

Description This command enables the PPP debug context.

event

event

Syntax [**no**] **event**

Context debug>service>id>ppp

Description This command enables the PPP event debug context.

dhcp-client

Syntax **dhcp-client** [**terminate-only**]
no dhcp-client

Context debug>service>id>ppp>event

Description This command enable PPP event debug for DHCP client.

Parameters **terminate-only —** Enables debug for local terminated PPP session

l2tp

Syntax **l2tp** [**terminate-only**]
no l2tp

Context debug>service>id>ppp>event

Description This command enables PPP L2TP event debug.

Parameters **terminate-only —** Enables debug for local terminated PPP session

ppp

Syntax **ppp** [**terminate-only**]
no ppp

Context debug>service>id>ppp>event

Description This command enables PPP event debug.

Parameters **terminate-only —** Enables debug for local terminated PPP session

packet

Syntax [**no**] **packet**

Context debug>service>id>ppp

Description This command enables the PPP packet debug context.

detail-level

Syntax **detail-level** {**low**|**medium**|**high**}
no detail-level

Context debug>service>id>ppp>packet

Description This command specify the detail level of PPP packet debug output.

dhcp-client

Syntax [**no**] **dhcp-client**

Context debug>service>id>ppp>packet

This command enables packet debug output for DHCP client of the PPP session

discovery

Syntax **discovery** [**padi**] [**pado**] [**padr**] [**pads**] [**padt**]
no discovery

Context debug>service>id>ppp>packet

Description This command enables PPP discovery packet debug output.

Parameters **padi/pado/padr/pads/padt —** Enables the corresponding type of PPP discovery packet

mode

Syntax **mode** {**dropped-only**|**ingr-and-dropped**|**egr-ingr-and-dropped**}
no mode

Context debug>service>id>ppp>packet

Description This command specifies PPP packet debug mode

Parameters **dropped-only —** Only displays dropped packet.

ingr-and-dropped: — Only displays ingress packet and dropped packet.

egr-ingr-and-dropped: — Displays ingress, egress and dropped packet.

ppp

Syntax **ppp** [**lcp**] [**pap**] [**chap**] [**ipcp**] [**ipv6cp**]
no ppp

Context debug>service>id>ppp>packet

Description This command enables PPP discovery packet debug output for the specified PPP protocol.

Parameters **lcp/pap/chap/ipcp/ipv6cp —** Enables debug for the specified protocol.

sap

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

Context debug>service>id>ppp>packet

Description This command enables PPP debug output for the specified SAP, this command allow multiple instances.

Parameters *sap-id* — Specifies the SAP ID.

username

Syntax [**no**] **username** *username*

Context debug>service>id>ppp

Description This command enable PPP debug for the specified username. since not all PPP packets contain username, so a mac debug filter will be created automatically when system sees a PPP packet contain the specified username.

Multiple username filters can be specified in the same debug command.

Parameters *user-name* — Specifies the ppp username.

circuit-id

Syntax [**no**] **circuit-id** *circuit-id*

Context debug>service>id>ppp

Description This command enable PPP debug for the specified circuit-id.

Multiple circuit-id filters can be specified in the same debug command.

Parameters *circuit-id* — Specifies the circuit-id in PADI.

remote-id

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

Context debug>service>id>ppp

Description This command enable PPP debug for the specified remote-id.

Multiple remote-id filters could be specified in the same debug command.

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

msap

Syntax [**no**] **msap** <*msap-id*

Context debug>service>id>ppp

Description This command enable PPP debug for the specified managed SAP.

Multiple msap filters could be specified in the same debug command.

Parameters *msap-id* — Specifies the managed SAP ID.

authentication

Syntax **authentication** [**policy** *policy-name*] [**mac-addr** *ieee-address*] [**circuit-id** *circuit-id*]

Context debug>subscr-mgmt

Description This command debugs subscriber authentication.

Parameters **policy** *policy-name* **—** Specify an existing subscriber management authentication policy name.

mac-addr *ieee-address* **—** Specifies the 48-bit MAC address xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx.

circuit-id *circuit-id* **—** Specify the circuit-id, up to 256 characters.

sub-ident-policy

Syntax [**no**] **sub-ident-policy** *policy-name*

Context debug>subscr-mgmt

Description This command debugs subscriber identification policies.

Parameters *policy-name* — Specifies the subscriber identification policy to debug.

script-compile-error

Syntax [**no**] **script-compile-error**

Context debug>subscr-mgmt>sub-ident-plcy

Description This command send the traceback of the compile error to the logger. The traceback contains detailed information about where and why the compilation fails. The compilation takes place when the CLI user changes the admin state of the Python URL from shutdown to no-shutdown.

script-export-variables

Syntax [**no**] **script-export-variables**

Context debug>subscr-mgmt>sub-ident-plcy

Description This command sends the result (the three output variables) of the Python script to the logger when the script ran successfully.

script-output

Syntax [**no**] **script-output**

Context debug>subscr-mgmt>sub-ident-plcy

Description This command sends the output (such as from 'print' statements) of the Python script to the logger.

script-output-on-error

Syntax [**no**] **script-output-on-error**

Context debug>subscr-mgmt>sub-ident-plcy

Description This command sends the output (such as from 'print' statements) of the Python script to the logger, but only when the script fails.

script-runtime-error

Syntax [**no**] **script-runtime-error**

Context debug>subscr-mgmt>sub-ident-plcy

Description This command sends the traceback of the Python script failure to the logger.

script-all-info

Syntax **script-all-info**

Context debug>subscr-mgmt>sub-ident-plcy

Description This command enables the script-compile-error, script-export-variables, script-output, script-output-on-error, and script-runtime-error functionalities.

srrp

Syntax [**no**] **srrp**

Context debug>router

Description This command enables debugging for SRRP packets.

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

events

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

Context debug>router>srrp

Description This command enables debugging for SRRP packets.

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

packets

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

Context debug>router>srrp

Description This command enables debugging for SRRP packets.

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

radius

Syntax [**no**] **radius**

Context debug>router

Description This command enables the debug router RADIUS context.

detail-level

Syntax **detail-level** {**low**|**medium**|**high**}
no detail-level

Context debug>router>radius

Description This command specifies the output detail level of command **debug router radius**.

Default medium

Parameters **low —** Output includes packet type, server address, length, radius-server-policy name

medium — All output in low level plus RADIUS attributes in the packet

high — All output in medium level plus hex packet dump

packet-type

Syntax **packet-type** [**authentication**] [**accounting**] [**coa**]
no packet-type

Context debug>router>radius

Description This command specifies the RADIUS packet type filter of command **debug router radius**

Default authentication accounting coa

Parameters **authentication —** RADIUS authentication packet.

accounting — RADIUS accounting packet.

coa — RADIUS change of authorization packet.

radius-attr

radius-attr type *attribute-type* [**transaction**]
radius-attr type *attribute-type* [**transaction**] {**address**|**hex**|**integer**|**string**} **value** *attribute-value*
radius-attr vendor *vendor-id* **type** *attribute-type* [**transaction**] [**encoding** *encoding-type*]
radius-attr vendor *vendor-id* **type** *attribute-type* [**transaction**] [**encoding** *encoding-type*] {**address**|**hex**|**integer**|**string**} **value** *attribute-value*
no radius-attr type *attribute-type*
no radius-attr type *attribute-type* {**address**|**hex**|**integer**|**string**} **value** *attribute-value*
no radius-attr vendor *vendor-id* **type** *attribute-type*
no radius-attr vendor *vendor-id* **type** *attribute-type* {**address**|**hex**|**integer**|**string**} **value** *attribute-value*

Context debug>router>radius

Description This command specifies the RADIUS attribute filter of command **debug router radius**.

Default none

Parameters *attribute-type* — Specifies the RADIUS attribute type.

address — Specifies the value is a IPv4 or IPv6 address/prefix/subnet

string — Specifies the value is a ASCII string

integer — Specifies the value is a integer

hex — Specifies the value is a binary string in hex format, e.g: "\0xAB01FE"

attribute-value — Specifies the value of the RADIUS attribute.

transaction — With this parameter, system will output both request and response packets in the same session even in case response packet doesn't include the filter attribute.

vendor-id — Specifies the vendor id for the vendor specific attribute.

encoding-type — Specifies the size of vendor-type and vendor-length in bytes. It is a two digitals string: "xy", x is the size of vendor-type, range from 1-4; y is the size of vendor-length of vendor-length, range from 0-2; it is "11" by default.

Monitor Commands

subscriber

Syntax **subscriber** *sub-ident-string* **sap** *sap-id* **sla-profile** *sla-profile-name* [**base** | **ingress-queue-id** *ingress-queue-id* | **egress-queue-id** *egress-queue-id*] [**interval** *seconds*] [**repeat** *repeat*] [**absolute** | **rate**]

Context monitor>service

Description This command monitors statistics for a subscriber.

Parameters **sub-ident-string** — Specifies an existing subscriber identification profile to monitor.

sap *sap-id* — Specifies the physical port identifier portion of the SAP definition. See Common Service Commands on page 1740 for *sap-id* command syntax.

sla-profile ***sla-profile-name*** — Specifies an existing SLA profile.

interval *seconds* — Configures the interval for each display in seconds.

Default 11

Values 11 — 60

repeat *repeat* — Configures how many times the command is repeated.

Default 10

Values 1 — 999

absolute — When the **absolute** keyword is specified, the raw statistics are displayed, without processing. No calculations are performed on the delta or rate statistics.

Default mode delta

rate — When the **rate** keyword is specified, the rate-per-second for each statistic is displayed instead of the delta.

base — Monitor base statistics.

ingress-queue-id *ingress-queue-id* — Monitors statistics for this queue.

Values 1 — 32

egress-queue-id *egress-queue-id* — Monitors statistics for this queue.

Values 1 — 8

Sample Output

```
A:Dut-A# monitor service subscriber alcatel_100 sap 1/2/1:101 sla-profile sla_default
===============================================================================
Monitor statistics for Subscriber alcatel_100
===============================================================================
At time t = 0 sec (Base Statistics)
-------------------------------------------------------------------------------
```

```
SLA Profile Instance statistics
-------------------------------------------------------------------------------
                          Packets                 Octets
Off. HiPrio           : 0                       0
Off. LowPrio          : 94531                   30704535
Off. Uncolor          : 0                       0

Queueing Stats (Ingress QoS Policy 1000)
Dro. HiPrio           : 0                       0
Dro. LowPrio          : 7332                    2510859
For. InProf           : 0                       0
For. OutProf          : 87067                   28152288

Queueing Stats (Egress QoS Policy 1000)
Dro. InProf           : 880                     127660
Dro. OutProf          : 0                       0
For. InProf           : 90862                   12995616
For. OutProf          : 0                       0
-------------------------------------------------------------------------------
SLA Profile Instance per Queue statistics
-------------------------------------------------------------------------------
                          Packets                 Octets
Ingress Queue 1 (Unicast) (Priority)
Off. HiPrio           : 0                       0
Off. LowPrio          : 0                       0
Off. Uncolor          : 0                       0
Dro. HiPrio           : 0                       0
Dro. LowPrio          : 0                       0
For. InProf           : 0                       0
For. OutProf          : 0                       0

Ingress Queue 2 (Unicast) (Priority)
Off. HiPrio           : 0                       0
Off. LowPrio          : 94531                   30704535
Off. Uncolor          : 0                       0
Dro. HiPrio           : 0                       0
Dro. LowPrio          : 7332                    2510859
For. InProf           : 0                       0
For. OutProf          : 87067                   28152288

Ingress Queue 3 (Unicast) (Priority)
Off. HiPrio           : 0                       0
Off. LowPrio          : 0                       0
Off. Uncolor          : 0                       0
Dro. HiPrio           : 0                       0
Dro. LowPrio          : 0                       0
For. InProf           : 0                       0
For. OutProf          : 0                       0

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

Egress Queue 1
Dro. InProf           : 880                     127660
Dro. OutProf          : 0                       0
```

```
For. InProf           : 90862                  12995616
For. OutProf          : 0                      0

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

Egress Queue 3
Dro. InProf           : 0                      0
Dro. OutProf          : 0                      0
For. InProf           : 0                      0
For. OutProf          : 0                      0
===============================================================================
A:Dut-A#


A:Dut-A# monitor service subscriber alcatel_100 sap 1/2/1:101 sla-profile sla_default
base rate
===============================================================================
Monitor statistics for Subscriber alcatel_100
===============================================================================
At time t = 0 sec (Base Statistics)
-------------------------------------------------------------------------------
SLA Profile Instance statistics
-------------------------------------------------------------------------------
                        Packets                  Octets
Off. HiPrio           : 0                      0
Off. LowPrio          : 109099                 35427060
Off. Uncolor          : 0                      0

Queueing Stats (Ingress QoS Policy 1000)
Dro. HiPrio           : 0                      0
Dro. LowPrio          : 8449                   2894798
For. InProf           : 0                      0
For. OutProf          : 100523                 32489663

Queueing Stats (Egress QoS Policy 1000)
Dro. InProf           : 880                    127660
Dro. OutProf          : 0                      0
For. InProf           : 105578                 15104553
For. OutProf          : 0                      0
-------------------------------------------------------------------------------
At time t = 11 sec (Mode: Rate)
-------------------------------------------------------------------------------
SLA Profile Instance statistics
-------------------------------------------------------------------------------
                        Packets                  Octets                   % Port
                                                                          Util.
Off. HiPrio           : 0                      0                        0.00
Off. LowPrio          : 1469                   477795                   0.38
Off. Uncolor          : 0                      0                        0.00

Queueing Stats (Ingress QoS Policy 1000)
Dro. HiPrio           : 0                      0                        0.00
Dro. LowPrio          : 119                    40691                    0.03
For. InProf           : 0                      0                        0.00
For. OutProf          : 1349                   437350                   0.34

Queueing Stats (Egress QoS Policy 1000)
```

```
Dro. InProf           : 0                      0                      0.00
Dro. OutProf          : 0                      0                      0.00
For. InProf           : 1469                   209129                 0.16
For. OutProf          : 0                      0                      0.00
===============================================================================
A:Dut-A#
A:Dut-A# monitor service subscriber alcatel_100 sap 1/2/1:101 sla-profile sla_default
ingress-queue-id 1
===============================================================================
Monitor statistics for Subscriber alcatel_100
===============================================================================
At time t = 0 sec (Base Statistics)
-------------------------------------------------------------------------------
                        Packets                Octets
Ingress Queue 1 (Unicast) (Priority)
Off. HiPrio           : 0                      0
Off. LowPrio          : 0                      0
Off. Uncolor          : 0                      0
Dro. HiPrio           : 0                      0
Dro. LowPrio          : 0                      0
For. InProf           : 0                      0
For. OutProf          : 0                      0
===============================================================================
A:Dut-A#


A:Dut-A# monitor service subscriber alcatel_100 sap 1/2/1:101 sla-profile sla_default
egress-queue-id 1
===============================================================================
Monitor statistics for Subscriber alcatel_100
===============================================================================
At time t = 0 sec (Base Statistics)
-------------------------------------------------------------------------------
                        Packets                Octets
Egress Queue 1
Dro. InProf           : 880                    127660
Dro. OutProf          : 0                      0
For. InProf           : 164366                 23506178
For. OutProf          : 0                      0
===============================================================================
A:Dut-A#
```

host

Syntax **host** [**sap** *sap-id*] [**wholesaler** *service-id*] [**port** *port-id*] [**inter-dest-id** *intermediate-destination-id*] [**detail**]
host [**sap** *sap-id*] [**wholesaler** *service-id*] [**port** *port-id*] **no-inter-dest-id** [**detail**]
host summary
host [**detail**] **wholesaler** *service-id* (VPRN only)

Context show>service>id

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

Parameters **sap** *sap-id* — Displays SAP information for the specified SAP ID. See Common Service Commands on page 1740 for *sap-id* command syntax.

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

Values Up to 32 characters maximum

summary — Displays summary host information.

wholesaler *service-id* **—** The VPRN service ID of the wholesaler. When specified in this context, SAP, SDP, interface, IP address and MAC parameters are ignored.

Values *service-id*: 1 — 2147483647
svc-name: 64 characters maximum