



CHAPTER 6

Configuring the Line Interface

This module describes how to configure the physical line interfaces (ports) as well as how to configure those interfaces for tunneling, VLAN translation, DSCP marking, and traffic rules.

- [Line Interfaces, page 6-1](#)
- [Tunneling Protocols, page 6-2](#)
- [Configuring Traffic Rules and Counters, page 6-8](#)
- [DSCP Marking, page 6-16](#)
- [Counting Dropped Packets, page 6-16](#)

Line Interfaces

The Line Interfaces (Subscriber and Network) are used to connect the SCE platform to the network. See the description of network topologies in the *Topology* section of the [Cisco SCE8000 Installation and Configuration Guide](#).

Flow Control and Bandwidth Considerations



Note

By design, the SCE platform reacts to Ethernet flow control and does not activate it. Therefore, it is possible for a situation to arise in which flow control actually stalls the SCE platform by overflowing the SCE platform queues, thereby causing traffic to be dropped on the Rx interfaces. If this situation persists for more than five seconds, it may trigger the internal sanity checks mechanism within the SCE platform, which may in turn trigger a reload of the SCE platform in an attempt to recover

Maximum Packet Size

The MTU value for the Cisco SCE8000 traffic processing is 9238 bytes. However, in the current version, packets larger than 1600 bytes are bypassed and are not handled by the service control application.

How to Configure the Ten Gigabit Ethernet Line Interfaces

The 10GBE line interfaces are configured in TenGigabitEthernet mode.

-
- Step 1** At the SCE# prompt, type **configure** and press **Enter**.
Enters Global Configuration mode.
- Step 2** At the SCE(config)# prompt, type **interface TenGigabitEthernet 3/bay number/0** and press **Enter**, where *bay number* is the number of the selected SPA bay (0-3). Currently, the slot number is always 3 and the port number is always 0.
Enters Interface Configuration mode for the selected 10 GBE interface.
- Step 3** At the SCE(config-if)# prompt, type **exit** and press **Enter**.
Exits to global configuration mode, from which you can access a different Ten Gigabit Ethernet interface.
-

Tunneling Protocols

- [Selecting the Tunneling Mode, page 6-3](#)
- [Asymmetric L2 Support, page 6-7](#)
- [How to Display the Tunneling Configuration, page 6-8](#)

Tunneling technology is used across various telecommunications segments to solve a wide variety of networking problems. The SCE platform is designed to recognize and process various tunneling protocols in several ways. The SCE platform is able to either ignore the tunneling protocols ("skip" the header) or treat the tunneling information as subscriber information ("classify").

The following table shows the support for the various tunneling protocols (the default behavior for each protocol is in bold):

Table 6-1 Tunneling Protocol Summary

Protocol	Supported handling	Mode name	Symmetric/ Asymmetric
L2TP	Ignore tunnel	IP-tunnel L2TP skip	asymmetric
	Don't ignore tunnel – classify by external IP	No IP-Tunnel	symmetric
IPinIP	Ignore tunnel	ip-tunnel IPinIP skip	symmetric
	Don't ignore tunnel – classify by external IP	no ip-tunnel IPinIP skip	symmetric
VLAN	Ignore tunnel	VLAN symmetric skip	symmetric
	Ignore tunnel – asymmetric	VLAN a-symmetric skip	asymmetric
	VLAN tag used for VPN classification	VLAN symmetric classify	symmetric
MPLS	Ignore tunnel (inject unlabeled)	MPLS traffic-engineering skip	symmetric

When the tunneling information is ignored, the subscriber identification is the subscriber IP of the IP packet carried inside the tunnel.

Asymmetric Tunneling

Some tunneling modes are symmetric and some are asymmetric (see [Table 6-1](#)). Any time that one of the asymmetric tunneling modes is enabled, the entire system is automatically set to asymmetric flow open mode. In this mode, flows are opened earlier than in symmetric flow open mode, and the first packet of each direction of the flow (upstream and downstream) reaches the software. This is required to support redirect and block operations over asymmetric layer 2 protocols. However, it also has some performance impact, so that a certain performance degradation should be expected in any asymmetric mode.

It is also possible to explicitly configure the system to treat all flows as having asymmetric layer 2 characteristics (including Ethernet, VLAN, MPLS, and L2TP), for the purpose of packet injection (such as block flow and redirect flow operations).

To view the effective flow open mode, use the **show interface linecard 0 flow-open-mode** command.



Note

For directions on how to configure the asymmetric tunneling option, see [Asymmetric L2 Support, page 6-7](#)

L2TP

L2TP is an IP-based tunneling protocol, therefore the system must be specifically configured to recognize the L2TP flows, given the UDP port used for L2TP. The SCE platform can then skip the external IP, UDP, and L2TP headers, reaching the internal IP, which is the actual subscriber traffic. If L2TP is not configured, the system treats the external IP header as the subscriber traffic, thus all the flows in the tunnel are seen as a single flow.

VLAN

A single VLAN tag is supported per packet (no QinQ support).

Subscriber classification by VLAN tag is supported only in symmetric VLAN environments – i.e. where the upstream and downstream tags of a flow are identical.

Selecting the Tunneling Mode

- [How to Configure IP Tunnels, page 6-4](#)
- [IPinIP Tunneling, page 6-4](#)
- [How to Configure the VLAN Environment, page 6-6](#)
- [How to Configure the L2TP Environment, page 6-6](#)

Use these commands to configure tunneling:

- **ip-tunnel**
- **vlan**
- **L2TP identify-by**

How to Configure IP Tunnels

By default, IP tunnel recognition is disabled. Use this command to configure recognition of L2TP tunnels and skipping into the internal IP packet.

Step 1 From the SCE(config-if)# prompt, type **ip-tunnel L2TP skip** and press **Enter**.

To disable IP tunnels, use the following command:

Step 1 From the SCE(config-if)# prompt, type **no ip-tunnel** and press **Enter**.



Note

The IP tunneling mode is mutually exclusive with VPN-based classification.

IPinIP Tunneling

- [Enabling IPinIP Tunneling, page 6-4](#)
- [Disabling IPinIP Tunneling, page 6-5](#)
- [DSCP Marking for IPinIP Tunnels, page 6-5](#)

IPinIP is an IP-based tunneling protocol; therefore the system must be specifically configured to recognize the flows inside the tunnel. The SCE platform will then skip the external IP header, reaching the internal IP, which is the actual subscriber traffic. When IPinIP skip is disabled, the system treats the external IP header as the subscriber traffic, resulting in all IPinIP traffic being reported as generic IP.

Guidelines for configuring IPinIP tunnels:

- IPinIP and other tunnels: IPinIP is supported simultaneously with plain IP traffic and any other tunneling protocol supported by the SCE platform.
- Overlapping IP addresses: There is no support for overlapping IP addresses within different IPinIP tunnels.
- DSCP marking: For IPinIP traffic, DSCP marking can be done on either the external or the internal IP header exclusively. See [Disabling IPinIP Tunneling, page 6-5](#).

Fragmentation

Fragmentation should be avoided whenever possible. If it is not possible to avoid fragmentation, it is recommended to opt for internal fragmentation. If that is also not possible, the SCE platform can be operated under conditions of external fragmentation

Enabling IPinIP Tunneling

By default, IP tunnel recognition is disabled. Use this command to configure recognition of IPinIP tunnels and skipping into the internal IP packet.

Step 1 From the SCE(config-if)# prompt, type **ip-tunnel IPinIP skip** and press **Enter**.

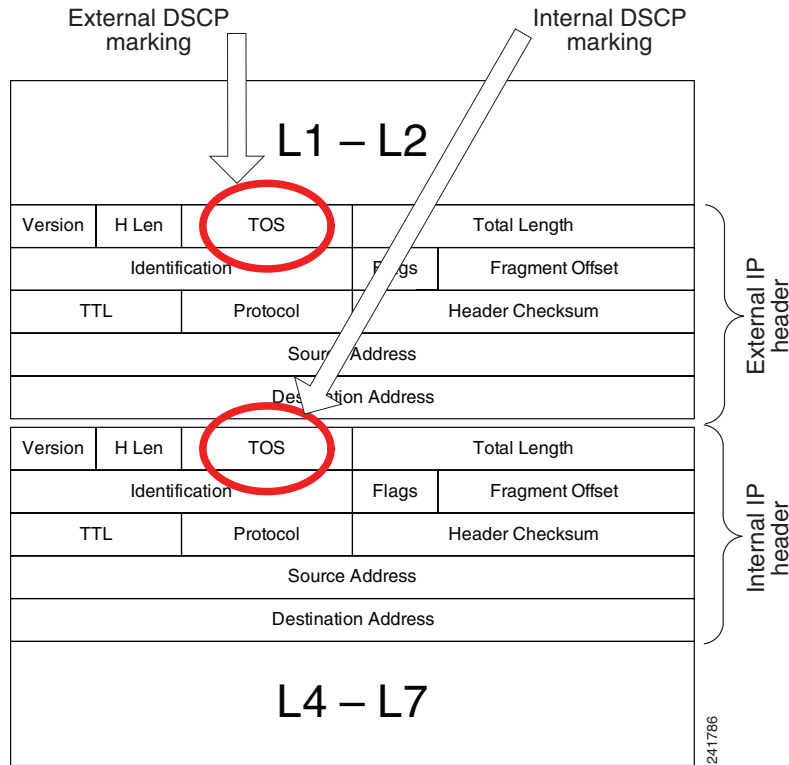
Disabling IPinIP Tunneling

Step 1 From the SCE(config-if)# prompt, type **no ip-tunnel IPinIP skip** and press **Enter**.

DSCP Marking for IPinIP Tunnels

DSCP marking modifies the DSCP bits of the IPv4 header. In IPinIP tunnels there are at least two IP headers. By default, DSCP marking is performed only on the external IP header (refer to the illustration below).

Figure 6-1 DSCP Marking for IPinIP Tunnels



Note

DSCP marking should be enabled and configured through SCA BB console. Refer to the section “How to Manage DSCP Marker Values” in the chapter “Using the Service Configuration Editor: Traffic Control” in the *Cisco Service Control Application for Broadband User Guide* for further information.

Use this command to configure the SCE platform to mark the DSCP bits of the internal IP header. This command takes effect only when *IPinIP skip* is enabled.

Step 1 From the SCE(config-if)# prompt, type **ip-tunnel IPinIP DSCP-marking-skip** and press **Enter**. Enables DSCP marking on the internal IP header of IPinIP traffic.

To perform DSCP marking on the external IP header, use the following command:

-
- Step 1** From the SCE(config-if)# prompt, type **no ip-tunnel IPinIP DSCP-marking-skip** and press **Enter**.
-

How to Configure the VLAN Environment

Use this command to configure the VLAN environment.

- [Options, page 6-6](#)
- [Configuring the VLAN Environment: Example, page 6-6](#)

Options

There are three options:

- **symmetric classify**
- **symmetric skip** (default)
- **a-symmetric skip**

Symmetric environment refers to an environment in which the same VLAN tags are used for carrying a transaction in the upstream and downstream directions

Setting the mode to classify means that VPN and flow classification will use the VLAN tag. Using VLAN classification is mutually exclusive with other tunnel-based classification or IP tunnels.

An a-symmetric environment is an environment in which the VLAN tags might not be the same in the upstream and downstream directions of the same flow.

The SCE platform is configured by default to work in symmetric environments. A specific command should be used to allow correct operation of the SCE platform in asymmetric environments and instruct it to take into consideration that the upstream and downstream of each flow has potentially different VLAN tags.



Note

Using the a-symmetric skip value incurs a performance penalty.

- Step 1** From the SCE(config if)# prompt, type **vlan {symmetric classify | symmetric skip | a-symmetric skip}** and press **Enter**.

Specify the desired VLAN mode.

Configuring the VLAN Environment: Example

The following example selects VLAN-based classification.

```
SCE(config if)#vlan symmetric classify
```

How to Configure the L2TP Environment

- [External Fragmentation in the L2TP Environment, page 6-7](#)
- [Options, page 6-7](#)

External Fragmentation in the L2TP Environment

If external fragmentation exists in the L2TP environment, it is required to configure a *quick-forwarding-ignore* traffic rule (see [Configuring Traffic Rules and Counters, page 6-8](#)) that bypasses all IP traffic targeted to either the LNS or LAC IP address. This will make sure that any packets not having the L2TP port indication (i.e. non-first fragments) will not require handling by the traffic processors.

In addition, to prevent reordering of L2TP tunneled fragments, it is advised to define a *quick-forwarding* traffic rule for all the L2TP traffic. This can be done based on the IP ranges in use by the internal IPs in the tunnel (as allocated by the LNS), or simply for all the traffic passing through the SCE platform.

**Note**

Flow redirection and flow blocking cannot be performed on quick-forwarded traffic.

Options

The following option is available:

- **portnumber** —The port number that the LNS and LAC use for L2TP tunnels.
Default port# = 1701

Step 1 From the SCE(config if)# prompt, type **L2TP identify-by port-number** *portnumber* and press **Enter**.

Asymmetric L2 Support

You should enable asymmetric layer 2 support in cases where the following conditions apply for any flows:

- Each direction of the flow has a different MAC address.
- The routers do not accept packets with the MAC address of the other link.

**Note**

'Asymmetric routing topology' support and 'asymmetric tunneling support' are two separate features. Asymmetric routing topology refers to topologies where the SCE platform might see some flows only in one direction (upstream/downstream). Asymmetric tunneling support (asymmetric L2 support) refers to the ability to support topologies where the SCE platform sees both directions of all flows, but some of the flows may have different layer 2 characteristics (like MAC addresses, VLAN tags, MPLS labels and L2TP headers), which the SCE platform must specifically take into account when injecting packets into the traffic (such as in block and redirect operations).

In order to support asymmetric layer 2, the SCE platform switches to *asymmetric flow open* mode, which incurs a certain performance penalty. This is NOT the case for asymmetric routing topology.

Step 1 From the SCE(config-if)# prompt, type **asymmetric-L2-support** and press **Enter**.

How to Display the Tunneling Configuration

-
- Step 1** From the SCE# prompt, type **show interface linecard 0 (VLANIL2TP|IP-tunnel)** and press **Enter**.
Displays the current configuration for the specified tunnel option.
-

Configuring Traffic Rules and Counters

- [Traffic Rules and Counters, page 6-8](#)
- [Configuring Traffic Counters, page 6-10](#)
- [Configuring Traffic Rules, page 6-11](#)
- [Managing Traffic Rules and Counters, page 6-14](#)

Traffic Rules and Counters

- [What are Traffic Rules and Counters?, page 6-8](#)
- [Traffic Rules, page 6-9](#)
- [Traffic Counters, page 6-10](#)

What are Traffic Rules and Counters?

Traffic rules and counters may be configured by the user. This functionality enables the user to define specific operations on the traffic flowing through the SCE Platform, such as blocking or ignoring certain flows or counting certain packets. The configuration of traffic rules and counters is independent of the application loaded by the SCE platform, and thus is preserved when the application being run by the SCE platform is changed.

Possible uses for traffic rules and counters include:

- Enabling the user to count packets according to various criteria. Since the traffic counters are readable via the *ciscoServiceControlTpStats* MIB, these might be used to monitor up to 32 types of packets, according to the requirements of the installation.
- Ignoring certain types of flows. When a traffic rules specifies an “ignore” action, packets matching the rule criteria will not open a new flow, but will pass through the SCE platform without being processed. This is useful when a particular type of traffic should be ignored by the SCE platform.
Possible examples include ignoring traffic from a certain IP range known to require no service, or traffic from a certain protocol.
- Blocking certain types of flows. When a traffic rules specifies a “block” action, packets matching the rule criteria (and not belonging to an existing flow) will be dropped and not passed to the other interface. This is useful when a particular type of traffic should be blocked by the SCE platform.
Possible examples include performing ingress source address filtering (dropping packets originating from a subscriber port whose IP address does not belong to any defined subscriber-side subnet), or blocking specific ports.

It should be noted that using traffic rules and counters does not affect performance. It is possible to define the maximum number of both traffic rules and counters without causing any degradation in the SCE platform performance.

Traffic Rules

A traffic rule specifies that a defined action should be taken on packets processed by the SCE Platform that meet certain criteria. The maximum number of rules for the Cisco SCE8000 is 64, which includes not only traffic rules configured via the SCE platform CLI, but also any additional rules configured by external management systems, such as SCA BB. Each rule is given a name when it is defined, which is then used when referring to the rule.

Packets are selected according to user-defined criteria, which may be any combination of the following:

- **IP address** — A single address or a subnet range can be specified for each of the line ports (Subscriber / Network).
- **Protocol** — TCP/UDP/ICMP/IGRP/EIGRP/IS-IS/OSPF/Other
- **TCP/UDP Ports** — A single port or a port range can be specified for each of the line ports (Subscriber / Network). Valid for the TCP/UDP protocols only.
- **Direction (Upstream/Downstream)** (TCP only).

The possible actions are:

- **Count** the packet by a specific traffic counter
- **Block** the packet (do not pass it to the other side)
- **Ignore** the packet (do not provide service for this packet — No bandwidth metering, transaction reporting etc. is done)
- **Quick-forward** the packet **with service** — forward delay-sensitive packets through the fast path while maintaining serviceability for these packets
- **Quick-forward** the packet **with no service (quick-forwarding-ignore)**— forward delay-sensitive packets through the fast path with no service provided for these packets

Block and **Ignore** actions affect only packets that are not part of an existing flow.



Note

Block and **Ignore** are mutually exclusive. However, blocked or ignored packets can also be counted.

It is possible for a single packet to match more than one rule (The simplest way to cause this is to configure two identical rules with different names). When this happens, the system operates as follows:

- Any counter counts a specific packet only once. This means that:
 - If two rules specify that the packet should be counted by the same counter, it is counted only once.
 - If two rules specify that the packet should be counted by different counters, it is counted twice, once by each counter.
- **Block** takes precedence over **Ignore** — If one rule specifies **Block**, and another rule specifies **Ignore**, the packet is blocked.

Traffic Counters

Traffic counters count the traffic as specified by the traffic rules. The maximum number of counters is 32. Each counter is given a name when it is defined, which is then used when referring to the counter.

A traffic counter can be configured in one of two ways:

- **Count packets** — the counter is incremented by 1 for each packet it counts.
- **Count bytes** — the counter is incremented by the number of bytes in the packet for each packet it counts.

Configuring Traffic Counters

A traffic counter must be created before it can be referenced in a traffic rule. Use the following commands to create and delete traffic counters.

- [How to Create a Traffic Counter, page 6-10](#)
- [How to Delete a Traffic Counter, page 6-10](#)
- [How to Delete all Existing Traffic Counters, page 6-11](#)

How to Create a Traffic Counter

Options

The following options are available:

- **name** —The name of the counter
- **Count packets** — the counter is incremented by 1 for each packet it counts.
- **Count bytes** — the counter is incremented by the number of bytes in the packet for each packet it counts.

Step 1 From the SCE(config if)# prompt, type **traffic-counter name name count-bytes|count-packets** and press **Enter**.

Adds a traffic counter with the specified name and counting mode.

How to Delete a Traffic Counter

Step 1 From the SCE(config if)# prompt, type **no traffic-counter name name** and press **Enter**.



Note A traffic counter cannot be deleted if it is used by any existing traffic rule.

How to Delete all Existing Traffic Counters

- Step 1** From the SCE(config if)# prompt, type **no traffic-counter all** and press **Enter**.
Removes all traffic counters.



Note A traffic counter cannot be deleted if it is used by any existing traffic rule.

Configuring Traffic Rules

Use the following commands to create and delete traffic rules.

- [How to Create a Traffic Rule, page 6-11](#)
- [How to Delete a Traffic Rule, page 6-14](#)
- [How to Delete all Traffic Rules, page 6-14](#)

How to Create a Traffic Rule

Options

The following options are available:

IP specification:

```
all|([all-but] (ip-address|ip-range))
```

- *ip-address* is a single IP address in dotted-decimal notation, such as 10.1.2.3
- *ip-range* is an IP subnet range, in the dotted-decimal notation followed by the number of significant bits, such as 10.1.2.0/24.
- Use the **all-but** keyword to exclude the specified IP address or range of IP addresses

protocol:

Any one of the following protocols:

```
TCP/UDP/ICMP/IGRP/EIGRP/IS-IS/OSPF/all
```

port specification:

```
all|([all-but] (port#|port-range))
```

- Specify the ports only if the protocol is either TCP or UDP.
- Specify the port or port range for both the subscriber-side and the network-side.
- Specify a range of ports using the form MinPort:MaxPort.
- Use the **all-but** keyword to exclude the specified port or range of ports

id specification:

```
all|([all-but] tunnel id)
```

- tunnel id is an 8-bit Hex value range, in the format '(HEX) *Tunnel-id*' or '(HEX) *MinTunnelId*:(HEX) *MaxTunnelId*', which reflects the lower eight bits of the VLAN tag.
- Tunnel-ID-based rules can only be used in " *VLAN symmetric classify* " mode (see [How to Configure the VLAN Environment, page 6-6](#), and only when *tunnel id* mode is enabled.

Use the **traffic-rule tunnel-id-mode** command.



Note The VLAN tag itself is a 12-bit value, and therefore aliasing of the lower 8 bits can occur, depending on the VLAN tags used.

direction:

Any of the following:

```
upstream/downstream/both
```

traffic-counter:

Either of the following:

- **name** *<name of an existing traffic counter>*— Packets meeting the criteria of the rule are to be counted in the specified counter. If a counter name is defined, the “count” action is also defined implicitly. The keyword **name** must appear as well as the actual name of the counter.
- **none** — If **none** is specified, then an action must be explicitly defined via the action option.

action: (not required if the action is count only)

One of the following:

- **block** — Block the specified traffic
- **ignore** — Bypass the specified traffic; traffic receives no service
- **quick-forwarding** — Forward delay-sensitive packets through the fast path while maintaining serviceability for these packets
- **quick-forwarding-ignore** — Forward delay-sensitive packets through the fast path with no service provided for these packets

Step 1 From the SCE(config if)# prompt, type **traffic-rule name name IP-addresses (all(subscriber-side <IP specification> network-side <IP specification>)) protocol protocol [ports subscriber-side <port specification> network-side <port specification>] [tunnel-id <tunnel-id specification>] direction direction traffic-counter <traffic-counter>[action action]**

Configuring Traffic Rules: Examples

- [Example 1, page 6-13](#)
- [Example 2, page 6-13](#)
- [Example 3, page 6-13](#)

Example 1

This example creates the following traffic rule:

- Name = rule1
- IP addresses: subscriber side = all IP addresses, network side = 10.10.10.10 only
- Protocol = all
- Direction = both
- Traffic counter = counter1
- The only action performed will be counting

```
SCE(config if)# traffic-rule name rule1 IP-addresses subscriber-side all network-side
10.10.10.10 protocol all direction both traffic-counter name counter1
```

Example 2

This example creates the following traffic rule:

- Name = rule2
- IP addresses: subscriber side = all IP addresses, network side = all IP addresses EXCEPT the subnet 10.10.10.0/24
- Protocol = TCP
- Ports: subscriber-side = 100-200, network-side = all
- Tunnel id = all
- Direction = downstream
- Traffic counter = counter2
- Action = Block
- The actions performed will be counting and blocking

The first command enables tunnel id mode.

```
SCE(config if)#traffic-rule tunnel-id-mode
SCE(config if)# traffic-rule name rule2 IP-addresses subscriber-side all network-side
all-but 10.10.10.0/24 protocol tcp ports subscriber-side 100:200 network-side all
tunnel-id all direction downstream traffic-counter name counter2 action block
```

Example 3

This example creates the following traffic rule:

- Name = rule3
- IP addresses: all
- Protocol = IS-IS
- Direction = upstream
- Traffic counter = none
- Action = ignore (required since traffic-counter = none)
- The only action performed will be **Ignore**.

```
SCE(config if)# traffic-rule name rule3 IP-addresses all protocol IS-IS direction upstream
traffic-counter none action ignore
```

How to Delete a Traffic Rule

-
- Step 1** From the SCE(config if)# prompt, type **no traffic-rule name** *name* and press **Enter**.
Removes the specified traffic rule.
-

How to Delete all Traffic Rules

-
- Step 1** From the SCE(config if)# prompt, type **no traffic-rule all** and press **Enter**.
Removes all existing traffic rules.
-

Managing Traffic Rules and Counters

Use these commands to display existing traffic rule configuration, as well as traffic counter configuration (packets/bytes and the name of the rule using the counter) and traffic counter value.

You can also reset a specific counter or all counters.

- [How to View a Specified Traffic Rule, page 6-14](#)
- [How to View all Traffic Rules, page 6-14](#)
- [How to View a Specified Traffic Counter, page 6-15](#)
- [How to View all Traffic Counters, page 6-15](#)
- [How to Reset a Specified Traffic Counter, page 6-15](#)
- [How to Reset all Traffic Counters, page 6-15](#)

How to View a Specified Traffic Rule

-
- Step 1** From the SCE# prompt, type **show interface linecard 0 traffic-rule name** *rule-name* and press **Enter**.
Displays the configuration of the specified traffic rule.
-

How to View all Traffic Rules

-
- Step 1** From the SCE# prompt, type **show interface linecard 0 traffic-rule all** and press **Enter**.
Displays the configuration of all existing traffic rules.
-

How to View a Specified Traffic Counter

-
- Step 1** From the SCE# prompt, type **show interface linecard 0 traffic-counter name** *counter-name* and press **Enter**.
Displays the value of the specified counter and lists the traffic rules that use it.
-

Viewing a Traffic Counter: Example

The following example displays information for the traffic counter “cnt”.

```
SCE# show interface linecard 0 traffic-counter name cnt
Counter 'cnt' value: 0 packets. Rules using it: None.
```

How to View all Traffic Counters

-
- Step 1** From the SCE# prompt, type **show interface linecard 0 traffic-counter all** and press **Enter**.
Displays the value of the each counter and lists the traffic rules that use it.
-

Viewing the Traffic Counters: Example

The following example displays information for all existing traffic counters.

```
SCE# show interface linecard 0 traffic-counter all
Counter 'cnt' value: 0 packets. Rules using it: None.
Counter 'cnt2' value: 0 packets. Rules using it: Rule2.
2 counters listed out of 32 available.
```

How to Reset a Specified Traffic Counter

-
- Step 1** From the SCE# prompt, type **clear interface linecard 0 traffic-counter name** *counter-name* and press **Enter**.
Resets the specified traffic counter.
-

How to Reset all Traffic Counters

-
- Step 1** From the SCE# prompt, type **clear interface linecard 0 traffic-counter all** and press **Enter**.
Resets all traffic counters.
-

DSCP Marking

DSCP marking is used in IP networks as a means to signal the priority of a packet. The Cisco Service Control solution supports the DSCP classification on a per-service, per-package level via the SCA BB application. The SCE platform DSCP marking feature enables marking the DSCP field in the IP header of each packet according to the policy configured via the SCA BB console. The actual DSCP value set in the IP header is determined according to the value defined in a configurable DSCP translation table.

DSCP marking configuration is performed via the SCA BB console, The SCE platform CLI allows you to view the state of DSCP marking (enabled or disabled) for each interface and to display the DSCP translation table.

For information on configuring DSCP marking, please refer to the *Cisco Service Control Application for Broadband User Guide* (for any SCA BB release 3.1.5 or later).



Note

DSCP marking in release 3.1.5 or later is not backwards compatible with any SCOS version prior to release 3.1.5.

How to Display the DSCP Marking Configuration

Use this command to display the state of DSCP marking (enabled or disabled) per interface and the DSCP translation table.

Step 1 From the SCE> prompt, type **show interface linecard 0 ToS-marking** and press **Enter**.

Counting Dropped Packets

- [About Counting Dropped Packets, page 6-16](#)
- [Disabling the Hardware Packet Drop, page 6-17](#)

About Counting Dropped Packets

By default, the SCE platform hardware drops WRED packets (packets that are marked to be dropped due to BW control criteria). However, this presents a problem for the user who needs to know the number of dropped packets per service. To be able to count dropped packets per service, the traffic processor must see all dropped packets for all flows. However, if the hardware is dropping red packets, the traffic processor will not be able to count all dropped packets and the user will not get proper values on the relevant MIB counters (*tpTotalNumWredDiscardedPackets*).



Note

The MIB object *tpTotalNumWredDiscardedPackets* counts dropped packets. The value in this counter is absolute only when hardware packet drop is disabled (not the default mode). When hardware packet drop is enabled (default mode), this MIB counter provides only a relative value indicating the trend of the number of packet drops, with a factor of approximately 1:6.

The user can disable the drop-wred-packets-by-hardware mode. This allows the application to access existing per-flow counters. The application can then retrieve the number of dropped packets for every flow and provide the user with better visibility into the exact number of dropped packets and their distribution.

**Note**

Counting all dropped packets has a considerable effect on system performance, and therefore, by default, the drop-wred-packets-by-hardware mode is enabled.

Disabling the Hardware Packet Drop

Use this command to disable the drop-wred-packets-by-hardware mode, enabling the software to count all dropped packets.

By default hardware packet drop is enabled.

**Note**

Disabling this feature may have both delay and performance implications.

- Step 1** From the SCE(config if)# prompt, type **no accelerate-packet-drops** and press **Enter**.
Disables hardware packet drop.

To enable hardware packet drop, use the following command:

- Step 1** From the SCE(config if)# prompt, type **accelerate-packet-drops** and press **Enter**.

