



Configuring Line Interfaces

Revised: August 26, 2014

Introduction

This chapter describes how to configure the physical line interfaces (ports) as well as how to configure those interfaces for tunneling, DSCP marking, and traffic rules:

- [Line Interfaces, page 7-2](#)
- [Tunneling Protocols, page 7-5](#)
- [Managed VPNs, page 7-21](#)
- [Configuring Traffic Rules and Counters, page 7-25](#)
- [DSCP Marking, page 7-35](#)
- [Counting the Dropped Packets, page 7-36](#)

Line Interfaces

- [Information About Line Interfaces, page 7-2](#)
- [Configuring the Line Interfaces, page 7-2](#)

The Line Interfaces (Subscriber and Network) are used to connect the SCE platform to the network. See the description of network topologies in the “Cisco SCE8000 GBE Topology and Topology-Related Parameters” chapter of the *Cisco SCE8000 GBE Installation and Configuration Guide*.

Information About Line Interfaces

The Cisco SCE 8000 GBE line interfaces are found on the 8-port Gigabit Ethernet SPAs installed in subslots 0 and 1 of slot 3. Each 8-port Gigabit Ethernet SPA provides eight GBE ports, which interface with either subscriber or network traffic. These interfaces can be individually configured using CLI commands in this section.

Flow Control and Bandwidth Considerations

**Note**

By design, the SCE platform reacts to Ethernet flow control and does not activate it. Therefore, a situation could arise in which flow control 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. However, the technical support file may not be generated when the SCE platform recovers after reloading several times.

Maximum Packet Size

The MTU value for the Cisco SCE 8000 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.

Configuring the Line Interfaces

The GBE line interfaces are configured in GigabitEthernet mode. You must enter GigabitEthernet mode for the desired interface. This gives you access to the following configuration commands for that interface:

- **auto-negotiate**
- **global-controller bandwidth**
- **global-controller name**

You can also configure a range of GigabitEthernet line interfaces using the **interface range** command. This command allows you to specify a range of interfaces. Any of the three configuration commands listed above are then applied to all interfaces in the specified range.

How to Configure a Specific Gigabit Ethernet Line Interface

-
- Step 1** At the SCE8000# prompt, type **configure** and press **Enter**.
Enters Global Configuration mode.
- Step 2** At the SCE8000(config)# prompt, type **interface GigabitEthernet 3/bay number/port number** and press **Enter**.
Enters Interface Configuration mode for the selected GBE interface.
- *bay number* is the number of the selected SPA bay (0 or 1)
 - *port number* is the number of the selected interface (0-7)
 - Currently, the slot number is always 3.
- Step 3** At the SCE8000(config if)# prompt, type **exit** and press **Enter**.
Exits to global configuration mode, from which you can access a different Gigabit Ethernet interface.
-

How to Configure a Range of Gigabit Ethernet Line Interfaces

-
- Step 1** At the SCE8000# prompt, type **configure** and press **Enter**.
Enters Global Configuration mode.
- Step 2** At the SCE8000(config)# prompt, type **interface range GigabitEthernet 3/bay range/port range** and press **Enter**.
Enters Interface Configuration mode for the selected range of GBE interfaces.
- *bay range* can be any of the following: 0,1, or 0-1.
 - *port range* can be any range of the interface numbers between 0 and 7. It can also be any specific port number in that range.
 - Currently, the slot number is always 3.
- Step 3** At the SCE8000(config if range)# prompt, type **exit** and press **Enter**.
Exits to global configuration mode, from which you can access a different Gigabit Ethernet interface.
-

Configuring a Range of Gigabit Ethernet Line Interfaces: Example

This example illustrates how to configure ports 3-6 on both the Cisco SCE 8000 SPA modules.

```
SCE8000# configure
SCE8000(config)# interface range GigabitEthernet 3/0-1/3-6
SCE8000 (config if range)#
```

How to Configure the Gigabit Ethernet Line Interfaces for a Specified Cisco SCE 8000 of a Cascaded Pair

-
- Step 1** At the SCE8000# prompt, type **configure** and press **Enter**.
Enters Global Configuration mode.
- Step 2** Specify the ID of the SCE 8000 platform in either the **interface range GigabitEthernet** or **interface GigabitEthernet** command, as follows:
- **interface range GigabitEthernet** *sce-id/3/bay range/port range*
 - **interface GigabitEthernet** *sce-id/3/bay number/port number*
- Here, sce-id is the ID of the SCE 8000 platform in the cascaded pair (0 or 1).
-

Configuring the Gigabit Ethernet Line Interfaces for a Specified Cisco SCE 8000: Example

This example illustrates how to configure ports 3-6 on both Cisco SCE 8000 SPA modules on Cisco SCE platform #1 of a cascaded pair:

```
SCE8000# configure
SCE8000(config)# interface range GigabitEthernet 1/3/0-1/3-6
SCE (config if range)#
```

Tunneling Protocols

- [Tunneling IPv6 Traffic, page 7-7](#)
- [Selecting the Tunneling Mode, page 7-8](#)
- [Asymmetric L2 Support, page 7-19](#)
- [Displaying the Tunneling Configuration, page 7-19](#)

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*). A special case of classification by tunneling information is VPN with private IP support.

Only 6to4, 6rd, and DS-Lite tunnels are supported for IPv6 traffic. There is no **skip** command for IPv6 tunnels.

Table 7-1 shows the support for the various tunneling protocols (the default behavior for each protocol is in bold):

Table 7-1 Tunneling Protocol Summary

Protocol	Supported Handling	Mode Name	Symmetric or Asymmetric
6to4	Enables the 6to4 mode.	IP-tunnel 6to4	—
	Disables the 6to4 mode.	no IP-tunnel 6to4	—
6rd	Enables the 6to4 mode.	IP-tunnel 6to4	—
	Disables the 6to4 mode.	no IP-tunnel 6to4	—
DS-Lite	Enables the DS-Lite mode.	IP-tunnel DS-Lite	—
	Disables the DS-Lite mode.	no IP-tunnel DS-Lite	—
L2TP	Ignores the tunnel.	IP-tunnel L2TP skip	Asymmetric
	Does not ignore tunnel – Classify by external IP.	no IP-tunnel	Symmetric
GRE	Ignores the tunnel.	ip-tunnel GRE skip	Symmetric
	Does not ignore tunnel – Classify by external IP.	no ip-tunnel GRE skip	Symmetric
IPinIP	Ignores the tunnel	ip-tunnel IPinIP skip	Symmetric
	Does not ignore tunnel – Classify by external IP	no ip-tunnel IPinIP skip	Symmetric
VLAN	Ignores tunnel	VLAN symmetric skip	Symmetric
	Ignores tunnel – asymmetric	VLAN asymmetric skip	Asymmetric
	VLAN tag is used for VPN classification	VLAN symmetric classify	Symmetric
MPLS	Ignores the tunnel (injects unlabeled)	MPLS traffic-engineering skip	Symmetric
	Ignores the tunnel (injects labeled)	MPLS VPN skip	Asymmetric

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 7-1](#)). 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 has some impact on both performance and capacity, so that a certain performance degradation should be expected in any asymmetric mode.

You can explicitly configure the system to treat all flows as having asymmetric layer 2 characteristics (including Ethernet, VLAN, MPLS, and L2TP).

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” section on page 7-19](#)

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, that is, where the upstream and downstream tags of a flow are identical.

Tunneling IPv6 Traffic

The following tunnels are supported for IPv6 traffic:

- 6to4 and 6rd

Tunneling can be done in two modes:

- IPv6 mode—Traffic is handled as native IPv6-classified IPv6-over-IPv4 service. All the IPv6 flows work in a subscriber-less mode.

In the IPv6 mode, you can configure the IPv6 hash or the IPv4 hash. When the IPv6 hash is configured, the IPv6 traffic is handled by the traffic processor card configured to handle the IPv6 traffic based on the IPv6 hash. When the IPv4 hash is configured, the IPv6 traffic is handled by the traffic processor card configured to handle the IPv6 traffic based on the IPv4 IP hash.

- IPv4 mode—Traffic is handled as IPv4-classified IPv6-over-IPv4 service. Traffic is accounted for and enforced based on the IPv4 subscriber. Hashing is not significant.

- DS-Lite

DS-Lite is a light-weight dual stack that is used when only public IPv6 addresses and private IPv4 addresses can be assigned. The extension header is supported on the DS-Lite tunnels.

When DS-Lite is enabled, the IPv6 traffic is handled as TCP/UDP by the traffic processor configured to handle the IPv6 traffic. If DS-Lite is disabled, the IPv6 traffic is bypassed as non-TCP/UDP by the traffic processor configured to handle the IPv6 traffic.

DS-Lite bundling is supported for FTP traffic. IPv6 addresses and L4 ports are used for binding FTP control flows and FTP data flows. The FTP control flows are classified, and IPv6 addresses and L4 ports are used to create a binding context to bind the data flows to the control flow.

- L2TP

Cisco SCE supports IPv6 over IPv4 L2TP tunnels. In L2TP IPv6 over IPv4 tunnels, the internal L3 header is IPv6 and the external L3 header is IPv4. The Cisco SCE uses internal IPv6 addresses for tasks such as subscriber awareness, classification, load-balancing, congestion management.

Cisco SCE process IPv6 over IPv4 L2TP tunnels when L2TP skip is enabled.

The L2TP IPv6 over IPv4 is considered as tunneled traffic only in Generic Usage RDRs.

Some of the tunneling protocols can be configured in both modes, but some of the protocols work in tunneled mode only. For example, the 6to4 tunneling protocol is supported in both the modes.

The following tunnels are not supported for IPv6 traffic:

- GRE
- VLAN
- MPLS

Selecting the Tunneling Mode

- [Configuring the 6to4 Tunnels, page 7-9](#)
- [Configuring the DS-Lite Tunnels, page 7-10](#)
- [Configuring the L2TP Tunnels, page 7-11](#)
- [Configuring GRE Tunneling, page 7-12](#)
- [Configuring IPinIP Tunneling, page 7-13](#)
- [Configuring DSCP Marking, page 7-14](#)
- [Configuring the 6to4 Environment, page 7-16](#)
- [Configuring the VLAN Environment, page 7-17](#)
- [Configuring the MPLS Environment, page 7-18](#)
- [Configuring the L2TP Environment, page 7-18](#)

Use these commands to configure tunneling:

- **ip-tunnel 6to4**
 - ip-tunnel DS-Lite
 - **ip-tunnel l2tp**
 - ip-tunnel gre
 - **ip-tunnel IPinIP**
 - **ip-tunnel (GRE|IPinIP) DSCP-marking-skip**
 - **vlan**
 - **mpls**
 - **L2TP identify-by**
-

Configuring the 6to4 Tunnels

Before configuring the 6to4 tunnels or 6rd tunnels, verify whether you have configured the 6to4 environment.

You must configure the 6to4 environment to use the 6to4 tunnels. For details on configuring the 6to4 environment, see the [“Configuring the 6to4 Environment” section on page 7-16](#).



Caution

IP tunneling must be enabled or disabled only when no application is loaded or the linecard is shut down.

Enabling 6to4 Tunneling

By default, IP tunnel recognition is disabled. Use the following steps to enable 6to4 tunnels:

-
- Step 1** Shut down the linecard. (This is a root-level command.)
From the SCE8000(config if)#> prompt, enter **shutdown** and press **Enter**.
- Step 2** Enable 6to4 tunneling.
From the SCE8000(config if)#> prompt, enter **ip-tunnel 6to4** and press **Enter**.
- Step 3** Restart the linecard.
From the SCE8000(config if)#> prompt, enter **no shutdown** and press **Enter**.
-

Disabling 6to4 Tunneling

Use these steps to disable 6to4 tunneling:

-
- Step 1** Shut down the linecard. (This is a root level command.)
From the SCE8000(config if)#> prompt, enter **shutdown** and press **Enter**.
- Step 2** Disable 6to4 tunneling.
From the SCE8000(config if)#>prompt, enter **no ip-tunnel 6to4** and press **Enter**.
-

- Step 3** Restart the linecard.
From the SCE8000(config if)#> prompt, enter **no shutdown** and press **Enter**.
-

Configuring the DS-Lite Tunnels



Caution

IP tunneling must be enabled or disabled only when no application is loaded or the linecard is shut down.

Enabling DS-Lite Tunneling

By default, IP tunnel recognition is disabled. Use the following steps to enable the recognition of DS-Lite tunnels:

- Step 1** Shut down the linecard. (This is a root-level command.)
From the SCE8000(config if)#> prompt, enter **shutdown** and press **Enter**.
- Step 2** Enable DS-Lite tunneling.
From the SCE8000(config if)#> prompt, enter **ip-tunnel DS-Lite** and press **Enter**.
- Step 3** Restart the linecard.
From the SCE8000(config if)#> prompt, enter **no shutdown** and press **Enter**.
-

Disabling DS-Lite Tunneling

Use these steps to disable DS-Lite tunneling:

- Step 1** Shut down the linecard. (This is a root-level command.)
From the SCE8000(config if)#> prompt, enter **shutdown** and press **Enter**.
- Step 2** Disable DS-Lite tunneling.
From the SCE8000(config if)#> prompt, enter **no ip-tunnel DS-Lite** and press **Enter**.
- Step 3** Restart the linecard.
From the SCE8000(config if)#> prompt, enter **no shutdown** and press **Enter**.
-

Enabling DS-Lite Extension Header Support

Use the following steps to enable the DS-Lite extension header support:

- Step 1** Shut down the linecard. (This is a root level command.)
From the SCE8000(config if)#> prompt, enter **shutdown** and press **Enter**.
- Step 2** Enable DS-Lite tunneling.
From the SCE8000(config if)#> prompt, enter **ip-tunnel DS-Lite** and press **Enter**.

- Step 3** Enable DS-Lite extension header support.
From the SCE8000(config if)#> prompt, enter **ip-tunnel DS-Lite Extention-Header-Support** and press **Enter**.
- Step 4** Restart the linecard.
From the SCE8000(config if)#> prompt, enter **no shutdown** and press **Enter**.
-

Disabling DS-Lite Extension Header Support

Use these steps to disable DS-Lite extension header support:

- Step 1** Shut down the linecard. (This is a root level command.)
From the SCE8000(config if)#> prompt, enter **shutdown** and press **Enter**.
- Step 2** Disable DS-Lite extension header support.
From the SCE8000(config if)#> prompt, enter **no ip-tunnel DS-Lite Extention-Header-Support** and press **Enter**.
- Step 3** Restart the linecard.
From the SCE8000(config if)#> prompt, enter **no shutdown** and press **Enter**.
-

Enabling DS-lite Extend-IANA to Set 128 Bit Hash Mode

Use these steps to enable DS-lite extend-IANA to set 128 bit hash mode:

- Step 1** Shut down the linecard. (This is a root level command.)
From the SCE8000(config if)#> prompt, enter **shutdown** and press **Enter**.
- Step 2** Enable DS-Lite tunneling.
From the SCE8000(config if)#> prompt, enter **ip-tunnel DS-Lite** and press **Enter**.
- Step 3** Enable Extend-IANA to set 128 bit hash mode for DS-Lite environment.
From the SCE8000(config if)#> prompt, enter **ip-tunnel DS-Lite Extend-IANA** and press **Enter**.
The DS-Lite Extension IANA support is significant only if DS-Lite mode is enabled.
- Step 4** Restart the linecard.
From the SCE8000(config if)#> prompt, enter **no shutdown** and press **Enter**.
-

Configuring the L2TP Tunnels



Caution

IP tunneling must be enabled or disabled only when no application is loaded or the linecard is shut down.

Enabling L2TP Tunneling

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** Shut down the linecard. (This is a root level command.)
From the SCE8000(config if)#> prompt, enter **shutdown** and press **Enter**.
- Step 2** Enable L2TP tunneling.
From the SCE8000(config if)#> prompt, enter **ip-tunnel l2tp skip** and press **Enter**.
- Step 3** Restart the linecard.
From the SCE8000(config if)#> prompt, enter **no shutdown** and press **Enter**.
-

Disabling L2TP Tunneling

Disables all IP tunnels except IPinIP and GRE.

-
- Step 1** Shut down the linecard. (This is a root level command.)
From the SCE8000(config if)#> prompt, enter **shutdown** and press **Enter**.
- Step 2** Disable L2TP tunneling.
From the SCE8000(config if)#> prompt, enter **no ip-tunnel l2tp skip** and press **Enter**.
- Step 3** Restart the linecard.
From the SCE8000(config if)#> prompt, enter **no shutdown** and press **Enter**.
-

Configuring GRE Tunneling

- [Enabling GRE Tunneling, page 7-13](#)
- [Disabling GRE Tunneling, page 7-13](#)

GRE tunneling 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 GRE skip is disabled, the system treats the external IP header as the subscriber traffic, resulting in all GRE traffic being reported as generic IP.

Guidelines for configuring GRE tunnels:

- GRE and other tunnels: GRE tunnels are 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 GRE tunnels.
- DSCP marking: For GRE traffic, DSCP marking can be done on either the external or the internal IP header exclusively. (See [“Configuring DSCP Marking” section on page 7-14.](#))

**Caution**

IP tunneling can be configured (enabled, disabled or DSCP marking configuration) only when there is no application loaded or the linecard is shut down.

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 GRE Tunneling

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

-
- Step 1** Shut down the linecard. (This is a root level command.)
From the SCE8000(config if)#> prompt, enter **shutdown** and press **Enter**.
- Step 2** Enable GRE tunneling.
From the SCE8000(config if)#> prompt, enter **ip-tunnel gre skip** and press **Enter**.
- Step 3** Restart the linecard.
From the SCE8000(config if)#> prompt, enter **no shutdown** and press **Enter**.
-

Disabling GRE Tunneling

-
- Step 1** Shut down the linecard. (This is a root level command.)
From the SCE8000(config if)#> prompt, enter **shutdown** and press **Enter**.
- Step 2** Disable GRE tunneling.
From the SCE8000(config if)#> prompt, enter **no ip-tunnel gre skip** and press **Enter**.
- Step 3** Restart the linecard.
From the SCE8000(config if)#> prompt, enter **no shutdown** and press **Enter**.
-

Configuring IPinIP Tunneling

- [Enabling IPinIP Tunneling, page 7-14](#)
- [Disabling IPinIP Tunneling, page 7-14](#)

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 “[Configuring DSCP Marking](#)” section on page 7-14).

**Caution**

IP tunneling can be configured (enabled, disabled or DSCP marking configuration) only when there is no application loaded or the linecard is shut down.

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** Shut down the linecard. (This is a root level command.)
From the SCE8000(config if)#> prompt, type **shutdown** and press **Enter**.
- Step 2** Enable IPinIP tunneling.
From the SCE8000(config if)#> prompt, type **ip-tunnel IPinIP skip** and press **Enter**.
- Step 3** Restart the linecard.
From the SCE8000(config if)#> prompt, type **no shutdown** and press **Enter**.
-

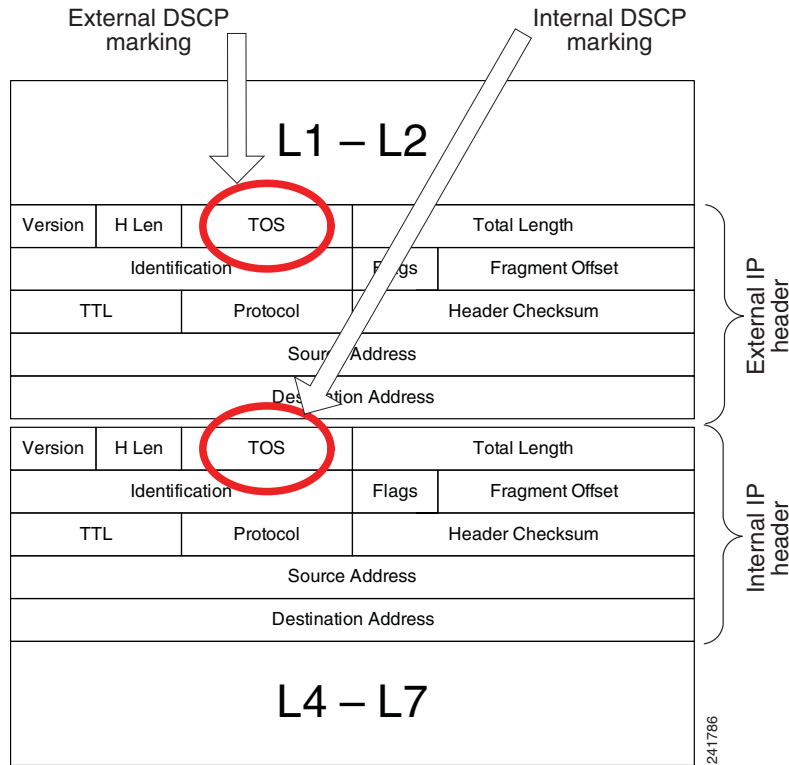
Disabling IPinIP Tunneling

-
- Step 1** Shut down the linecard. (This is a root level command.)
From the SCE8000(config if)#> prompt, type **shutdown** and press **Enter**.
- Step 2** Disable IPinIP tunneling.
From the SCE8000(config if)#> prompt, type **no ip-tunnel IPinIP skip** and press **Enter**.
- Step 3** Restart the linecard.
From the SCE8000(config if)#> prompt, type **no shutdown** and press **Enter**.
-

Configuring DSCP Marking

DSCP marking modifies the DSCP bits of the IPv4 header. In GRE and IPinIP tunnels there are at least two IP headers. By default, DSCP marking is performed only on the external IP header (refer to [Figure 7-1](#)). You can configure whether DSCP marking will be performed in the internal or external header.

Figure 7-1 DSCP Marking for IPinIP or GRE Tunnels



Note

DSCP marking should be enabled and configured through SCA BB console. See the *Cisco Service Control Application for Broadband User Guide* for further information.

Configuring DSCP Marking on the Internal IP Header

Use this command to configure the SCE platform to mark the DSCP bits of the internal IP header. This command takes effect only when the relevant tunneling mode (*GRE skip* or *IPinIP skip*) is enabled.

- Step 1** Shut down the linecard. (This is a root level command.)
From the SCE8000(config if)#> prompt, type **shutdown** and press **Enter**.
- Step 2** Configure the DSCP marking (.)
From the SCE8000(config if)#> prompt, type **ip-tunnel (GRE|IPinIP) DSCP-marking-skip** and press **Enter**.
Enables DSCP marking on the internal IP header of IPinIP traffic.
- Step 3** Restart the linecard.
From the SCE8000(config if)#> prompt, type **no shutdown** and press **Enter**.

Configuring DSCP Marking on the External IP Header

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

-
- Step 1** Shut down the linecard. (This is a root level command.)
From the SCE8000(config if)#> prompt, type **shutdown** and press **Enter**.
- Step 2** Configure the DSCP marking.
From the SCE8000(config if)#> prompt, type **no ip-tunnel (GRE|IPinIP) DSCP-marking-skip** and press **Enter**.
Enables DSCP marking on the external IP header of IPinIP traffic.
- Step 3** Restart the linecard.
From the SCE8000(config if)#> prompt, type **no shutdown** and press **Enter**.
-

Configuring the 6to4 Environment

Configure the 6to4 environment after you configure the IPv6 environment. To configure the 6to4 environment, complete the following procedure:

-
- Step 1** Enter the Global Configuration mode:
SCE8000# **configure**
- Step 2** Configure the 6to4-IPv6 mode.
From the Global Configuration prompt, enter **debug const-db name seCommonConstDb.box.isIPv6to4IPv6Mode value true** and press **Enter**. By default, the value is false.



Note If IPv6 mode is set to true, the IPv6 hash value should also be set to true.

- Step 3** (Only for IPv6 mode) Configure the 6to4-IPv6 hash.
From the Global Configuration prompt, enter **debug const-db name seCommonConstDb.box.isIPv6to4IPv6Hash value true** and press **Enter**. By default, the value is false.
- Step 4** Copy the running configuration to the startup configuration.
SCE8000#> **copy running-config startup-config**
- Step 5** Reboot the Cisco SCE 8000 device.
After Cisco SCE 8000 restarts, you can use the following **configuration** and **show** commands to configure the 6to4 and 6rd tunnels:
- **configure interface linecard 0 IP-tunnel 6to4**
 - **configure interface linecard 0 no IP-tunnel 6to4**
 - **show interface linecard 0 IP-tunnel6to4**
-

Configuring the VLAN Environment



Note

The SCE 8000 supports a maximum of 4096 VLAN tags.

Use this command to configure the VLAN environment.

- [Options, page 7-17](#)
- [Configuring the VLAN Environment: Example, page 7-17](#)

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. This is the only mode that supports private IP addresses. 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, affecting both performance and capacity.

From the SCE8000(config if)# prompt, type:

Command	Purpose
vlan {symmetric classify symmetric skip a-symmetric skip}	Configures the VLAN environment. Specify the desired VLAN mode.

Configuring the VLAN Environment: Example

The following example selects VLAN-based classification:

```
SCE8000(config if)#vlan symmetric classify
```

Configuring the MPLS Environment

Use this command to set the MPLS environment.

Options

The following options are available:

- **traffic-engineering skip** (default)—Use when all IP addresses are unique and MPLS labels are not mandatory for routing.
- **VPN skip**—Use when all IP addresses are unique, but MPLS labels are mandatory for routing.

Use the *VPN* keyword when the labels are mandatory in the traffic, otherwise use *traffic-engineering* (default).

Note that using the *VPN* value incurs a performance penalty.

From the SCE8000(config if)# prompt, type:

Command	Purpose
mpls {traffic-engineering skip vpn skip}	Sets the MPLS environment. Specify the desired MPLS mode.

Configuring the L2TP Environment

- [External Fragmentation in the L2TP Environment, page 7-18](#)
- [Options, page 7-18](#)

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](#)” section on [page 7-25](#)) 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

By enabling quick-forwarding, the Cisco SCE can only perform traffic monitoring for externally-fragmented L2TP traffic. It cannot perform flow redirection, flow blocking, or rate-limiting.

Options

The following option is available:

- **portnumber**—The port number that the LNS and LAC use for L2TP tunnels.
The default is 1701.

From the SCE8000(config if)# prompt, type:

Command	Purpose
L2TP identify-by port-number <i>portnumber</i>	Configures the L2TP environment.

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 pair of MAC addresses.
- 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). Note as well, that in order to support asymmetric layer 2, the SCE platform switches to *asymmetric flow open* mode, which incurs a certain performance penalty, as well as reducing capacity. This is NOT the case for asymmetric routing topology.

From the SCE8000(config if)# prompt, type:

Command	Purpose
asymmetric-L2-support	Enables asymmetric L2 support.

Displaying the Tunneling Configuration

From the SCE8000# prompt, type:

Command	Purpose
show interface linecard 0 (MPLS VLAN L2TP IP-tunnel)	Displays the current configuration for the specified tunnel option.

How to Display the 6to4 Configuration

From the SCE8000# prompt, type:

Command	Purpose
show interface linecard 0 IP-tunnel 6to4	Displays the current configuration for the 6to4 tunnel option.

How to Display the DS-Lite Configuration

From the SCE8000# prompt, type:

Command	Purpose
<code>show interface linecard 0 IP-tunnel DS-Lite</code>	Displays the current configuration for the DS-Lite tunnel option.

How to Display the IPinIP Configuration

From the SCE8000# prompt, type:

Command	Purpose
<code>show interface linecard 0 ip-tunnel IPinIP</code>	Displays the current configuration for the specified tunnel option.

How to Display the Logged-In VPNs

Options

The following options are available

- **vpn-name**—The name of a specific currently logged-in VPN for which to display details.
- **all-names**—Use this keyword to display all the VPN names that are currently logged into the system.

From the SCE8000> prompt, type:

Command	Purpose
<code>show interface linecard 0 VPN {name <i>vpn-name</i> all-names}</code>	Displays the logged-in VPNs.

How to Display the Asymmetric L2 Support Mode

From the SCE8000# prompt, type:

Command	Purpose
<code>show interface linecard 0 asymmetric-L2-support</code>	Displays asymmetric L2 support mode.

Managed VPNs

- [Private IP Addresses, page 7-21](#)
- [Capacity, page 7-21](#)
- [Limitations for VPN mode, page 7-21](#)

A managed VPN is a named entity, introduced similarly to the same way that a subscriber is introduced, and containing VPN mappings.

A managed VPN contains a single VLAN mapping. A VPN-based subscriber contains a set of mappings of the form: IP@VpnName, where IP can be either a single IP address or a range of addresses.

Managed VPN entities can be configured only via the SM. The SCE platform CLI can be used to view VPN-related information, but not to configure the VPNs.

Private IP Addresses

Private IP addresses are supported only in the following mode, as this mode provides information regarding the higher-level entity (VLAN or VPN) to which the IP addresses of the flow belong:

- VLAN symmetric classify

Capacity

The system supports:

- 2048 VPNs
- 80,000 IP mappings over VPNs

Limitations for VPN mode

Mutually exclusive system modes

When the system is working in VPN mode, the following modes are not supported:

- DDoS
- Value Added Services (VAS) mode

Subscriber-related limitations

- The SM must be configured to operate in Push mode.
- Introduced subscriber aging is not supported when using VPN-based subscribers

TCP-related requirements

- Number of Upstream TCP Flows – There must be enough TCP flows opening from the subscriber side on each PE-PE route in each period of time. The higher the rate of TCP flows from the subscriber side, the higher the accuracy of the mechanism can be.

VPN configuration requirements

- In VLAN-based VPNs (VLAN symmetric classify mode), a subscriber may have IP mappings over more than one VPN, but only if the IP mappings are the full range of the VPN (0.0.0.0/0). (This option is provided for backwards compatibility, supporting legacy multi-VLAN subscribers.)

Monitoring VPN Support

The SCE platform CLI allows you to do the following:

- Display VPN-related mappings
- Monitor subscriber counters

Displaying VPN-Related Mappings

Use the following Viewer commands to display subscriber mappings. These commands display the following information:

- All the mappings for a specified VPN
- A listing of all currently logged-in VPNs
- A listing of all subscribers mapped to an IP range on a specified VPN
- The number of subscribers mapped to an IP range on a specified VPN

How to Display Mappings for a Specified VPN

- [Options, page 7-22](#)
- [Displaying Mappings for a Specified VPN: Examples, page 7-22](#)

Options

The following option is available:

- **vpn-name**—The name of the VPN for which to display mappings.

From the SCE8000> prompt, type:

Command	Purpose
show interface linecard 0 VPN name <i>vpn-name</i>	Displays mappings for a specified VPN.

Displaying Mappings for a Specified VPN: Examples

The following example illustrates the output of this command for a VLAN-based VPN:

```
SCE8000> show interface linecard 0 VPN name vpn3
VPN name: Vpn3
VLAN: 2
Number of subscriber mappings: 0
Explicitly introduced VPN
```

The following example illustrates the output of this command for an automatically created VLAN VPN:

```
SCE8000> show interface linecard 0 VPN name 2
VPN name: 2
VLAN: 2
Number of subscriber mappings: 1
Automatically created VPN
```

How to Display a Listing of all VPNs

Use this command to display a listing of all currently logged-in VPNs.

From the SCE8000> prompt, type:

Command	Purpose
show interface linecard 0 VPN all-names	Displays a listing of all currently logged-in VPNs.

Displaying a Listing of All VPNs: Example

```
SCE8000> show interface linecard 0 VPN all-names
```

How to Display Subscriber Mappings for an IP Range on a Specified VPN

- [Options, page 7-23](#)
- [Displaying Subscribers Mapped to a IP range on a Specified VPN: Example, page 7-23](#)

Options

The following options are available:

- **ip-range**—The IP range for which to display mapped subscribers
- **vpn-name**—The name of the VPN for which to display mappings.

From the SCE8000> prompt, type:

Command	Purpose
show interface linecard 0 subscriber mapping included-in IP <i>ip-range</i> VPN <i>vpn-name</i>	Displays subscriber mappings for an IP range on a specified VPN. The VPN option allows you to search for subscribers with a private IP mapping.

Displaying Subscribers Mapped to a IP range on a Specified VPN: Example

```
SCE8000> show interface linecard 0 subscriber mapping included-in IP 10.0.0.0/0 VPN vpn1
Subscribers with IP mappings included in IP range '10.0.0.0/0@vpn1':
Subscriber 'Sub10', mapping '10.1.4.150/32@vpn1'.
Subscriber 'Sub10', mapping '10.1.4.149/32@vpn1'.
Subscriber 'Sub10', mapping '10.1.4.145/32@vpn1'.
Subscriber 'Sub11', mapping '10.1.4.146/32@vpn1'.
Total 2 subscribers found, with 4 matching mappings
```

How to Display the Number of Subscribers Mapped to an IP range on a Specified VPN

- [Options, page 7-23](#)
- [Displaying the Number of Subscribers Mapped to Range on a Specified VPN: Example, page 7-24](#)

Options

The following options are available:

- **ip-range**—The IP range for which to display mapped subscribers
- **vpn-name**—The name of the VPN for which to display mappings.

Use the **amount** keyword to display the number of subscribers rather than a listing of subscriber names.

From the SCE8000> prompt, type:

Command	Purpose
show interface linecard 0 subscriber amount mapping included-in IP <i>ip-range</i> VPN <i>vpn-name</i>	Displays the number of subscribers mapped to an IP range on a specified VPN.

Displaying the Number of Subscribers Mapped to Range on a Specified VPN: Example

```
SCE8000> show interface linecard 0 subscriber amount mapping included-in IP 0.0.0.0/0 VPN
vpn1
There are 2 subscribers with 4 IP mappings included in IP range '0.0.0.0/0'.
```


Configuring Traffic Rules and Counters

- [Traffic Rules and Counters, page 7-25](#)
- [Configuring Traffic Counters, page 7-27](#)
- [Configuring Traffic Rules, page 7-28](#)
- [Managing Traffic Rules and Counters, page 7-33](#)

Traffic Rules and Counters

- [What are Traffic Rules and Counters?, page 7-25](#)
- [Traffic Rules, page 7-26](#)
- [Traffic Counters, page 7-26](#)

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 SCE 8000 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 a combination of the following:

- **IP address**—A single address or a subnet range can be specified for each of the line ports (Subscriber / Network) for IPv4 IP addresses. A specific IP address or a CIDR notation prefix for IPv6 IP addresses.
- **Protocol**—For IPv4, TCP/UDP/ICMP/IGRP/EIGRP/IS-IS/OSPF/Other. For IPv6, TCP and UDP.
- **TCP/UDP Ports**—A single port or a port range can be specified for each of the line ports (Subscriber / Network) for IPv4 IP address. A single port only can be specified for each of the line ports (Subscriber / Network) for IPv4 IP address. Valid only for TCP/UDP.
- **Direction (Upstream/Downstream)**—Valid only for TCP.

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 and so on are performed.)
- **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

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

Note that **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 7-27](#)
- [How to Delete a Traffic Counter, page 7-27](#)
- [How to Delete all Existing Traffic Counters, page 7-27](#)

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.

From the SCE8000(config if)# prompt, type:

Command	Purpose
traffic-counter name <i>name</i> count-bytes count-packets	Adds a traffic counter with the specified name and counting mode.

How to Delete a Traffic Counter

From the SCE8000(config if)# prompt, type:

Command	Purpose
no traffic-counter name <i>name</i>	Deletes a traffic counter. Note that a traffic counter cannot be deleted if it is used by any existing traffic rule.

How to Delete all Existing Traffic Counters

From the SCE8000(config if)# prompt, type:

Command	Purpose
no traffic-counter all	Removes all traffic counters. Note that 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 for IPv4 Addresses, page 7-28](#)
- [How to Create a Traffic Rule for IPv6 Addresses, page 7-31](#)
- [How to Delete a Traffic Rule, page 7-32](#)
- [How to Delete All Traffic Rules, page 7-32](#)
- [How to Delete All Flow Control Traffic Rules, page 7-33](#)

How to Create a Traffic Rule for IPv4 Addresses

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 [“Configuring the VLAN Environment” section on page 7-17](#), and only when *tunnel id* mode is enabled.

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

Note that 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
- **flow-capture**—Capture the flow configured by this rule. No service to this flow

From the SCE8000(config if)# prompt, type:

Command	Purpose
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]	Creates a traffic rule.

Configuring Traffic Rules: Examples

- [Example 1, page 7-29](#)
- [Example 2, page 7-30](#)
- [Example 3, page 7-30](#)
- [Example 4, page 7-30](#)

Example 1

This example creates the following traffic rule:

```
SCE8000(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
```

- 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

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.

```
SCE8000(config if)#traffic-rule tunnel-id-mode
SCE8000(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**.

```
SCE8000(config if)# traffic-rule name rule3 IP-addresses all protocol IS-IS direction
upstream traffic-counter none action ignore
```

Example 4

The following example illustrates how to configure a traffic rule that will be used as a recording rule using the flow-capture option. All flows that match this rule will be recorded when the flow capture process is in operation.

1. Name—FlowCaptureRule
2. IP addresses: subscriber side—all IP addresses, network side—all IP addresses
3. Direction—both
4. Protocol—250
5. Traffic counter name—counter2

6. Action—flow-capture

7. The actions performed will be counting and flow capture.

```
SCE8000>enable 10
Password:<cisco>
SCE8000#configure
SCE8000(config)#interface linecard 0
SCE8000(config if)#traffic-rule name FlowCaptureRule ip-addresses subscriber-side all
network-side all protocol 250 direction both traffic-counter name counter2 action
flow-capture
SCE8000(config if)#
```

How to Create a Traffic Rule for IPv6 Addresses

Options

The following options are available:

IP specification:

```
all | (ip-address | ip-prefix)
```

- *ip-address* is a single IP address in CIDR notation, such as A:B:C:D:E:F:G:H/I
- *ip-prefix* is an IP subnet range, in the CIDR notation, such as 2001:DB8:0:1:FFFF:1234::5.

protocol:

Any one of the following protocols:

```
TCP/UDP/all
```

port specification:

```
port
```

- Specify the ports only if the protocol is either TCP or UDP.
- Specify the port for both the subscriber side and the network side.
- Create multiple rules if you plan to use multiple ports.

direction:

One of the following:

```
upstream/downstream/both
```

traffic-counter:

Either of the following:

- **name** <name of an existing traffic counter>—Packets that meet the traffic rule criteria are counted in the specified counter. If a counter name is defined, the count action is also defined implicitly. The **name** keyword must also be displayed with the actual name of the counter.
- **none**—If **none** is specified, an action must be explicitly defined using the action option.

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

One of the following:

- *block*—Block the specified traffic.
- *classical-open-flow-mode*—Use the classical open flow mode for the specified flow.
- *ignore*—Bypass the specified traffic; traffic receives no service.

From the SCE8000(config if)# prompt, type:

Command	Purpose
traffic-rule name <i>name</i> ipv6 IP-addresses (all((subscriber-side <IP specification> network-side <IP specification>)) protocol <i>protocol</i> [ports subscriber-side <port specification> network-side <port specification>] [tunnel-id <tunnel-id specification>] direction <i>direction</i> traffic-counter <traffic-counter>[action <i>action</i>]	Creates a traffic rule.

Example for Configuring a Traffic Rule

This example creates a traffic rule called rule2:

- Name—rule2
- IP addresses—subscriber side:all, network side: all
- Protocol—TCP
- Tunnel ID all
- Direction—downstream
- Traffic counter—counter2
- Action—block

The actions that are performed are counting and blocking.

```
SCE8000> enable 10
Password:<cisco>
SCE # config
SCE8000(config)# interface linecard 0
SCE8000(config if)# traffic-rule name rule2 ipv6 ip-addresses subscriber-side all
network-side all protocol tcp tunnel-id all direction downstream traffic-counter name
counter2 action block
SCE8000(config if)
```

How to Delete a Traffic Rule

From the SCE8000(config if)# prompt, type:

Command	Purpose
no traffic-rule name <i>name</i>	Removes the specified traffic rule.

How to Delete All Traffic Rules

From the SCE8000(config if)# prompt, type:

Command	Purpose
no traffic-rule all	Removes all existing traffic rules.

How to Delete All Flow Control Traffic Rules

From the SCE8000(config if)# prompt, type:

Command	Purpose
no traffic-rule capture	Removes all flow capture 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 7-33](#)
- [How to View All Traffic Rules, page 7-33](#)
- [How to View a Specified Traffic Counter, page 7-33](#)
- [How to View all Traffic Counters, page 7-34](#)
- [How to Reset a Specified Traffic Counter, page 7-34](#)
- [How to Reset All Traffic Counters, page 7-34](#)

How to View a Specified Traffic Rule

From the SCE8000# prompt, type:

Command	Purpose
show interface linecard 0 traffic-rule name <i>rule-name</i>	Displays the configuration of the specified traffic rule.

How to View All Traffic Rules

From the SCE8000# prompt, type:

Command	Purpose
show interface linecard 0 traffic-rule all	Displays the configuration of all existing traffic rules.

How to View a Specified Traffic Counter

From the SCE8000# prompt, type:

Command	Purpose
show interface linecard 0 traffic-counter name <i>counter-name</i>	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”.

```
SCE8000# show interface linecard 0 traffic-counter name cnt
Counter 'cnt' value: 0 packets. Rules using it: None.
```

How to View all Traffic Counters

From the SCE8000# prompt, type:

Command	Purpose
show interface linecard 0 traffic-counter all	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.

```
SCE8000# 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

From the SCE8000# prompt, type:

Command	Purpose
clear interface linecard 0 traffic-counter name <i>counter-name</i>	Resets the specified traffic counter.

How to Reset All Traffic Counters

From the SCE8000# prompt, type:

Command	Purpose
clear interface linecard 0 traffic-counter all	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.


Note

DSCP marking is supported only for IPv4 addresses.

DSCP marking configuration is performed via the Cisco SCA BB console. The Cisco 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, see the [Cisco Service Control Application for Broadband User Guide](#).


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.

From the SCE8000> prompt, type:

Command	Purpose
<code>show interface linecard 0 ToS-marking</code>	Displays the state of DSCP marking (enabled or disabled) per interface and the DSCP translation table.

Counting the Dropped Packets

- [About Counting the Dropped Packets, page 7-36](#)
- [Disabling the Hardware Packet Drop, page 7-36](#)

About Counting the 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 that 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 the following 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.

From the SCE8000(config if)# prompt, type:

Command	Purpose
no accelerate-packet-drops	Disables hardware packet drop.
accelerate-packet-drops	Enables hardware packet drop.