

Cisco SCE 8000 10GBE Software Configuration Guide

Release 5.2.x
September 23, 2015

Cisco Systems, Inc.
www.cisco.com

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco website at www.cisco.com/go/offices.

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

Any Internet Protocol (IP) addresses used in this document are not intended to be actual addresses. Any examples, command display output, and figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses in illustrative content is unintentional and coincidental.

Cisco SCE 8000 10GBE Software Configuration Guide

© 2015 - 2016 Cisco Systems, Inc. All rights reserved.



CONTENTS

About this Guide 1

Introduction 1

Document Revision History 2

Organization 3

Related Publications 5

Conventions 6

Obtaining Documentation and Submitting a Service Request 7

CHAPTER 1 **Cisco Service Control Overview** 1-1

Introduction 1-1

Cisco Service Control Solution 1-2

Service Control for Broadband Service Providers 1-2

Cisco Service Control Capabilities 1-3

Cisco SCE Platform Description 1-4

Bandwidth Management of P2P Traffic 1-5

Management and Collection 1-6

Network Management 1-6

Subscriber Management 1-7

Service Configuration Management 1-7

Data Collection 1-7

IPv6 Support 1-8

CHAPTER 2 **Command-Line Interface** 2-1

Introduction 2-1

Authorization and Command Mode Levels (Hierarchy) 2-2

CLI Authorization Levels 2-2

CLI Command Mode Hierarchy 2-3

Prompt Indications 2-6

Navigating Between Authorization Levels and Command Modes 2-6

The do Command: Executing Commands Without Exiting 2-7

CLI Help Features 2-9

Partial Help 2-9

Argument Help 2-9

Navigational and Shortcut Features 2-11
Command History 2-11
Keyboard Shortcuts 2-11
Auto-Completion 2-12
FTP User Name and Password 2-13
Managing Command Output 2-14
Scrolling the Screen Display 2-14
Filtering Command Output 2-14
Redirecting Command Output to a File 2-15
Creating a CLI Script 2-16

CHAPTER 3 **Basic Cisco SCE 8000 Platform Operations** 3-1
Introduction 3-1
Starting the Cisco SCE 8000 Platform 3-2
Checking Conditions Prior to System Startup 3-2
Starting the System and Observing Initial Conditions 3-2
Final Tests 3-3
How to Verify Operational Status 3-3
How to View the User Log Counters 3-3
Managing Configurations 3-5
Viewing Configurations 3-5
How to Save or Change the Configuration Settings 3-6
Example for Saving or Changing the Configuration Settings 3-7
Restoring a Previous Configuration 3-8
Example for Restoring a Previous Configuration 3-8
How to Display the Cisco SCE Platform Version Information 3-10
Example for Displaying the Cisco SCE Platform Version Information 3-10
How to Display the Cisco SCE Platform Inventory 3-13
Examples for Displaying the Cisco SCE Platform Inventory 3-13
Displaying the Cisco SCE Platform Inventory: FRUs Only 3-13
Displaying the Complete Cisco SCE Platform Inventory 3-14
How to Display the System Uptime 3-17
Example for Displaying the System Uptime 3-17
Configuring the System Mode 3-17
Configuring the IPv6 Prefix Length 3-18
Monitoring Control Processor CPU Utilization 3-20
CLI Commands for Monitoring Control Processor CPU Utilization 3-20
Example for Monitoring Control Processor CPU Utilization 3-21

Rebooting and Shutting Down the Cisco SCE Platform 3-23
Rebooting the Cisco SCE Platform 3-23
Examples for Rebooting the Cisco SCE Platform 3-23
How to Shut Down the Cisco SCE Platform 3-23
Example for Shutting Down the Cisco SCE Platform 3-24

CHAPTER 4 **Utilities** 4-1
Introduction 4-1
Working with Cisco SCE Platform Files 4-2
Working with Directories 4-2
How to Create a Directory 4-2
How to Delete a Directory 4-2
How to Change Directories 4-3
How to Display your Working Directory 4-3
How to List the Files in a Directory 4-3
Working with Files 4-4
How to Rename a File 4-4
How to Delete a File 4-4
Copying Files 4-5
How to Display File Contents 4-6
How to Unzip a File 4-6
The User Log 4-7
The Logging System 4-7
Copying the User Log 4-7
Enabling and Disabling the User Log 4-8
Viewing the User Log Counters 4-8
Viewing the User Log 4-9
Clearing the User Log 4-9
Generating a File for Technical Support 4-9
Generating a File for Technical Support: Example 4-9
Managing Syslog 4-10
Enabling and Disabling Syslog 4-10
Configuring Remote Syslog Hosts 4-10
How to Add a Remote Syslog Host 4-11
How to Remove a Remote Syslog Host 4-11
Configuring the Minimum Severity Level to be Logged to Syslog 4-11
How to Configure the Minimum Severity Level for Syslog Messages 4-12
How to Restore the Default Minimum Severity Level for Syslog Messages 4-12
Configuring the Syslog Facility 4-12

How to Configure the Syslog Facility 4-13
How to Restore the Default Syslog Facility 4-13
Configuring the Syslog Logging Rate Limit 4-13
How to Configure the Syslog Rate Limit 4-14
How to Restore the Default Syslog Rate Limit 4-14
Configuring the Syslog Time Stamp Format 4-14
How to Configure the Syslog Time Stamp Format 4-15
How to Restore the Default Syslog Time Stamp Format 4-15
Enabling and Disabling the Syslog Message Counter 4-15
Monitoring Syslog 4-15
How to Display the Syslog Configuration 4-16
How to Display the Syslog Counters 4-16
Flow Capture 4-17
Limitations 4-17
The Flow Capture Process 4-18
Configuring a Flow Capture Traffic Rule 4-18
Configuring the Flow Capture Settings 4-18
Performing the Flow Capture 4-20
Monitoring the Flow Capture 4-20

CHAPTER 5 **Configuring the Management Interface and Security** 5-1
Introduction 5-1
Management Interface and Security 5-2
Configuring the Management Ports 5-3
Entering the Management Interface Configuration Mode 5-3
Configuring the Management Port Physical Parameters 5-4
Setting the IP Address and Subnet Mask of the Management Interface 5-4
Configuring the Management Interface Speed and Duplex Parameters 5-5
Specifying the Active Management Port 5-7
Management Interface Redundancy 5-8
Configuring the Management Ports for Redundancy 5-8
Configuring the Fail-Over Mode 5-9
Monitoring the Management Interface 5-10
Configuring Management Interface VLANs 5-11
Monitoring Management VLANs 5-14
TACACS+ Authentication, Authorization, and Accounting 5-15
Information About TACACS+ Authentication, Authorization, and Accounting 5-15
Login Authentication 5-16
Accounting 5-16

Privilege-Level Authorization 5-16
Command-Level Authorization 5-17
General AAA Fallback and Recovery Mechanism 5-17
About Configuring TACACS+ 5-18
Configuring the Cisco SCE Platform TACACS+ Client 5-19
Adding a New TACACS+ Server Host 5-19
Removing a TACACS+ Server Host 5-20
Configuring the Global Default Key 5-20
Configuring the Global Default Timeout 5-21
Managing the User Database 5-22
Adding a New User to the Local Database 5-22
Defining the User Privilege Level 5-24
Adding a New User with Privilege Level and Password 5-24
Deleting a User 5-26
Configuring AAA Login Authentication 5-26
Configuring Maximum Login Attempts 5-26
Configuring the AAA Login Authentication Methods 5-27
Configuring AAA Privilege-Level Authorization Methods 5-28
Configuring AAA Command-Level Authorization Methods 5-28
Configuring AAA Accounting 5-29
Monitoring TACACS+ 5-30
Displaying Statistics for TACACS+ Servers 5-30
Displaying Statistics, Keys, and Timeouts for TACACS+ Servers 5-30
Monitoring TACACS+ Users 5-31
Configuring Access Control Lists (ACLs) 5-32
Adding Entries to an ACL 5-33
Removing an ACL 5-33
Defining a Global ACL 5-34
Managing the Telnet Interface 5-35
Preventing Telnet Access 5-35
Assigning an ACL to the Telnet Interface 5-35
Removing ACL Assignment from the Telnet Interface 5-36
Configuring Telnet Timeout 5-36
Configuring the SSH Server 5-37
The SSH Server 5-37
Key Management 5-37
Managing the SSH Server 5-38
Generating a Set of SSH Keys 5-38
Enabling the SSH Server 5-38

Disabling the SSH Server 5-38
Running Only SSHv2 5-39
Assigning an ACL to the SSH Server 5-39
Removing the ACL Assignment from the SSH Server 5-39
Deleting the Existing SSH Keys 5-39
Monitoring the Status of the SSH Server 5-40
Configuring and Managing the SNMP Interface 5-41
About the SNMP Interface 5-41
SNMP Protocol 5-41
Security Considerations 5-42
About CLI 5-43
About MIBs 5-44
Configuration via SNMP 5-44
Enabling the SNMP Interface 5-45
How to Enable the SNMP Interface 5-45
How to Disable the SNMP Interface 5-45
Configuring SNMP Community Strings 5-45
Defining a Community String 5-46
Removing a Community String 5-46
Displaying the Configured Community Strings 5-47
Configuring SNMP Notifications 5-47
Configuring SNMP Server Group 5-48
Configuring SNMP Server View 5-48
Configuring SNMP Server User 5-49
Defining SNMP Hosts 5-50
Configuring SNMP Traps 5-51
SNMP Walk Acceleration for linkServiceUsage Queries 5-53
How to Enable SNMP Query Acceleration 5-53
SNMPv3 Configuration Example 5-53

CHAPTER 6 **Global Configuration** 6-1
Introduction 6-1
IP Routing Configuration 6-2
Configuring the IP Routing Table 6-2
How to Configure the Default Gateway 6-2
How to Add an Entry to the IP Routing Table 6-3
How to Display the IP Routing Table 6-3
IP Advertising 6-4
Configuring IP Advertising 6-4

How to Display the Current IP Advertising Configuration 6-5
Configuring Time Clocks and Time Zone 6-6
Displaying the System Time 6-6
Displaying the System Time: Example 6-6
Displaying the Calendar Time 6-7
Displaying the Calendar Time: Example 6-7
Setting the System Clock 6-7
Options 6-7
Setting the System Clock: Example 6-7
Setting the Calendar 6-7
Options 6-8
Setting the Calendar: Example 6-8
Setting the Time Zone 6-8
Options 6-8
Setting the Time Zone: Example 6-9
Removing the Current Time Zone Setting 6-9
Configuring Daylight Saving Time 6-9
Options 6-9
Guidelines 6-10
How to Define Recurring Daylight Saving Time Transitions 6-11
How to Define Non-Recurring Daylight Saving Time Transitions 6-11
How to Cancel the Daylight Saving Time Configuration 6-11
How to Display the Current Daylight Saving Time Configuration 6-12
Configuring SNTP 6-13
How to Enable the SNTP Multicast Client 6-13
How to Disable the SNTP Multicast Client 6-14
How to Enable the SNTP Unicast Client 6-14
Options 6-14
Enabling SNTP Unicast Client: Example 6-14
Disabling the SNTP Unicast Client 6-14
How to Disable the SNTP Unicast Client and Remove All Servers 6-14
How to Remove One SNTP Server 6-15
How to Define the SNTP Unicast Update Interval 6-15
Options 6-15
How to Display SNTP Information 6-15
Domain Name Server (DNS) Settings 6-17
Configuring DNS Lookup 6-17
How to Enable DNS Lookup 6-17
How to Disable DNS Lookup 6-17

Configuring Name Servers 6-18
Options 6-18
How to Define Domain Name Servers 6-18
How to Remove a Domain Name Server 6-18
How to Remove All Domain Name Servers 6-19
How to Add a Host to the Host Table 6-19
Options 6-19
Adding Hosts to Removing them from the Host Table: Example 6-19
How to Display Current DNS Settings 6-19
Displaying Current DNS Settings: Example 6-19
Configuring Cisco Discovery Protocol 6-20
Cisco Discovery Protocol 6-20
Cisco Discovery Protocol on the Cisco SCE 8000 Platform 6-21
CDP Operational Modes on the Cisco SCE 8000 6-21
CDP Limitations on the Cisco SCE 8000 6-22
Configuring CDP on the Cisco SCE 8000 Platform 6-22
Enabling CDP Globally 6-22
Setting CDP Mode 6-23
Enabling CDP on a Specific Traffic Interface 6-23
Setting the Hold Time 6-24
Setting the Timer 6-24
Monitoring and Maintaining CDP 6-25
CDP Configuration Examples 6-27
Example: Setting the CDP Mode 6-27
Example: Monitoring and Maintaining CDP 6-27
Enabling the CLI Interface Warning Banner 6-29
OS Fingerprinting and NAT Detection 6-30
Restrictions and Limitations 6-30
Configuring OS Fingerprinting 6-31
Monitoring OS Fingerprinting 6-33
Using the Bursty Traffic Convergence 6-33
Restrictions and Known Limitations in Bursty Traffic Convergence 6-35
DNS Assisted Classification/Sampling 6-36
Configuring DNS Assisted Classification 6-36

CHAPTER 7 **Configuring Line Interfaces** 7-1
Introduction 7-1
Line Interfaces 7-2
Information About Line Interfaces 7-2

Flow Control and Bandwidth Considerations 7-2
Maximum Packet Size 7-2
How to Configure the Ten Gigabit Ethernet Line Interfaces 7-2
Changing the Traffic Direction on the Ten Gigabit Ethernet Line Interfaces 7-3
How to Specify the Traffic Direction on Link 1 7-3
Tunneling Protocols 7-4
Tunneling IPv6 Traffic 7-6
Selecting the Tunneling Mode 7-7
Configuring the 6to4 Tunnels 7-8
Configuring DS-Lite Tunnels 7-9
Configuring L2TP Tunnels 7-10
Configuring GRE Tunneling 7-11
Configuring IPinIP Tunneling 7-13
Configuring DSCP Marking 7-14
Configuring the 6to4 Environment 7-15
Configuring the VLAN Environment 7-16
Configuring the MPLS Environment 7-17
Configuring the L2TP Environment 7-18
Asymmetric L2 Support 7-18
Displaying the Tunneling Configuration 7-19
How to Display the 6to4 Configuration 7-19
How to Display the DS-Lite Configuration 7-19
How to Display the IPinIP Configuration 7-20
How to Display the Logged-In VPNs 7-20
Options 7-20
How to Display the Asymmetric L2 Support Mode 7-21
Managed VPNs 7-22
Monitoring VPN Support 7-23
Displaying VPN-related Mappings 7-23
Configuring Traffic Rules and Counters 7-26
Traffic Rules and Counters 7-26
What are Traffic Rules and Counters? 7-26
Traffic Rules 7-27
Traffic Counters 7-27
Configuring Traffic Counters 7-28
How to Create a Traffic Counter 7-28
How to Delete a Traffic Counter 7-28
How to Delete all Existing Traffic Counters 7-28
Configuring Traffic Rules 7-29

How to Create a Traffic Rule for IPv4 Addresses 7-29
How to Create a Traffic Rule for IPv6 Addresses 7-32
How to Delete a Traffic Rule 7-33
How to Delete All Traffic Rules 7-33
How to Delete All Flow Control Traffic Rules 7-34
Managing Traffic Rules and Counters 7-34
How to View a Specified Traffic Rule 7-34
How to View All Traffic Rules 7-34
How to View a Specified Traffic Counter 7-34
How to View All Traffic Counters 7-35
How to Reset a Specified Traffic Counter 7-35
How to Reset All Traffic Counters 7-35
DSCP Marking 7-36
How to Display the DSCP Marking Configuration 7-36
Counting Dropped Packets 7-37
About Counting Dropped Packets 7-37
Disabling the Hardware Packet Drop 7-37

CHAPTER 8 **Configuring the Connection** 8-1
Introduction 8-1
Configuring the Connection Mode 8-2
Options 8-2
Configuring the Connection Mode Examples 8-3
Monitoring the Connection Mode and Related Parameters 8-4
Connection Mode Examples 8-4
Configuring the Link Mode 8-6
About the Link Mode 8-6
Options 8-6
External Optical Bypass 8-7
How to Activate the External Bypass 8-7
How to Deactivate the External Bypass 8-8
How to Set the External Bypass to the Default State 8-8
How to Display the State of the External Bypass 8-8
Hardware Bypass 8-9
How to Enable the Hardware Bypass Mode 8-9
How to Disable the Hardware Bypass Mode 8-10
How to Display the Status of the Hardware Bypass Mode 8-10
How to Set the Hardware Bypass Status for a Static Party 8-10
How to Reset the Hardware Bypass Staus of a static party 8-10

How to Display the Hardware Bypass Status of a Static Party 8-11
How to Display the Startup Configuration Party Database 8-11
How to Display the Currently Running Party Database Configuration 8-12
How to Copy the Running Configuration Party Database to the Startup Configuration Party Database 8-12
How to Copy the Startup Configuration Party Database and Create a Backup File 8-12
Configuring the Static Subscribers 8-13
How to the Set the IP Address for a Static Subscriber 8-13
How to Set the IP Range for a Static Subscriber 8-13
How to Display All Mappings to Dual Stack Static Subscriber 8-14
How to Display IPv6 Mappings to Dual Stack Static Subscriber 8-14
How to Display Dual Stack Static Subscriber 8-14
Link Failure Reflection 8-15
How to Enable Link Failure Reflection 8-15
How to Disable Link Failure Reflection 8-15
Enabling and Disabling Link Failure Reflection on All Ports 8-15
Options 8-16
How to Enable Link Failure Reflection on All Ports 8-16
How to Disable Link Failure Reflection on All Ports 8-16
Configuring Link Failure Reflection in Linecard-Aware Mode 8-16
How to Enable Linecard-Aware Mode 8-17
How to Disable Linecard-Aware Mode 8-17
Asymmetric Routing Topology 8-18
Asymmetric Routing and Other Service Control Capabilities 8-18
Enabling Asymmetric Routing 8-18
Monitoring Asymmetric Routing 8-19
Monitoring Asymmetric Routing: Example 8-19
Configuring a Forced Failure 8-20
Configuring the Failure Recovery Mode 8-21
Options 8-21
Configure the Failure Recovery Mode: Examples 8-21
Configuring the Cisco SCE Platform/SM Connection 8-22
Configuring the Behavior of the Cisco SCE Platform in Case of Failure of the SM 8-22
Options 8-22
Configuring the SM-SCE Platform Connection Timeout 8-23
Options 8-23

CHAPTER 9 **Raw Data Formatting: The RDR Formatter and NetFlow Exporting** 9-1

Introduction 9-1

RDR Formatter and NetFlow Exporting Support 9-2
The RDR Formatter 9-2
NetFlow 9-2
NetFlow Terminology 9-2
NetFlow Exporting Support 9-3
Data Destinations 9-4
Categories 9-5
Priority 9-5
Setting DSCP for NetFlow 9-6
Forwarding Modes 9-6
Protocol 9-6
Transport Type 9-6
Configuring Data Destinations and Categories 9-7
Configuring a Data Destination 9-7
Options 9-7
Configuring the Data Destinations: Examples 9-7
Configuring the Data Categories 9-8
Configuring the Buffer Size 9-9
Configuring a Destination and Assigning Categories 9-9
Configuring the Forwarding Mode 9-13
Options 9-13
Configuring the Forwarding Mode: Example 9-13
Configuring the RDR Formatter 9-14
Options 9-14
How to Configure the Size of the RDR Formatter History Buffer 9-14
Configuring NetFlow Exporting Support 9-15
Options 9-15
How to Configure a DSCP Value for NetFlow 9-15
Options 9-15
How to Configure the Template Refresh Interval 9-15
Options 9-15
Configuring the Flexible Redirection Across SCOS Login 9-16
Configuring Redirection Rate Limit 9-17
Configuring Total Link Limit 9-17
Configuring Dynamic Mapping of RDRs to Categories 9-20
Configuring Mappings 9-20
Options 9-21
How to Restore the Default Mapping for a Specified RDR Tag 9-21
Displaying Data Destination Configuration and Statistics 9-22

How to the Display the Current RDR Formatter Configuration 9-22
Displaying the RDR Formatter Configuration: Example 9-23
How to the Display the Current RDR Formatter Statistics 9-23
Displaying the Current RDR Formatter Statistics: Example 9-23
Disabling the Linecard from Sending RDRs 9-25
Disabling RDR Aggregation 9-26

CHAPTER 10 **Managing Subscribers** 10-1
Introduction 10-1
Information About Subscribers 10-2
What is a Subscriber? 10-2
Subscriber Modes in Service Control Solutions 10-3
Subscriber Database: Capacity and Limits 10-4
Working with Large Numbers of Subscribers 10-5
Actual Maximum Number of Subscribers 10-5
Subscriber Mapping Limits 10-5
Rate of Creating Anonymous Subscribers 10-5
Aging Subscribers 10-5
VPN-Based Subscribers 10-6
Automatic VLAN VPNs 10-6
Synchronizing Subscriber Information in a Cascade System 10-6
Anonymous Groups and Subscriber Templates 10-7
Subscriber Files 10-8
Subscriber Default csv File Format 10-8
IPv6 Subscriber csv File Format 10-9
Here is an example of a subscriber csv file in the default format: 10-9
Subscriber Anonymous Groups csv File Format 10-9
Importing and Exporting Subscriber Information 10-10
Editing the subaware.pro File 10-10
Options 10-11
How to Import Subscriber Information 10-11
How to Export Subscriber Information 10-11
How to Import a Subscriber Template 10-12
How to Export a Subscriber Template 10-12
Removing Subscribers and Templates 10-13
How to Remove a Specific Subscriber 10-13
Options 10-13
How to Remove All Introduced Subscribers 10-14
How to Remove a Specific Anonymous Subscriber Group 10-14

Options 10-14
How to Remove All Anonymous Subscriber Groups 10-14
How to Remove All the Anonymous Subscribers 10-14
How to Remove All Subscriber Templates 10-15
Removing VPN-based Subscribers 10-15
How to Remove Subscribers by Device 10-15
How to Remove Subscribers from the SM 10-15
How to Remove Subscribers from a Specified SCMP Peer Device 10-17
Creating Anonymous Groups 10-18
Defining Anonymous Groups 10-18
How to Define an Anonymous Group 10-18
Importing and Exporting Anonymous Groups 10-19
How to Import Anonymous Groups 10-19
How to Export Anonymous Groups 10-19
Monitoring Subscribers 10-20
How to Monitor the Subscriber Database 10-20
How to Display the Subscriber Database Counters 10-21
Clearing the Subscriber Database Counters 10-22
Displaying Subscribers 10-22
Displaying Subscribers: All Current Subscriber Names 10-22
Displaying Subscribers: By Subscriber Property or Prefix 10-23
How to Display Subscribers: By Mapping (IP Address, VPN, or VLAN ID) 10-25
How to Display Subscriber Information 10-27
How to Display a Listing of Subscriber Properties 10-28
How to Display Complete Information for a Specified Subscriber 10-28
How to Display Values of Subscriber Properties for a Specified Subscriber 10-28
How to Display Mappings for a Specified Subscriber 10-28
How to Display OS Counters for a Specified Subscriber 10-29
Displaying Anonymous Subscriber Information 10-29
How to Display Currently Configured Anonymous Groups 10-29
How to Display Currently Configured Templates for Anonymous Groups 10-30
How to Display Current Configuration for a Specified Anonymous Group 10-30
How to Display Subscribers in a Specified Anonymous Group 10-30
How to Display All Subscribers Currently in Anonymous Groups 10-30
How to Display the Number of Subscribers in a Specified Anonymous Group 10-31
How to Display the Total Number of Subscribers in All Anonymous Groups 10-31
Configuring the Actual Maximum Number of Subscribers 10-32
How to Override the Configured Capacity Option 10-32
How to Override the Configured Capacity Option in a Cascade Setup 10-32

How to Restore the Configured Capacity Option 10-33
How to Monitor the Maximum Number of Subscribers 10-33
Configuring Subscriber Aging 10-34
How to Enable Aging for Anonymous Group Subscribers 10-34
How to Enable Aging for Introduced Subscribers 10-34
How to Disable Aging for Anonymous Group Subscribers 10-34
How to Disable Aging for Introduced Subscribers 10-35
How to Set the Aging Timeout Period for Anonymous Group Subscribers 10-35
Options 10-35
How to Set the Aging Timeout Period for Introduced Subscribers 10-35
Options 10-35
How to Display Aging for Anonymous Group Subscribers 10-35
How to Display Aging for Introduced Subscribers 10-36
Managing VPNs and VPN Subscriber Mappings 10-37
How to Display VPN-Related Mappings 10-37
How to Clear Automatic VPNs 10-37
Configuring the Cisco SCE Platform/SM Connection 10-39
Configuring the Behavior of the Cisco SCE Platform in Case of Failure of the SM 10-39
Options 10-39
Configuring the SM-SCE Platform Connection Timeout 10-40
Options 10-40

CHAPTER 11 **Redundancy and Failover** 11-1
Introduction 11-1
Redundancy and Failover 11-2
Terminology and Definitions 11-2
Redundant Topologies 11-2
External Bypass 11-3
Hardware Bypass 11-3
In-line Dual Link Redundant Topology 11-3
Failure Detection 11-3
Link Failure Reflection 11-5
Hot Standby and Failover 11-6
Hot Standby 11-6
Failover 11-6
Hardware Crash Mode 11-8
Failure in the Cascade Connection 11-9
Installing a Cascaded System 11-9
Recovery 11-11

Replacing the Cisco SCE Platform (Manual Recovery) 11-11
Manual Steps 11-11
Automatic Steps (in parallel with the manual steps, requires no user intervention): 11-12
Reboot Only (Fully Automatic Recovery) 11-12
CLI Commands for Cascaded Systems 11-13
Topology-Related Parameters for Redundant Topologies 11-13
Configuring the Connection Mode 11-13
Examples 11-14
Monitoring the System 11-14
How to View the Current Connection Mode 11-14
How to View the Cisco SCE-ID 11-15
How to View the Current Redundancy Status of the Cisco SCE Platform 11-15
How to View Information about the Peer Cisco SCE Platform 11-15
How to View Information about the Cascade Connections 11-16
How to View the Current Link to Port Mappings 11-16
How to View the Current Link Mode 11-17
Configuring Forced Failure 11-18
System Upgrades 11-19
Firmware Upgrade (package installation) 11-19
Application Upgrade 11-19
Simultaneous Upgrade of Firmware and Application 11-20

CHAPTER 12 **Identifying and Preventing Distributed Denial-of-Service Attacks** 12-1
Introduction 12-1
Attack Filtering and Attack Detection 12-2
Attack Filtering 12-2
Specific Attack Filtering 12-2
Attack Detection 12-4
Attack Detection Thresholds 12-4
Attack Handling 12-5
Subscriber Notification 12-6
Hardware Filtering 12-6
Configuring Attack Detectors 12-8
Enabling Specific-IP Detection 12-10
Options 12-10
How to Enable Specific-IP Detection 12-10
How to Enable Specific-IP Detection for the TCP Protocol Only for all Attack Directions 12-11
How to Enable Specific-IP Detection for the TCP Protocol for Port-Based Detections Only for Dual-Sided Attacks 12-11

How to Disable Specific-IP Detection for Protocols Other than TCP, UDP, and ICMP for all Attack Directions 12-11
How to Disable Specific-IP Detection for ICMP for Single-Sided Attacks Defined by the Source IP 12-11
Configuring the Default Attack Detector 12-11
Options 12-12
How to Define the Default Action and Optionally, the Default Thresholds 12-13
How to Reinstate the System Defaults for a Selected Set of Attack Types 12-13
How to Reinstate the System Defaults for All Attack Types 12-14
Specific Attack Detectors 12-14
Options 12-14
How to Enable a Specific Attack Detector and Assign it an ACL 12-15
How to Define the Action and Optionally the Thresholds for a Specific Attack Detector 12-16
How to Define the Subscriber Notification Setting for a Specific Attack Detector 12-16
How to Define the SNMP Trap Setting for a Specific Attack Detector 12-16
How to Define the List of Destination Ports for TCP or UDP Protocols for a Specific Attack Detector 12-17
How to Delete User-Defined Values 12-17
How to Disable a Specific Attack Detector 12-17
How to Disable All Non-default Attack Detectors 12-18
How to Disable All Attack Detectors 12-18
Sample Attack Detector Configuration 12-18
Subscriber Notifications 12-20
Configuring the Subscriber Notification Port 12-20
Options 12-20
How to Remove the Subscriber Notification Port 12-20
Preventing and Forcing Attack Detection 12-21
Options 12-21
Preventing Attack Filtering 12-21
How to Remove All dont-filter Settings 12-22
Forcing Attack Filtering 12-22
How to Remove All force-filter Settings 12-23
Monitoring Attack Filtering 12-24
Monitoring Attack Filtering Using SNMP Traps 12-24
Monitoring Attack Filtering Using CLI Commands 12-26
How to Display a Specified Attack Detector Configuration 12-26
How to Display the Default Attack Detector Configuration 12-28
How to Display All Attack Detector Configurations 12-28
How to Display Filter State (Enabled or Disabled) 12-29
How to Display Configured Threshold Values and Actions 12-29

How to Display the Current Counters 12-30
How to Display all Currently Handled Attacks 12-31
How to Display all Existing Force-Filter Settings 12-31
How to Display all Existing Don't-Filter Settings 12-31
How to Display the List of Ports Selected for Subscriber Notification 12-31
How to Find out Whether Hardware Attack Filtering has been Activated 12-32
Viewing the Attack Log 12-32
The Attack Log 12-32
How to View the Attack Log 12-33
How to Copy the Attack Log to a File 12-33

CHAPTER 13 **Managing the SCMP** 13-1
Introduction 13-1
About SCMP 13-2
SCMP Terminology 13-3
Deployment Scenarios 13-3
Single ISG Router with a Single Cisco SCE Platform (1xISG – 1xCisco SCE) 13-4
Single ISG Router with Two Cascaded Cisco SCE Platforms (1xISG – 2xCisco SCE) 13-4
Multiple ISG Routers with Two Cascaded Cisco SCE Platforms (NxISG – 2xCisco SCE) 13-5
Multiple ISG Routers with Multiple Cisco SCE Platforms via Load Balancing (NxISG – MxCisco SCE) 13-6
SCMP Peer Devices 13-7
Connection Management 13-7
SCMP Subscriber Management 13-8
GUID and Subscriber ID 13-8
Configuring the SCMP 13-9
Configuring SCMP Parameters 13-9
How to Enable the SCMP 13-9
How to Disable the SCMP 13-10
How to Configure the SCMP Peer Device to Push Sessions 13-10
Configuring the SCMP Peer Device to Force Each Subscriber to Single Cisco SCE Platform 13-10
Defining the Keep-alive Interval Parameter 13-11
Defining the Reconnect Interval Parameter 13-11
Defining the Loss-of-Sync Timeout Parameter 13-12
Adding an SCMP Peer Device 13-12
How to Define an SCMP Peer Device 13-12
How to Assign the SCMP Peer Device to an Anonymous Group 13-13
Deleting Subscribers Managed by an SCMP Peer Device 13-14
Options 13-14

Deleting an SCMP Peer Device 13-14
Defining the Subscriber ID 13-14
Options 13-15
Configuring the RADIUS Client 13-15
Options 13-16
Monitoring the SCMP Environment 13-17
How to Monitor the SCMP 13-17
Options 13-17
How to display the general SCMP configuration 13-17
How to display the configuration all currently defined SCMP peer devices 13-18
How to display the configuration for a specified SCMP peer device 13-18
How to display the statistics for all SCMP peer devices 13-18
How to display the statistics for a specified SCMP peer device 13-19
Monitoring the RADIUS Client 13-19

CHAPTER 14 **Value-Added Services (VAS) Traffic Forwarding** 14-1
Introduction 14-1
Information About VAS Traffic Forwarding 14-2
VAS Service Goals 14-2
How VAS Traffic Forwarding Works 14-3
Requirements for VAS Servers 14-4
VAS Traffic Forwarding and SCA BB 14-5
VLAN Tags for VAS Traffic Forwarding 14-5
Service Flow 14-5
Data Flow 14-6
Non-VAS Data Flow 14-7
VAS Data Flow 14-7
Load Balancing 14-8
Load Balancing and Subscribers 14-8
Load Balancing and Subscriber Mode 14-9
VAS Redundancy 14-10
VAS Server Failure 14-10
VAS Server Group Failure 14-10
Ethernet Switch Failure 14-11
Disabling a VAS Server 14-11
VAS Status and VAS Health Check 14-12
VAS Server States 14-13
VAS Traffic Forwarding Topologies 14-14
Single Cisco SCE Platform, Multiple VAS Servers 14-14

Data Flow 14-15
Multiple Cisco SCE Platforms, Multiple VAS Servers 14-15
SNMP Support for VAS 14-17
Interactions Between VAS Traffic Forwarding and Other Cisco SCE Platform Features 14-18
Incompatible Cisco SCE Platform Features 14-18
VAS Traffic Forwarding and DDoS Processing 14-18
Specific IP DDoS Attack Detection 14-18
Specific IP Attack Filter 14-18
VAS Traffic Forwarding and Bandwidth Management 14-19
Global Controllers and VAS Flows 14-19
Configuring VAS Traffic Forwarding 14-20
Configuring VAS Traffic Forwarding from the SCA BB Console 14-21
Global Options 14-21
Enabling VAS Traffic Forwarding 14-21
Options 14-21
Disabling VAS Traffic Forwarding 14-22
Configuring the VAS Traffic Link 14-22
Options 14-23
How to Select the Link for VAS Traffic 14-23
How to Revert to the Default Link for VAS Traffic 14-23
Configuring a VAS Server 14-23
Options 14-24
How to Enable a VAS Server 14-24
How to Disable a VAS Server 14-24
How to Restore all VAS Server Properties to Default 14-24
Assigning a VLAN ID to a VAS Server 14-24
Options 14-25
How to Configure the VLAN Tag Number for a Specified VAS Server 14-25
How to Remove the VLAN Tag Number from a Specified VAS Server 14-25
Configuring the Health Check 14-25
How to Enable VAS Server Health Check 14-26
How to Disable VAS Server Health Check 14-27
How to Define the UDP Ports to be Used for Health Check 14-27
How to Remove the UDP Ports Configuration 14-27
Configuring Pseudo IP Addresses for the Health Check Packets 14-27
Configuring a VAS Server Group 14-28
Adding and Removing Servers 14-29
Configuring VAS Server Group Failure Parameters 14-29
VAS Configuration Example 14-31

Monitoring VAS Traffic Forwarding 14-32
How to Display Global VAS Status and Configuration 14-32
Example 14-32
How to Display Operational and Configuration Information for a Specific VAS Server Group 14-33
Example 14-33
How to Display Operational and Configuration Information for All VAS Server Groups 14-33
How to Display Operational and Configuration Information for a Specific VAS Server 14-33
Example 14-33
How to Display Operational and Configuration Information for All VAS Servers 14-34
How to Display the VAS Servers Used by a Specified Subscriber 14-34
How to Display Health Check Counters for a Specified VAS Server 14-34
Example 14-34
How to Display Health Check Counters for All VAS Servers 14-35
How to Clear the Health Check Counters for a Specified VAS Server 14-35
How to Clear the Health Check Counters for All VAS Servers 14-35
Intelligent Traffic Mirroring 14-36
Using Traffic Mirroring for Behavioral Targeting 14-36
How Traffic Mirroring Works 14-37
Traffic Mirroring and SCA BB 14-37
Mirroring Termination 14-38
Mirroring Exceptions 14-38
Mirroring the TCP-Segmented HTTP GET Packets 14-38
Cisco SCE Connectivity 14-39
Traffic Mirroring and Bandwidth Management 14-41
Configuring Traffic Mirroring 14-41
Monitoring Traffic Mirroring 14-42
Traffic Mirroring Sample Configuration 14-42

APPENDIX A **Cisco Service Control MIBs** A-1
Introduction A-1
MIB Files A-2
Loading MIBs A-4
pcube to Cisco MIB Mapping A-5
Pcube Engage MIB (CISCO-SCAS-BB-MIB) A-6
pcube to Cisco MIB Mapping: Detailed OID Mappings A-7
Cisco SCE Platform-Specific MIB Information A-26
CISCO-ENTITY-ALARM-MIB A-26
MIB Updates A-27
Release 3.5.5 MIB Updates A-27

CISCO-SERVICE-CONTROL-TP-STATS-MIB A-27
Release 3.6.0 MIB Updates A-27
CISCO-PROCESSOR-MIB A-27
ENTITY-MIB A-27
Index Changes A-28
Release 3.6.5 MIB Updates A-29
Temperature Sensor Traps Updated A-29
Release 3.7.0 MIB Updates A-29
SNMP Support for Aggregative Global Controllers A-29
linkUp and linkDown Notification Traps A-29
Release 4.1.0 MIB Updates A-30
SNMP TRAP for Global Attacks A-30
SNMP Walk Functionality for Temperature MIBs A-30

APPENDIX B **Monitoring Cisco SCE Platform Utilization** B-1
Introduction B-1
Cisco SCE Platform Utilization Indicators B-2
CPU Utilization B-2
Flows Capacity B-2
Subscribers Capacity B-2
Service Loss B-3
Monitoring Service Loss B-3

APPENDIX C **Cisco SCE 8000 Licensing Information** C-1
OpenSSH License C-1
NetSNMP License C-9



About this Guide

Revised: September 23, 2015

Introduction

This preface describes who should read *Cisco SCE 8000 10GBE Software Configuration Guide,* how it is organized, and its document conventions.

This guide is for experienced network administrators who are responsible for configuring and maintaining the Cisco SCE platform.

Document Revision History

The following Document Revision History records the changes made to this document.

Table 1 Document Revision History

Revision	Cisco Service Control Release and Date	Change Summary
OL-30621-02	Release 4.1.x February 07, 2014	Updated "MIB Updates" section on page A-27 with limitations on linkUp/linkDown trap.
OL-30621-01	Release 4.1.x December 23, 2013	First version of this document (new for the Release 4.1.x train). The following changes were made from the last release of the 4.0.x train: • Updated "Configuring and Managing the SNMP Interface" section on page 5-41 with SNMPv3 details. • Updated the "Tunneling Protocols" section on page 7-4. • Added "Release 4.1.0 MIB Updates" section on page A-30.

Organization

This guide contains the following sections.

Table 2 Document Organization

Section	Title	Description
Chapter 1	Cisco Service Control Overview	Overview of Cisco SCE platform management.
Chapter 2	Command-Line Interface	Detailed explanation of how to use the Cisco SCE Command-line Interface.
Chapter 3	Basic Cisco SCE 8000 Platform Operations	Explanation of how to manage configurations, install applications and upgrade the system software.
Chapter 4	Utilities	Explanation of the setup wizard and the user log, as well as of file operations.
Chapter 5	Configuring the Management Interface and Security	Explanation of how to configure the various management options: Telnet, SSH, and SNMP. Also how to configure the system time, Domain Name Settings, management IP address, and passwords.
Chapter 6	Global Configuration	Explanation of how to configure various global settings, such as system time, Domain Name Settings, and IP routing.
Chapter 7	Configuring Line Interfaces	Explanation of how to configure tunneling, TOS marking, and traffic rules.
Chapter 8	Configuring the Connection	Explanation of how to configure the connection mode, link mode, and failure behaviors
Chapter 9	Raw Data Formatting: The RDR Formatter and NetFlow Exporting	Explanation of how to configure the connection mode, link mode, and failure behaviors.
Chapter 10	Managing Subscribers	Explanation of how to import and export subscriber information and how to monitor subscribers.
Chapter 11	Redundancy and Failover	Explanation of how to configure and manage a redundant system.
Chapter 12	Identifying and Preventing Distributed Denial-of-Service Attacks	Explanation of how to configure attack filtering.
Chapter 13	Managing the SCMP	Explanation of Service Control Management Protocol (SCMP), which is a protocol that integrates the Cisco SCE platform and the ISG (Intelligent Service Gateway) functionality of the Cisco routers. It also explains how to configure and manage SCMP, SCMP peer devices and the RADIUS client.

Table 2 Document Organization (continued)

Section	Title	Description
Chapter 14	Value-Added Services (VAS) Traffic Forwarding	Explanation of Value Added Services (VAS) traffic forwarding and how to configure it. Also explains how the same capabilities are used for traffic mirroring.
Appendix A	Cisco Service Control MIBs	Explanation of how to map the proprietary pcube MIB supported in previous releases to the new MIB structure.
Appendix B	Monitoring Cisco SCE Platform Utilization	Explanation of how to monitor Cisco SCE platforms that are installed in real traffic.
Appendix C	Cisco SCE 8000 Licensing Information	Copy of Open SSH and NetSNMP license information.

Related Publications

Your Cisco SCE platform and the software running on it contain extensive features and functionality, which are documented in the following resources:

- For further information regarding the Service Control CLI and a complete listing of all CLI commands, refer to the *Cisco SCE8000 CLI Command Reference*
- For initial installation and startup information, refer to the relevant installation guide:
  - *Cisco SCE8000 10GBE Installation and Configuration Guide*
- For international agency compliance, safety, and statutory information for wide-area network (WAN) interfaces for the Cisco SCE 2000 platform, refer to the regulatory and safety information document:
  - *Regulatory Compliance and Safety Information for Cisco SCE8000*
- For installation and configuration of the other components of the Service Control Management Suite refer to:
  - *Cisco SCMS Subscriber Management User Guide*
  - *Cisco SCMS Collection Manager User Guide*
  - *Cisco Service Control Application for Broadband User Guide*
  - *Cisco Insight User Guide*
- To view Cisco documentation or obtain general information about the documentation, refer to the following sources:
  - *Obtaining Documentation and Submitting a Service Request, page 7*
  - The Cisco Information Packet that shipped with your Cisco SCE 8000 platform.

Conventions

This document uses the following conventions.

Table 3 Conventions

Convention	Indication
bold font	Commands and keywords and user-entered text appear in **bold** font.
italic font	Document titles, new or emphasized terms, and arguments for which you supply values are in *italic* font.
[ ]	Elements in square brackets are optional.
{x | y | z}	Required alternative keywords are grouped in braces and separated by vertical bars.
[x | y | z]	Optional alternative keywords are grouped in brackets and separated by vertical bars.
string	A nonquoted set of characters. Do not use quotation marks around the string or the string will include the quotation marks.
`courier` font	Terminal sessions and information the system displays appear in `courier` font.
< >	Nonprinting characters such as passwords are in angle brackets.
[ ]	Default responses to system prompts are in square brackets.
!, #	An exclamation point (!) or a pound sign (#) at the beginning of a line of code indicates a comment line.



Note Means *reader take note*.



Tip Means *the following information will help you solve a problem*.



Caution Means *reader be careful*. In this situation, you might perform an action that could result in equipment damage or loss of data.



Timesaver Means *the described action saves time*. You can save time by performing the action described in the paragraph.



Warning **Means *reader be warned*. In this situation, you might perform an action that could result in bodily injury.**

Obtaining Documentation and Submitting a Service Request

For information on obtaining documentation, submitting a service request, and gathering additional information, see the monthly What's New in Cisco Product Documentation, which also lists all new and revised Cisco technical documentation, at:

http://www.cisco.com/en/US/docs/general/whatsnew/whatsnew.html

Subscribe to the *What's New in Cisco Product Documentation* as an RSS feed and set content to be delivered directly to your desktop using a reader application. The RSS feeds are a free service. Cisco currently supports RSS Version 2.0.



Cisco Service Control Overview

Revised: September 23, 2015

Introduction

This chapter provides a general overview of the Cisco Service Control solution. It introduces the Cisco service control concept and capabilities.

It also briefly describes the hardware capabilities of the service control engine (Cisco SCE) platform and the Cisco specific applications that together compose the complete Cisco service control solution.

- Cisco Service Control Solution, page 1-2
- Cisco Service Control Capabilities, page 1-3
- Cisco SCE Platform Description, page 1-4
- Management and Collection, page 1-6
- IPv6 Support, page 1-8

Cisco Service Control Solution

The Cisco service control solution is delivered through a combination of hardware and specific software solutions that address various operational and business-related challenges. Service providers can use the Cisco SCE platform to support classification, analysis, and control of Internet and IP traffic.

Service control enables service providers to:

- Capitalize on existing infrastructure.
- Analyze, charge for, and control IP network traffic at multigigabit wire line speeds.
- Identify and target high-margin content-based services and enable their delivery.

As access and bandwidth have become commodities where prices continually fall and profits disappear, service providers have realized that they must offer value-added services to derive more revenue from the traffic and services running on their networks.

Cisco service control solutions allow the service provider to capture profits from IP services through detailed monitoring, precise, real-time control, and awareness of applications as they are delivered.

Service Control for Broadband Service Providers

Service providers of any access technology (DSL, cable, mobile, and so on) targeting residential and business consumers must find new ways to get maximum leverage from their existing infrastructure, while differentiating their offerings with enhanced IP services.

The Cisco service control application for broadband adds a layer of service intelligence and control to existing networks that can:

- Report and analyze network traffic at subscriber and aggregate level for capacity planning
- Provide customer-intuitive tiered application services and guarantee application service level agreements (SLAs)
- Implement different service levels for different types of customers, content, or applications
- Identify network abusers who are violating the acceptable use policy (AUP)
- Identify and manage peer-to-peer traffic, NNTP (news) traffic, and spam abusers
- Enforce the AUP
- Integrate Service Control solutions easily with existing network elements and business support systems (BSS) and operational support systems (OSS)

Cisco Service Control Capabilities

The core of the Cisco service control solution is the network hardware device: the Service control engine (Cisco SCE). The core capabilities of the Cisco SCE platform, which support a wide range of applications for delivering service control solutions, include:

- Subscriber and application awareness—Application-level drilling into IP traffic for real-time understanding and controlling of usage and content at the granularity of a specific subscriber.
  - Subscriber awareness—The ability to map between IP flows and a specific subscriber to maintain the state of each subscriber transmitting traffic through the Cisco SCE platform and to enforce the appropriate policy on this subscriber's traffic.

    Subscriber awareness is achieved either through dedicated integrations with subscriber management repositories, such as a DHCP or a RADIUS server, or through sniffing of RADIUS or DHCP traffic.
  - Application awareness—The ability to understand and analyze traffic up to the application protocol layer (Layer 7).

    For application protocols implemented using bundled flows (such as FTP, which is implemented using Control and Data flows), the Cisco SCE platform understands the bundling connection between the flows and treats them accordingly.
- Application-layer, stateful, real-time traffic control—The ability to perform advanced control functions, including granular bandwidth (BW) metering and shaping, quota management, and redirection, using application-layer, stateful, real-time traffic transaction processing. This requires highly adaptive protocol and application-level intelligence.
- Programmability—The ability to quickly add new protocols and adapt to new services and applications in the service provider environment. Programmability is achieved using the Cisco Service Modeling Language (SML).

  Programmability allows new services to be deployed quickly and provides an easy upgrade path for network, application, or service growth.
- Robust and flexible back-office integration—The ability to integrate with existing third-party systems at the service provider, including provisioning systems, subscriber repositories, billing systems, and OSS systems. The Cisco SCE provides a set of open and well-documented APIs that allows a quick integration process.
- Scalable high-performance service engines—The ability to perform all of these operations at wire speed.

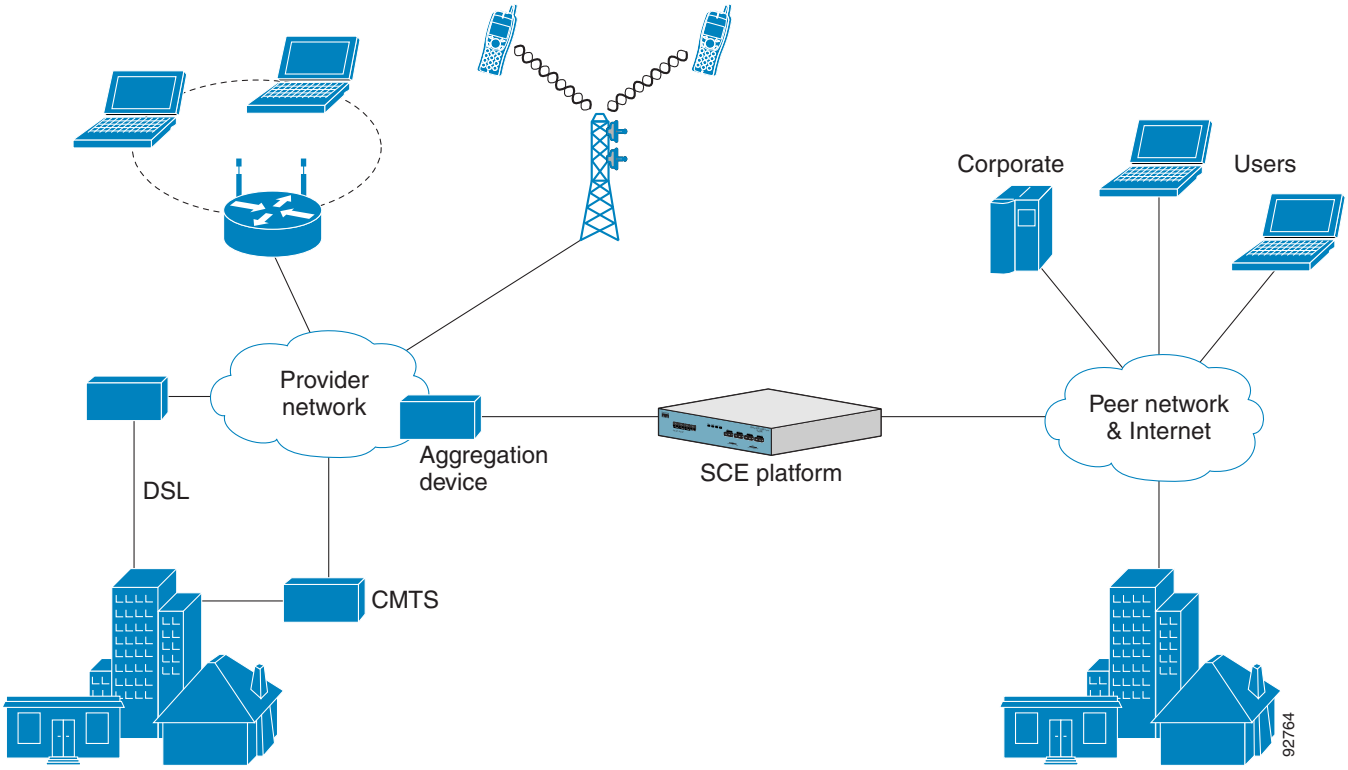
Cisco SCE Platform Description

The Cisco SCE family of programmable network devices performs application-layer stateful-flow inspection of IP traffic, and controls the traffic based on configurable rules. The Cisco SCE platform is a network device that uses ASIC components and reduced instruction set computer (RISC) processors to exceed beyond packet counting and expand into the contents of network traffic. Providing programmable, stateful inspection of bidirectional traffic flows, and mapping these flows with user ownership, Cisco SCE platforms provide real-time classification of network use. The classification provides the basis of the Cisco SCE platform advanced traffic-control and bandwidth-policing functionality. Where most bandwidth control functionality ends, the Cisco SCE platform provides further control and shaping options, including:

- Layer 7 stateful wire-speed packet inspection and classification
- Robust support for more than 600 protocols and applications, including:
  - General—HTTP, HTTPS, FTP, Telnet, Network News Transfer Protocol (NNTP), Simple Mail Transfer Protocol (SMTP), Post Office Protocol 3 (POP3), Internet Message Access Protocol (IMAP), Wireless Application Protocol (WAP), and others
  - Peer-to-Peer (P2P) file sharing—FastTrack-KazaA, Gnutella, BitTorrent, Winny, Hotline, eDonkey, DirectConnect, Piolet, and others
  - P2P VoIP—Skype, Skinny, DingoTel, and others
  - Streaming and Multimedia—Real Time Streaming Protocol (RTSP), Session Initiation Protocol (SIP), HTTP streaming, Real Time Protocol (RTP) and Real Time Control Protocol (RTCP), and others
- Programmable system core for flexible reporting and bandwidth control
- Transparent network and BSS and OSS integration into existing networks
- Subscriber awareness that relates traffic and usage to specific customers

Figure 1-1 illustrates a common deployment of a Cisco SCE platform in a network.

Figure 1-1 Cisco SCE Platform in the Network



Bandwidth Management of P2P Traffic

The Cisco SCE uses unique signatures to identify the networking flows of P2P, IM, and other applications. While defining packages to subscribers, you can create rules for different types of applications such as P2P, and IM and if required, associate these rules to separate Bandwidth Controls (BWCs). With BWC enforcement, you can limit the networking flows for all types of applications. There are three types of rules in the Cisco SCE which can be used for bandwidth enforcement at different levels.

- P2P based BWC

If the Cisco SCE is configured to enforce BWC based on peer-to-peer traffic, it detects the application based on its signature. Cisco SCE then includes the amount of network flows of P2P traffic and calculates the bandwidth accordingly. The consumed bandwidth is the sum of P2P data and the control traffic. Bandwidth limitation takes place as per the enforcement configured in the BWC.

- Default Service BWC

When an application is configured with discrete BWC, the Cisco SCE does not relate the amount of networking flows of the application when calculating the bandwidth consumed by it. The amount of networking flows consumed by the application is accounted with the Default Service. If there is any rate limit associated with Default Service BWC, this amount is accounted with Default Service BWC.

- No Cisco SCE Enforcement

No bandwidth control is enforced upon the subscribers. This results in unlimited bandwidth to the subscriber.

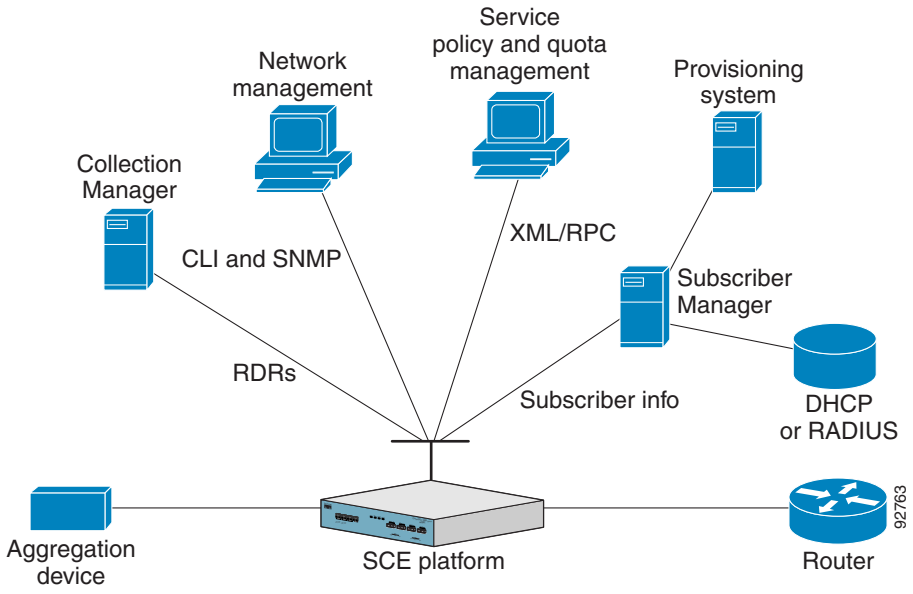
Management and Collection

The Cisco service control solution includes a complete management infrastructure that provides the following management components to manage all aspects of the solution:

- Network management
- Subscriber management
- Service Configuration management

These management interfaces are designed to comply with common management standards and to integrate easily with existing OSS infrastructure (Figure 1-2).

Figure 1-2 Service Control Management Infrastructure



Network Management

The Cisco service control solution provides complete network Fault, Configuration, Accounting, Performance, Security (FCAPS) Management.

Two interfaces provide network management:

- Command-line interface (CLI)—Accessible through the Console port or through a Telnet connection, the CLI is used for configuration and security functions.
- SNMP—Provides fault management (through SNMP traps) and performance-monitoring functionality.

Subscriber Management

Where the Cisco service control application for broadband (SCA BB) enforces policies on different subscribers and tracks usage on an individual subscriber basis, the Cisco service control management suite (SCMS) subscriber manager (SM) may be used as middleware software for bridging between OSS and Cisco SCE platforms. Subscriber information is stored in the SM database and can be distributed between multiple platforms according to actual subscriber placement.

The SM provides subscriber awareness by mapping network IDs to subscriber IDs. It can obtain subscriber information using dedicated integration modules that integrate with AAA devices, such as RADIUS or DHCP servers.

Subscriber information may be obtained in one of two ways:

- Push Mode—The SM pushes subscriber information to the Cisco SCE platform automatically upon logon of a subscriber.
- Pull Mode—The SM sends subscriber information to the Cisco SCE platform in response to a query from the Cisco SCE platform.

Service Configuration Management

Service configuration management is the ability to configure the general service definitions of a service control application. A service configuration file containing settings for traffic classification, accounting and reporting, and control is created and applied to a Cisco SCE platform. The SCA BB application provides tools to automate the distribution of these configuration files to Cisco SCE platforms. This standards-based approach makes it easy to manage multiple devices in a large network.

Service Control provides a GUI to edit and create these files and a complete set of APIs to automate their creation.

Data Collection

Data collection occurs as follows:

1. All analysis and data processing functions of the Cisco SCE platform result in the generation of Raw Data Records (RDRs), which the Cisco SCE platform forwards using a simple TCP-based protocol (RDR-Protocol).
2. RDRs are processed by the Cisco service control management suite collection manager.
3. The collection manager software is an implementation of a collection system that receives RDRs from one or more Cisco SCE platforms. It collects these records and processes them in one of its adapters. Each adapter performs a specific action on the RDR.

RDRs contain a variety of information and statistics, depending on the configuration of the system. Three main categories of RDRs include:

- Transaction RDRs—Records generated for each *transaction*, where a transaction is a single event detected in network traffic. The identification of a transaction depends on the particular application and protocol.
- Subscriber Usage RDRs—Records generated per subscriber, describing the traffic generated by that subscriber for a defined interval.
- Link RDRs—Records generated per link, describing the traffic carried over the link for a defined interval.

IPv6 Support

The Cisco SCE 8000 devices support processing of IPv6 traffic. The features that are available for IPv4, such as traffic processing, application classification and control, and management APIs, are available for IPv6 too.

The Cisco SCOS enhances the IPv6 capabilities of Cisco SCE 8000 devices by providing support to both IPv4 and IPv6 traffic simultaneously on all traffic processors.

Supported IPv6 features:

- Subscriber awareness.
- Virtual Gi
- Subscriber based classification for IPv6 traffic.
- Subscriber based bandwidth control.
- Content filtering for IPv6 traffic flows.
- Dual stack support on all traffic processors.
- Reporting and monitoring of IPv6 traffic.
- Support for IPv6 Fragmentation and Zones.
- IPv6 bundling support in DS-lite.
- Introduce IPv6 subscribers through Gx.
- Raw Data Records, except Attack RDRs and Malicious Traffic RDRs, are enhanced for IPv6.

  Redirection and blocking of traffic passing through 6to4, 6rd, and DS-Lite tunnels.

Cisco SCE devices work on three different system modes, namely, IPv4 only mode, IPv6 only mode, and dual stack mode. The default mode is dual stack mode. To configure the system mode, see the "Configuring the System Mode" section on page 3-17.

The following limitations are applicable to the IPv6 features on Cisco SCE Release 4.0.x:

- IPv6 addresses for connectivity to the management interfaces are not supported.
- Tunneling—The L2TP, GRE, GTP, MPLS, MPLS-VPN tunnels are not supported.
- Flow filter limitations:
  - For TCP flags, the PSH and URG fields are not considered, but SYN, ACK, RST, and FIN are considered.
  - Range options are not provided for subscriber side IP and network side IP address; instead, prefix length is used.
  - Inverse support is not provided for any of the fields.
- Value-Added Services (VAS) are not supported.
- Service classification and global bandwidth control are applicable for both IPv4 and IPv6. For example, there is only one browsing service for both IPv6 and IPv4 browsing traffic. Allocating the global bandwidth controller to a certain service limits the IPv6 and IPv4 traffic within that service.
- The Services over IPv6 are classified under the same service IDs as the corresponding services over IPv4.
- Flow capture is not supported for IPv6.

- Cisco SCE 8000 supports a maximum of 1M subscriber range. This means that the Cisco SCE can support a maximum of 1M subscribers with one mapping (either IPv4 or IPv6). But when the dual stack mode is enabled and all subscribers are dual stack subsribers—subscribers with one IPv4 and one IPv6 address, then the device supports only upto 500,000 dual stack subscribers.
- Cisco SCE 8000 devices identifies the IPv6 subscribers based on the MSB 64 bits of the subscriber IPv6 address. Cisco SCE 8000 devices support IPv6 subscribers within a range /32 to /64 and not less than /32.
- Virtual Gi is supported only in standalone configuration and not in cascade configuration.



Command-Line Interface

Revised: September 23, 2015

Introduction

This chapter describes how to use the Cisco SCE platform Command-Line Interface (CLI), its hierarchical structure, authorization levels and its help features. The Command-Line Interface is one of the Cisco SCE platform management interfaces.

The CLI is accessed through a Telnet session or directly via the console port on the front panel of the Cisco SCE platform. When you enter a Telnet session, you enter as the simplest level of user, in the User Exec mode.

The Cisco SCE platform supports up to eleven concurrent CLI sessions; five sessions initiated by Telnet connection, five sessions via SSH connection, and one session on the console port.

- Authorization and Command Mode Levels (Hierarchy), page 2-2
- CLI Help Features, page 2-9
- Navigational and Shortcut Features, page 2-11
- Managing Command Output, page 2-14
- Creating a CLI Script, page 2-16

Authorization and Command Mode Levels (Hierarchy)

When using the CLI there are two important concepts that you must understand to navigate:

- Authorization Level—Indicates the level of commands you can execute. A user with a simple authorization level can only view some information in the system, while a higher level administrator can actually make changes to configuration.

  This manual documents commands up to and including the admin authorization level.

- Command Hierarchy Level—Provides you with a context for initiating commands. Commands are broken down into categories and you can only execute each command within the context of its category. For example, to configure parameters related to the Line Card, you need to be within the Linecard Interface Configuration Mode. (See "CLI Command Mode Hierarchy" section on page 2-3.)

The following sections describe the available Authorization and Command Hierarchy Levels and how to maneuver within them.

The on-screen prompt indicates both your authorization level and your command hierarchy level, as well as the assigned hostname.

Note Throughout the manual, `SCE` is used as the sample host name.

CLI Authorization Levels

The Cisco SCE platform has four authorization levels, which represent the user access permissions. When you initially connect to the Cisco SCE platform, you automatically have the most basic authorization level, that is User, which allows minimum functionality.

To monitor the system, you must have Viewer authorization, while to perform administrative functions on the Cisco SCE platform, you must have Admin or Root authorization. A higher level of authorization is accessed by logging in with appropriate password, as described in the procedures below.

In each authorization level, all the commands of the lower authorization layers are available in addition to commands that are authorized only to the current level.

The following CLI commands are related to authorization levels:

- **enable**
- **disable**

Each authorization level has a value (number) corresponding to it. When using the CLI commands, use the values, not the name of the level, as shown in Table 2-1.

Table 2-1 Authorization Levels

Level	Description	Value	Prompt
User	Password required. This level enables basic operational functionality.	0	>
Viewer	Password required. This level enables monitoring functionality. All show commands are available to the Viewer authorization level, with the exception of those that display password information.	5	>

Table 2-1 Authorization Levels (continued)

Level	Description	Value	Prompt
Admin	Password required. For use by general administrators, the Admin authorization level enables configuration and management of the Cisco SCE platform.	10	#
Root	Password required. For use by technical field engineers, the Root authorization level enables configuration of all advanced settings, such as debug and disaster recovery. The Root level is used by technical engineers only.	15	#>

CLI Command Mode Hierarchy

The set of all CLI commands is grouped in hierarchical order, according to the type of the commands. The first two levels in the hierarchy are the User Exec and Privileged Exec modes. These are non-configuration modes in which the set of available commands enables the monitoring of the Cisco SCE platform, file system operations, and other operations that cannot alter the configuration of the Cisco SCE platform.

The next levels in the hierarchy are the Global and Interface configuration modes, which hold a set of commands that control the global configuration of the Cisco SCE platform and its interfaces. Any of the parameters set by the commands in these modes should be saved in the startup configuration, such that in the case of a reboot, the Cisco SCE platform restores the saved configuration.

Table 2-2 shows the available CLI modes.

Table 2-2 CLI Modes

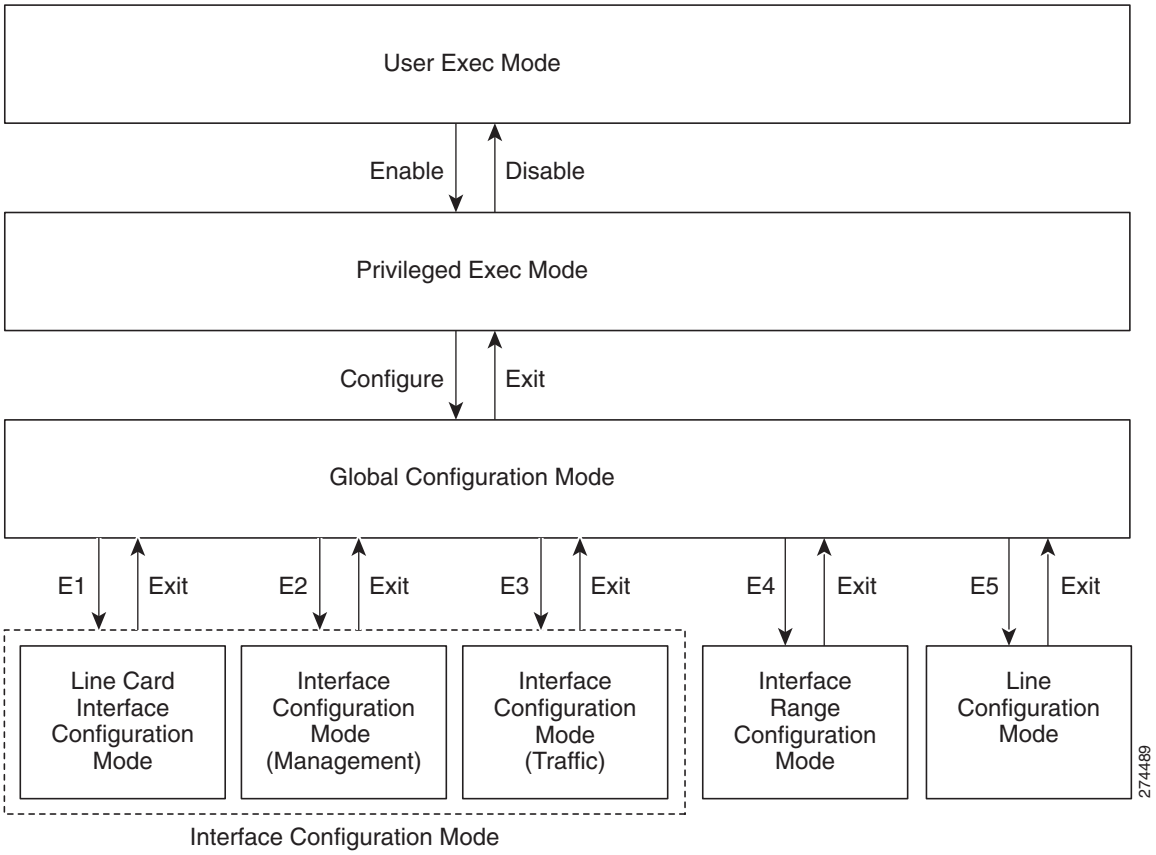
Mode	Description	Level	Prompt indication
User Exec	Initial mode. Also allows monitoring of the system (**show** commands).	User/Viewer	SCE>
Privileged Exec	General administration; file system manipulations and control of basic parameters that do not change the configuration of the Cisco SCE platform.	• Admin • Root	• SCE# • SCE#>
Global Configuration	Configuration of general system parameters, such as DNS, host name, and time zone.	• Admin • Root	• SCE(config)# • SCE(config)#>
Management Interface Configuration	Configuration of management interface parameters, such as the Ethernet interface properties and selection of the active port.	• Admin • Root	• SCE(config if)# • SCE(config if)#>
Interface Configuration	Configuration of specific system interface parameters, for the following interface modes. • linecard interface • specific traffic interface	• Admin • Root	• SCE(config if)# • SCE(config if)#>
Interface Range Configuration	Configuration of a range of traffic interfaces.	• Admin • Root	• SCE(config if range)# • SCE(config if range)#>
Line Configuration	Configuration of Telnet lines, such as an access-list.	• Admin • Root	• SCE(config-line)# • SCE(config-line)#>

When you login to the system, you have the User authorization level and enter User Exec mode. Changing the authorization level to Viewer does not change the mode. Changing the authorization level to Admin automatically moves you to Privileged Exec mode. To move to any of the configuration modes, you must enter command specific to that mode.

The list of available commands in each mode can be viewed using the question mark '?' at the end of the prompt.

Figure 2-1 illustrates the hierarchical structure of the CLI modes, and the CLI commands used to enter and exit a mode.

Figure 2-1 CLI Command Modes



The following commands are used to enter the various specific configuration modes from the global configuration mode:

- E1—**interface Linecard 0**
- E2 (management ports)— **interface Mng** *0/1, 0/2* **OR interface GigabitEthernet** *1/1, 1/2*
- E3 (traffic ports)—**interface TenGigabitEthernet** *3/0/0*, *3/1/0*, *3/2/0*, or *3/3/0*
- E4—**interface range TenGigabitEthernet** *3/<bay-range (any range between 0 and 3)>/0*
- E5—**line vty 0**

Note Although the system supports up to five concurrent Telnet connections, you cannot configure them separately. This means that any number you enter in the **line vty** command (**0**, **1**, **2**, **3** or **4**) will act as a **0** and configure all five connections together.

Note In order for the auto-completion feature to work, when you move from one interface configuration mode to another, you must first exit the current interface configuration mode (as illustrated in Figure 2-1).

Example:

This example illustrates moving into and out of configuration modes as follows:

- Enter global configuration mode
- Configure the Cisco SCE platform time zone
- Enter MNG (management) Interface configuration mode
- Configure the speed of the management interface
- Exit the MNG configuration mode to the global configuration mode
- Enter the Linecard Interface configuration
- Define the link mode
- Exit Linecard Interface configuration mode to user exec mode

```
SCE#configure
SCE(config)#clock timezone PST -10
SCE(config)#interface mng 0/1
SCE(config if)#speed 100
SCE(config if)#exit
SCE(config)#interface Linecard 0
SCE(config if)#link mode forwarding
SCE(config if)#end
sce>
```

Prompt Indications

The on-screen prompt indicates your authorization level, your command hierarchy level, and the assigned host name. The structure of the prompt is:

<hostname (mode-indication) level-indication>

Authorization levels are indicated as listed in Table 2-3.

Table 2-3 Prompt Indications: Authorization Levels

This prompt...	Indicates this...
>	User and Viewer levels
#	Admin level
#>	Root level

Command hierarchy levels are indicated as listed in Table 2-4.

Table 2-4 Prompt Indications: Command Hierarchy Levels

This command hierarchy...	Is indicated as...
User Exec	SCE>
Privileged Exec	sce#
Global Configuration	SCE (config)#
Interface Configuration	SCE (config if)#
Interface Range Configuration	SCE (config if range)#
Line Configuration	SCE (config-line)#

Example:

The prompt SCE1(config if)# indicates:

- The name of the Cisco SCE platform is SCE1
- The current CLI mode is Interface configuration mode
- The user has Admin authorization level

Navigating Between Authorization Levels and Command Modes

The authorization levels and command modes function together under one hierarchy. The User and Viewer authorization levels have only a single command mode. When you enter either the Admin authorization level or Root authorization level (these levels function in parallel), you enter the Privileged EXEC command mode. From this command mode you can access the following command modes:

- User EXEC authorization level
- Viewer authorization level
- Privileged EXEC command mode (in either the Admin authorization level or the Root authorization level)

- Global Configuration command mode

  From this command mode, the following Interface Command Modes can be accessed:

  - MNG Interface Configuration (management interface)
  - Linecard Interface Configuration
  - TenGigabitEthernet Interface Configuration (10GBE traffic interfaces)
  - Interface Range Configuration (range of traffic interfaces)
  - Line Configuration

Table 2-5 summarizes how to navigate the CLI command hierarchy.

Table 2-5 CLI Command Navigation Hierarchy

Authorization Level or Command Mode	Use this command to access	Use this command to exit
User Exec	Not applicable	**logout** or **exit** (exits the current CLI session)
Viewer	**enable 5**	**disable**
Privileged Exec	**enable 10** or **enable 15** (accesses root level)	**disable**
Global Configuration	**configure**	**exit** (exits to Privileged Exec) **end** (exits to User Exec)
MNG Interface Configuration (management)	**interface mng (0/1 | 0/2)** OR **interface gigabitethernet (1/1 | 1/2)**	**exit** (exits to Global Configuration) **end** (exits to User Exec)
Linecard Interface Configuration	**interface linecard 0**	**exit** (exits to Global Configuration) **end** (exits to User Exec)
TenGigabitEthernet Interface Configuration (10GBE traffic)	**interface tengigabitethernet** *3/<bay-number (0-3)>/0* OR **interface range tengigabitethernet** *3/<bay-range (any range between 0 and 3)>/0*	**exit** (exits to Global Configuration) **end** (exits to User Exec)
Line Configuration	**line vty 0**	**exit** (exits to Global Configuration) **end** (exits to User Exec)

The do Command: Executing Commands Without Exiting

When you are in either the global configuration mode or any of the interface configuration modes, it is possible to execute an EXEC mode command (such as a **show** command) or a privileged EXEC (such as **show running-config**) without exiting to the relevant command mode. Use the **do** command for this purpose.

How to execute an exec mode command from a configuration command mode

Step 1 At the SCE(config)# (or SCE(config if)# or SCE(config-line)#) prompt, type **do** <*command*> and press **Enter**.

The specified command executes without exiting to the appropriate exec command mode.

The following example shows how to display the running configuration while in interface configuration mode.

```
SCE(config if#) do show running-config
```

CLI Help Features

CLI provides context sensitive help. Two types of context sensitive help are supported:

- Partial Help, page 2-9
- Argument Help, page 2-9

Partial Help

To obtain a list of commands that begin with a particular character string, enter the abbreviated command entry immediately followed by a question mark (?). This form of help is called partial help, because it lists only the keywords or arguments that begin with the abbreviation you entered.

Example:

The following example illustrates how typing **c?** displays all available arguments that start with the letter c.

```
SCE(config)#snmp-server c?
Communitycontact
SCE(config)#snmp-server c
```

Argument Help

To obtain a list of keywords or parameters associated with a command, type a question mark (?) in place of a keyword or parameter on the command line.

Note that if **<Enter>** is acceptable input, the symbol <cr> represents the **Enter** key.

Example:

The following example illustrates how to get a list of all arguments or keywords expected after the command **snmp-server**.

```
SCE(config)#snmp-server?
community Define community string
contact Set system contact
enable Enable the SNMP agent
host Set traps destination
interface Set interface parameters
SCE(config)# snmp-server
```

When asking for help on particular parameter, the system informs you of the type of data that is an accepted legal value. The types of parameters supported are:

STRING	When a String is expected, you can enter any set of characters or digits. If the string has a space as one of its characters, use double-quote (“) marks to enclose the string.
DECIMAL	Any decimal number. Positive number is assumed, for negative numbers use the “–” symbol.
HEX	A hexadecimal number; must start with either 0x or 0X.

Example:

The following example illustrates the use of **?** to get help on commands syntax. In this example, you can enter either the word **running-config**, or any name of a file, after the word **copy**.

```
SCE#copy?
running-config Copy running configuration file
startup-config Backup the startup-config to a specified destination
STRING Source file
SCE#copy
```

Table 2-6 summarizes the CLI help features.

Table 2-6 Getting Help

Command	Purpose
?	List all commands available for a particular command mode

<*abbreviated-command-entry*>**?**
 	
Example:
`c?`
`calendar cd clear clock configure`
`copy copy-passive`
	Obtain a list of commands that begin with a particular character string.
	(Do not leave a space between the command and question mark.)

<*abbreviated-command-entry*>**<Tab>**
 	
Example:
`en <Tab>`
`enable`
	Complete a partial command name.

<*command*>**?**	List the keywords associated with the specified command.

<*command keyword*> **?**
 	
Example:
`show ?`
`access-lists Show all access-lists`
	List the arguments associated with the specified keyword.
	Leave a space between the keyword and question mark

Navigational and Shortcut Features

- Command History, page 2-11
- Keyboard Shortcuts, page 2-11
- Auto-Completion, page 2-12
- FTP User Name and Password, page 2-13

Command History

CLI maintains a history buffer of the most recent commands you used in the current CLI session for quick retrieval. Using the keyboard, you can navigate through your last commands, one by one, or all commands that start with a given prefix. By default, the system saves the last 30 commands you typed. You can change the number of commands remembered using the **history size** command.

To use the history functions, use the keys shown in Table 2-7.

Table 2-7 Keyboard Shortcuts for History Functions

Arrow	Shortcut	Description
Up arrow	Ctrl-P	Move cursor to the previous command with the same prefix.
Down arrow	Ctrl-N	Moves the cursor to the next command with the same prefix as original.
	Ctrl-L Ctrl-R	Re-display the current command line.

Keyboard Shortcuts

The Cisco SCE platform has several keyboard shortcuts that make it easier to navigate and use the system. Table 2-8 shows the keyboard shortcuts available.

You can get a display the keyboard shortcuts at any time by typing help bindings.

Table 2-8 Keyboard Shortcuts

Description	Shortcut key
Navigational shortcuts	
Move cursor one character to the right.	CTRL-F /->
Move cursor one character to the left.	CTRL-B /<-
Move cursor one word to the right (forward).	ESC-F
Move cursor one word to the left (backward).	ESC-B
Move cursor to the start of the line.	CTRL-A
Move cursor to the end of the line.	CTRL-E
Editing shortcuts	
Delete the character where the cursor is located.	CTRL-D

Table 2-8 Keyboard Shortcuts (continued)

Description	Shortcut key
Delete from the cursor position to the end of the word.	ESC-d
Delete the character before the current location of the cursor.	Backspace
Delete the character before the current location of the cursor.	CTRL-H
Deletes from the cursor position to the end of the line	CTRL-K
Deletes all characters from the cursor to the beginning of the line	CTRL-U
Delete the word to the left of the cursor.	CTRL-W
Recall the last item deleted.	CTRL-Y
Completes the word when there is only one possible completion.	<Tab>
Completes the word when there is only one possible completion. (Same functionality as **<Tab>**.)	CTRL-I

Auto-Completion

The CLI interface features tab completion. When you type in the first letters of a command and press **<Tab>**, the system automatically fills in the rest of the command or keyword. This feature works only when there is one command that could be possible using the starting letters.

Example:

The letters **snm** followed by **<Tab>** will be completed to the command **snmp-server**.

```
SCE(config)#snm <Tab>
SCE(config)#snmp-server
```

If you press **<Enter>** instead of **<Tab>**, and there is no ambiguity, the system actually carries out the command that is the result of the auto-completion.

Example: 1

The following example displays how the system completes a partial (unique) command for the **enable** command. The system carries out the command using the default authorization level (10) when you press **<Enter>**.

```
SCE>en <Enter>
Password:
sce#
```

Example: 2

The following example illustrates how to use the completion feature with a non-default value for the argument. In this example, the **enable** command is completed using the specified value (15) for the authorization level.

```
SCE>en 15 <Enter>
Password:
sce#
```

FTP User Name and Password

CLI enables saving FTP user name and password to be used in FTP operations—download and upload, per session.

These settings are effective during the current CLI session.

The following example illustrates how to set FTP password and user name and the use in these settings for getting a file named *config.tmp* from a remote station using FTP protocol.

```
sce#ip FTP password pw123
sce#ip FTP username user1
sce#copy ftp://@10.10.10.10/h:/config.tmp myconf.txt connecting 10.10.10.10 (user name
user1 password pw123) to retrieve config.tmp
sce#
```

Managing Command Output

- Scrolling the Screen Display, page 2-14
- Filtering Command Output, page 2-14
- Redirecting Command Output to a File, page 2-15

Some commands, such as many show commands, may have many lines of output. There are several ways of managing the command output:

- Scrolling options—When the command output is too large to be displayed all at once, you can control whether the display scrolls line by line or refreshes the entire screen.
- Filtering options—You can filter the output so that output lines are displayed only if they include or exclude a specified expression.
- Redirecting to a file—You can send the output to a specified file.

Note that by default, the **show** commands act the same as the **more** commands; that is, the output is displayed interactively a single screen at a time. Use the **no more** command to disable this feature so that **show** commands display the complete output all at one time.

Scrolling the Screen Display

The output of some **show** and **dir** commands is quite lengthy and cannot all be displayed on the screen at one time. Commands with many lines of output are displayed in chunks of 24 lines. You can choose to scroll the display line by line or refresh the entire screen. At the prompt after any line, you can type one of the following keys for the desired action:

- **<Enter>**—Show one more line
- **<Space>**—Show 24 more lines (a new chunk)
- **<g>**—Stop prompting for more
- **<?>**—Display a help string showing possible options
- Any other key—Quit showing the file

Filtering Command Output

You can filter the output of certain commands, such as **show**, **more**, and **dir**, so that output lines are displayed only if they include or exclude a specified expression. The filtering options are as follows:

- include—Shows all lines that include the specified text.
- exclude—Does not show any lines that include the specified text.
- begin—Finds the first line that includes the specified text, and shows all lines starting from that line. All previous lines are excluded.

The syntax of filtered commands is as follows:

- *command* **| include** *expression*
- *command* **| exclude** *expression*
- *command* **| begin** *expression*

Following is an example of how to filter the **show version** command to display only the last part of the output, beginning with the version information.

```
sce# show version | begin revision
```

Redirecting Command Output to a File

You can redirect the output of commands, such as **show**, **more**, and **dir**, to a file. When writing the output of these commands to a file, you can specify either of the following options:

- redirect—The new output of the command will overwrite the existing contents of the file.
- append—The new output of the command will be appended to the existing contents of the file.

The syntax of redirection commands is as follows:

- *command* | **redirect** *filename*
- *command* | **append** *filename*

This example illustrates how to do the following:

- Filter the **more** command to display only the gold package subscribers from a csv subscriber file.
- Redirect that output to a file named *current_gold_subscribers*. The output should not overwrite existing entries in the file, but should be appended to the end of the file.

```
sce# more subscribers_10.10.2008 include gold | append current_gold_subscribers
```

Creating a CLI Script

The CLI scripts feature allows you to record several CLI commands together as a script and play it back. This is useful for saving repeatable sequence of commands, such as software upgrade. For example, if you are configuring a group of Cisco SCE platforms and you want to run the same configuration commands on each platform, you could create a script on one platform and run it on all the other Cisco SCE platforms. The available script commands are:

- **script capture**
- **script stop**
- **script print**
- **script run**

Step 1 At the sce# prompt, type **script capture** ***filename.scr*** where *filename.scr* is the name of the script, with a scr file extension.

Step 2 Perform the actions you want to be included in the script.

Step 3 Type **script stop**.

The system saves the script.

The following is an example of recording a script for upgrading software.

```
sce#script capture upgrade.scr
sce#configure
SCE(config)#boot system new.pkg Verifying package file...
Package file verified OK.
SCE(config)#exit
sce#copy running-config startup-config
Writing general configuration file to temporary location...
Extracting files from '//apps/data/scos/images/new.pkg'...
Verifying package file...
Package file verified OK.
Device '//apps/data/scos/' has 81154048 bytes free, 21447973 bytes are needed for
extraction, all is well.
Extracting files to temp locations...
Renaming temp files...
Extracted OK.
Backing-up general configuration file...
Copy temporary file to final location...
sce#script stop
sce#
```



Basic Cisco SCE 8000 Platform Operations

Revised: September 23, 2015

Introduction

This chapter describes how to start up the Cisco SCE 8000 platform, reboot, and shutdown. It also describes how to manage configurations.

- Starting the Cisco SCE 8000 Platform, page 3-2
- Managing Configurations, page 3-5
- Example for Displaying the Cisco SCE Platform Version Information, page 3-10
- How to Display the Cisco SCE Platform Inventory, page 3-13
- How to Display the System Uptime, page 3-17
- Configuring the System Mode, page 3-17
- Configuring the IPv6 Prefix Length, page 3-18
- Rebooting and Shutting Down the Cisco SCE Platform, page 3-23

Starting the Cisco SCE 8000 Platform

The procedures for starting the Cisco SCE 8000 platform are explained in the following sections:

- Checking Conditions Prior to System Startup, page 3-2
- Starting the System and Observing Initial Conditions, page 3-2
- Final Tests, page 3-3

Checking Conditions Prior to System Startup

Check the following conditions before you start your Cisco SCE 8000 platform:

- Both power supply units are installed and connected. (If only one power supply is connected it will put the box in warning state.)
- First-time startup at installation:
  - Cisco SCE 8000 platform connected to local console (CON port)
  - The console terminal is turned on and properly configured
- Subsequent startups
  - Line interfaces are properly cabled (optional)
  - Cisco SCE 8000 platform is connected to at least one of the following types of management stations:
  - Direct connection to local console (CON port)
  - Remote management station via the LAN (Mng port)

Starting the System and Observing Initial Conditions

After installing your Cisco SCE 8000 platform and connecting cables, complete the following steps to start the Cisco SCE 8000 platform:

Step 1 Make sure the power cables are connected to the Cisco SCE 8000 platform.

Step 2 Plug the AC power supply cables into the AC power source, or make sure the circuit breakers at the DC panels are turned to the on position. Turn on the switches on both power supplies.

Step 3 Listen for the fans; you should immediately hear them operating.

Step 4 During the boot process, observe the following LEDs on the Cisco SCE 8000-SCM-E:

- The Power LEDs should be green.
- Optical Bypass LED should be green while the Cisco SCE 8000 is in bypass and unlit when the optical bypass is turned off.
- The Status LED should be a constant amber while booting. After a successful boot, the Status LED is steady green.

Note It takes a several minutes for the Cisco SCE 8000 to boot and for the status LED to change from amber to green.

Final Tests

The procedures for performing the final tests to verify that the Cisco SCE 8000 is functioning properly are explained in the following sections:

- How to Verify Operational Status, page 3-3
- How to View the User Log Counters, page 3-3

How to Verify Operational Status

After all the ports are connected, verify that the Cisco SCE 8000 is not in a Warning state.

Step 1 On the front panel of the Service Control module, examine the Status LED; it should be green.

Step 2 To display the operation status of the system, at the Cisco SCE8000# prompt, type **show system operation-status** and press **Enter**.

A message displaying the operation status of the system appears. If the system is operating in order, the following message appears:

```
System Operation status is Operational.
```

If the Status LED is red or flashing amber, the following message appears:

```
System Operation status is Warning
Description:
1. Power Supply problem
2. Line feed problem
3. Amount of External bypass devices detected is lower than expected amount
```

How to View the User Log Counters

View the user log for errors that occurred during the installation process.

At the SCE# prompt, type:

Command	Purpose
show logger device user-file-log counters	Shows the user log counters.

Examples for Viewing the User Log Counters

The following example shows the current User-File-Log device counters.

```
SCE#show logger device user-file-log counters
Logger device User-File-Log counters:
Total info messages: 1
Total warning messages: 0
Total error messages: 0
Total fatal messages: 0
```

If there are “Total error messages” or “Total fatal messages”, use the **show logger device user-file-log** command to display details about the errors.

Managing Configurations

- Viewing Configurations, page 3-5
- How to Save or Change the Configuration Settings, page 3-6
- Restoring a Previous Configuration, page 3-8

Viewing Configurations

When you enter configuration commands, it immediately affects the Cisco SCE platform operation and configuration. This configuration, referred to as the running-config, is saved in the Cisco SCE platform volatile memory and is effective while the Cisco SCE platform is up. After reboot, the Cisco SCE platform loads the startup-config, which includes the non-default configuration that was saved by the user, into the running-config.

The Cisco SCE platform provides commands for:

- Viewing the running configuration with only user-configured (non-default) values: **show running-config**
- Viewing the running configuration with all the Cisco SCE platform running configuration values, whether default or not: **show running-config all-data**
- Viewing the startup configuration: **show startup-config**

After configuring the Cisco SCE platform, you may query for the running configuration using the command **show running-config**.

Step 1 At the Cisco SCE8000# prompt, type show running-config.

The system shows the running configuration.

```
SCE8000#>show running-config
#This is a general configuration file (running-config).
#Created on 12:06:13 UTC SUN May 11 2008
#cli-type 1
#version 1
no management-agent notifications notification-list
1417,1418,804,815,1404,1405,1406,1407,1408,400
no management-agent notifications notification-list
402,421,440,441,444,445,446,450,437,457
no management-agent notifications notification-list 3593,3594,3595,10040
snmp-server community "public" ro
RDR-formatter forwarding-mode multicast
RDR-formatter destination 10.56.96.26 port 33000 category number 1 priority 100
RDR-formatter destination 10.56.96.26 port 33000 category number 2 priority 100
RDR-formatter destination 10.56.96.26 port 33000 category number 3 priority 100
RDR-formatter destination 10.56.96.26 port 33000 category number 4 priority 100
interface LineCard 0
connection-mode inline on-failure external-bypass
no silent
no shutdown
attack-filter subscriber-notification ports 80
replace spare-memory code bytes 3145728
interface GigabitEthernet 1/1
ip address 10.56.96.46 255.255.252.0
interface TenGigabitEthernet 3/0/0
bandwidth 10000000 burst-size 50000
global-controller 0 name "Default Global Controller"
interface TenGigabitEthernet 3/1/0
```

```
bandwidth 10000000 burst-size 50000
global-controller 0 name "Default Global Controller"
interface TenGigabitEthernet 3/2/0
bandwidth 10000000 burst-size 50000
global-controller 0 name "Default Global Controller"
interface TenGigabitEthernet 3/3/0
bandwidth 10000000 burst-size 50000
global-controller 0 name "Default Global Controller"

exit
ip default-gateway 10.56.96.1
line vty 0 4
exit
management-agent property "com.pcube.management.framework.install.activation.operation"
"Install"
management-agent property "com.pcube.management.framework.install.activated.package" "SCA
BB"
management-agent property "com.pcube.management.framework.install.activated.version"
"3.1.6 build 79"
management-agent property "com.pcube.management.framework.install.activation.date" "Sun
May 11 08:44:04 GMT+00:00 2008"
flow-filter partition name "ignore_filter" first-rule 4 num-rules 32
flow-filter partition name "udpPortsToOpenBySw" first-rule 40 num-rules 21
```

How to Save or Change the Configuration Settings

When you make changes to the current running configuration and you want those changes to continue to be in effect when the system restarts, you must save the changes before leaving the management session. You do that by saving the running configuration to the startup configuration file.

The Cisco SCE platform provides multiple interfaces for the purpose of configuration and management. All interfaces supply an API to the same database of the Cisco SCE platform and any configuration made through one interface is reflected through all interfaces. Furthermore, when saving the running configuration to the startup configuration from any management interface, all configuration settings are saved regardless of the management interface used to set the configuration.

For backup purposes, the old startup-config file is saved under the directory: `/system/prevconf`. Refer to "Restoring a Previous Configuration" section on page 3-8 for an explanation on how to restore a previous configuration.

To remove a configuration command from the running-config, use the **no** form of the command.

Step 1 At the SCE# prompt, type **show running-config** to view the running configuration.

The running configuration is displayed.

Step 2 Check the displayed configuration to make sure that it is set the way you want. If not, make the changes you want before saving.

Step 3 Type **copy running-config startup-config**.

The system saves all running configuration information to the configuration file, which is used when the system reboots.

The configuration file holds all information that is different from the system default in a file called *config.tx1* located in the directory: */system*.

Example for Saving or Changing the Configuration Settings

The following example shows how to save the running configuration file (first displaying the file to review the settings).

```
SCE# show running-config
#This is a general configuration file (running-config).
#Created on 12:06:13 UTC SUN May 11 2008
#cli-type 1
#version 1
no management-agent notifications notification-list
1417,1418,804,815,1404,1405,1406,1407,1408,400
no management-agent notifications notification-list
402,421,440,441,444,445,446,450,437,457
no management-agent notifications notification-list 3593,3594,3595,10040
snmp-server community "public" ro
RDR-formatter forwarding-mode multicast
RDR-formatter destination 10.56.96.26 port 33000 category number 1 priority 100
RDR-formatter destination 10.56.96.26 port 33000 category number 2 priority 100
RDR-formatter destination 10.56.96.26 port 33000 category number 3 priority 100
RDR-formatter destination 10.56.96.26 port 33000 category number 4 priority 100
interface LineCard 0
connection-mode inline on-failure external-bypass
no silent
no shutdown
attack-filter subscriber-notification ports 80
replace spare-memory code bytes 3145728
interface GigabitEthernet 1/1
ip address 10.56.96.46 255.255.252.0
interface TenGigabitEthernet 3/0/0
bandwidth 10000000 burst-size 50000
global-controller 0 name "Default Global Controller"
interface TenGigabitEthernet 3/1/0
bandwidth 10000000 burst-size 50000
global-controller 0 name "Default Global Controller"
interface TenGigabitEthernet 3/2/0
bandwidth 10000000 burst-size 50000
global-controller 0 name "Default Global Controller"
interface TenGigabitEthernet 3/3/0
bandwidth 10000000 burst-size 50000
global-controller 0 name "Default Global Controller"

exit
ip default-gateway 10.56.96.1
line vty 0 4
exit
management-agent property "com.pcube.management.framework.install.activation.operation"
"Install"
management-agent property "com.pcube.management.framework.install.activated.package" "SCA
BB"
management-agent property "com.pcube.management.framework.install.activated.version"
"3.1.6 build 79"
management-agent property "com.pcube.management.framework.install.activation.date" "Sun
May 11 08:44:04 GMT+00:00 2008"
flow-filter partition name "ignore_filter" first-rule 4 num-rules 32
flow-filter partition name "udpPortsToOpenBySw" first-rule 40 num-rules 21
SCE#copy running-config startup-config
Writing general configuration file to temporary location...
Backing-up general configuration file...
Copy temporary file to final location...
SCE#
```

Tip To remove a configuration command from the running-config, use the **no** form of the command.

The following example illustrates how to remove all DNS settings from the running configuration.

```
SCE(config)#no ip name-server
```

Restoring a Previous Configuration

When you save a new configuration, the system automatically backs up the old configuration in the directory `/system/prevconf/`. Up to nine versions of the startup configuration file are saved, namely `config.tx1-config.tx9`, where `config.tx1` is the most recently saved file.

You can view the old startup configuration files using the CLI command **more**.

Restoring a previous startup configuration means renaming the file so it overwrites the startup configuration (`config.txt`) file.

Step 1 At the SCE# prompt, type **more /system/prevconf/config.tx1** to view the configuration file.

The system displays the configuration information stored in the file.

Step 2 Read the configuration information to make sure it is the configuration you want to restore.

Note that you cannot undo the configuration restore command.

Step 3 Type **copy /system/config.tx1 /system/config.txt**.

The system sets the startup configuration to the configuration from *config.tx1*.

Example for Restoring a Previous Configuration

The following example displays a saved configuration file and then restores the file to overwrite the current configuration.

```
SCE#more /system/prevconf/config.tx1
#This is a general configuration file (running-config).
#Created on 12:07:41 UTC SUN May 11 2008
#cli-type 1
#version 1
no management-agent notifications notification-list
1417,1418,804,815,1404,1405,1406,1407,1408,400
no management-agent notifications notification-list
402,421,440,441,444,445,446,450,437,457
no management-agent notifications notification-list 3593,3594,3595,10040
snmp-server community "public" ro
RDR-formatter forwarding-mode multicast
RDR-formatter destination 10.56.96.26 port 33000 category number 1 priority 100
RDR-formatter destination 10.56.96.26 port 33000 category number 2 priority 100
RDR-formatter destination 10.56.96.26 port 33000 category number 3 priority 100
RDR-formatter destination 10.56.96.26 port 33000 category number 4 priority 100
interface LineCard 0
connection-mode inline on-failure external-bypass
no silent
no shutdown
attack-filter subscriber-notification ports 80
replace spare-memory code bytes 3145728
```

```
interface GigabitEthernet 1/1
ip address 10.56.96.46 255.255.252.0
interface TenGigabitEthernet 3/0/0
bandwidth 10000000 burst-size 50000
global-controller 0 name "Default Global Controller"
interface TenGigabitEthernet 3/1/0
bandwidth 10000000 burst-size 50000
global-controller 0 name "Default Global Controller"
interface TenGigabitEthernet 3/2/0
bandwidth 10000000 burst-size 50000
global-controller 0 name "Default Global Controller"
interface TenGigabitEthernet 3/3/0
bandwidth 10000000 burst-size 50000
global-controller 0 name "Default Global Controller"

exit
ip default-gateway 10.56.96.1
line vty 0 4
exit
management-agent property "com.pcube.management.framework.install.activation.operation"
"Install"
management-agent property "com.pcube.management.framework.install.activated.package" "SCA
BB"
management-agent property "com.pcube.management.framework.install.activated.version"
"3.1.6 build 79"
management-agent property "com.pcube.management.framework.install.activation.date" "Sun
May 11 08:44:04 GMT+00:00 2008"
flow-filter partition name "ignore_filter" first-rule 4 num-rules 32
flow-filter partition name "udpPortsToOpenBySw" first-rule 40 num-rules 21
SCE#copy /system/config.tx1 /system/config.txt
```

How to Display the Cisco SCE Platform Version Information

Use this command to display global static information on the Cisco SCE platform, such as software and hardware version, image build time, system uptime, last open packages names and information on the SLI application assigned.

From the SCE> prompt, type:

Command	Purpose
show version	Displays the Cisco SCE platform version information.

Example for Displaying the Cisco SCE Platform Version Information

The following example shows how to display the Cisco SCE platform version information:

```
SCE> show version

System version: Version 3.1.6S Build 279
Build time: Jun 10 2008, 19:27:47 (Change-list 335658)
Software version is: Version 3.1.6S Build 279
Hardware information is:
----------------
Firmware
----------------
kernel : [kernel] 1.0.0/5 (inactive: [kernel] 1.0.0/5)
u-boot : [uboot] 1.0.0/6 (field: [uboot] 0.8.1/13)
select : [ubs-cf1] 1.0.0/5 (secondary: [ubs-cf1] 1.0.0/5)
----------------
Slot 1: SCM-8000
----------------
serial-num : CAT1202G07D
part-num : 73-10598-01 38
cpld : 0x8162
vtpld : 0xc001
summit-0 : 0x10008
summit-1 : 0x10008
dpt/tx : 0x4837
cls/ff : 0x2047
cls flow cap: 33554432

----------------
TVR
----------------
#cpus : 1
cpu SVR: 0x80900120
cpu PVR : 0x80040202
cpu freq : 1000MHz
cpu (eeprom): 2.1, 1000MHz
cpld : 0xa1b7
cpld-ufm : 0xa803
summit : 0x10007
cf : Model=SMART CF, FwRev=0x20060811, Size=4062240KB
----------------
CFC-0
----------------
board type : P2
#cpus : 3
```

```
cpu-0 SVR : 0x80900121
cpu-0 PVR : 0x80040202
cpu-0 freq : 1500MHz
cpu-1 SVR : 0x80900121
cpu-1 PVR : 0x80040202
cpu-1 freq : 1500MHz
cpu-2 SVR : 0x80900121
cpu-2 PVR : 0x80040202
cpu-2 freq : 1500MHz
cpu (eeprom): 2.1, 1500MHz
cpld-0 : 0xb20e
cpld-1 : 0xb20e
cpld-2 : 0xb20e
cpld-0-ufm : 0xb803
cpld-1-ufm : 0xb803
cpld-2-ufm : 0xb803
summit-0 : 0x1000a
summit-1 : 0x1000a
fc : 0x1044
----------------
CFC-1
----------------
board type : P2
#cpus : 3
cpu-0 SVR : 0x80900121
cpu-0 PVR : 0x80040202
cpu-0 freq : 1500MHz
cpu-1 SVR : 0x80900121
cpu-1 PVR : 0x80040202
cpu-1 freq : 1500MHz
cpu-2 SVR : 0x80900121
cpu-2 PVR : 0x80040202
cpu-2 freq : 1500MHz
cpu (eeprom): 2.1, 1500MHz
cpld-0 : 0xb20e
cpld-1 : 0xb20e
cpld-2 : 0xb20e
cpld-0-ufm : 0xb803
cpld-1-ufm : 0xb803
cpld-2-ufm : 0xb803
summit-0 : 0x1000a
summit-1 : 0x1000a
fc : 0x1044
----------------
Slot 3: SIP-8000
----------------
serial-num : CAT1204G01H
part-num : 73-10947-01
cpld : 0x9162
summit-0 : 0x10006
summit-1 : 0x10006
dpt-0 : 0x3033
dpt-1 : 0x3033
spa[0] : SPA-1X10GE-L-V2
spa[1] : SPA-1XTENGE-XFP
spa[2] : SPA-1X10GE-L-V2
spa[3] : SPA-1XTENGE-XFP
----------------
SCE8000 Chassis
----------------
product-num : CISCO7604
serial-num : FOX10420BKZ
```

```
part-num : 73-9789-02
part-rev : A0
vid : V01
Part number: 73-10598-01 38
Revision:
Software revision:
LineCard S/N : CAT1202G07D
Power Supply type: AC
SML Application information is:
No application is configured.
Logger status: Enabled

Platform: SCE8000 - 4x10GBE
Management agent interface version: SCE Agent 3.1.6 Build 134
Software package file: ftp://ftpserver/simba.pkg
SCE8000 uptime is 9 minutes, 54 seconds
```

How to Display the Cisco SCE Platform Inventory

Unique Device Identification (UDI) is a Cisco baseline feature that is supported by all Cisco platforms. This feature allows network administrators to remotely manage the assets in their network by tracing specific devices through either CLI or SNMP. The user can display inventory information for a remote device via either:

- Entity MIB (see ENTITY-MIB)
- CLI **show inventory** command

  This command displays the UDIs only for field replaceable units (FRU).
- CLI **show inventory raw** command.

  This command displays all UDIs on the Cisco SCE 8000 platform.

The **show inventory** CLI commands display the following information:

- *Device name*
- *Description*
- *Product identifier*
- *Version identifier*
- *Serial number*

From the SCE> prompt, type:

Command	Purpose
show inventory [raw]	Displays the Cisco SCE platform inventory.

Examples for Displaying the Cisco SCE Platform Inventory

- Displaying the Cisco SCE Platform Inventory: FRUs Only, page 3-13
- Displaying the Complete Cisco SCE Platform Inventory, page 3-14

Displaying the Cisco SCE Platform Inventory: FRUs Only

The following example shows how to display the inventory (UDIs) for the FRUs only:

```
SCE> show inventory

NAME: "SCE8000 Chassis", DESCR: "CISCO7604"
PID: CISCO7604          , VID: V0 , SN: FOX105108X5
NAME: "SCE8000 Service Control Module (SCM) in slot 1", DESCR: "SCE8000-SCM-E"
PID: SCE8000-SCM-E      , VID: V0 , SN: CAT1122584N
NAME: "SCE8000 SPA Interface Processor (SIP) in slot 3", DESCR: "SCE8000-SIP"
PID: SCE8000-SIP        , VID: V0 , SN: CAT1150G07F

NAME: "SPA-1X10GE-L-V2", DESCR: "SPA-1X10GE-L-V2"
PID: SPA-1X10GE-L-V2    , VID: V02, SN: JAE11517RMR

NAME: "SPA-1X10GE-L-V2", DESCR: "SPA-1X10GE-L-V2"
PID: SPA-1X10GE-L-V2    , VID: V02, SN: JAE11496E1P

NAME: "SPA-1X10GE-L-V2", DESCR: "SPA-1X10GE-L-V2"
```

```
PID: SPA-1X10GE-L-V2   , VID: V02, SN: JAE11517RIO

NAME: "SPA-1X10GE-L-V2", DESCR: "SPA-1X10GE-L-V2"
PID: SPA-1X10GE-L-V2   , VID: V02, SN: JAE115295HH

NAME: "SCE8000 FAN 1", DESCR: "FAN-MOD-4HS"
PID: FAN-MOD-4HS       , VID: V0 , SN: DCH11013744

NAME: "SCE8000 AC or DC power supply 0", DESCR: "PWR-2700-AC/4"
PID: PWR-2700-AC/4     , VID: V0 , SN: APQ105000MV

NAME: "SCE8000 AC or DC power supply 1", DESCR: "PWR-2700-AC/4"
PID: PWR-2700-AC/4     , VID: V0 , SN: APQ105000MV

NAME: "XFP-10GLR-OC192SR ", DESCR: "XFP-10GLR-OC192SR "
PID: XFP-10GLR-OC192SR , VID: V02, SN: AGA1142N4B7

NAME: "XFP-10GLR-OC192SR ", DESCR: "XFP-10GLR-OC192SR "
PID: XFP-10GLR-OC192SR , VID: V02, SN: AGA1142N4AL

NAME: "XFP-10GLR-OC192SR ", DESCR: "XFP-10GLR-OC192SR "
PID: XFP-10GLR-OC192SR , VID: V02, SN: AGA1141N43R

NAME: "XFP-10GLR-OC192SR ", DESCR: "XFP-10GLR-OC192SR "
PID: XFP-10GLR-OC192SR , VID: V02, SN: AGA1143N4JN
```

Displaying the Complete Cisco SCE Platform Inventory

The following example shows how to display the complete inventory (UDIs) of the SCE platform:

```
SCE> show inventory raw

"SCE8000 Chassis", DESCR: "CISCO7604"
PID: CISCO7604         , VID: V01, SN: FOX105108X5


NAME: "SCE8000 Physical Slot 1", DESCR: "Container SCE8000 Service Control Module (SCM)
slot"
PID: ""                , VID: "" , SN: ""

NAME: "SCE8000 Physical Slot 2", DESCR: "Container SCE8000 Service Control Module (SCM)
slot"
PID: ""                , VID: "" , SN: ""

NAME: "SCE8000 Physical Slot 3", DESCR: "Container SCE8000 SPA Interface Processor (SIP)
slot"
PID: ""                , VID: "" , SN: ""

NAME: "SCE8000 Physical Slot 4", DESCR: "Container SCE8000 Optical Bypass slot"
PID: ""                , VID: "" , SN: ""

NAME: "SCE8000 Fan Module", DESCR: "Container SCE8000 Fan Module"
PID: ""                , VID: "" , SN: ""

NAME: "SCE8000 AC and DC power supply", DESCR: "Container SCE8000 AC and DC power supply"
PID: ""                , VID: "" , SN: ""

NAME: "SCE8000 Link", DESCR: "Container SCE8000 Link"
PID: ""                , VID: "" , SN: ""

NAME: "SCE8000 Backplane", DESCR: "Container SCE8000 Backplane "
PID: ""                , VID: "" , SN: ""
```

```
NAME: "SCE8000 Service Control Module (SCM) in slot 1", DESCR: "SCE8000-SCM-E"
PID: SCE8000-SCM-E     , VID: V01, SN: CAT1122584N

NAME: "SCE8000 SPA Interface Processor (SIP) in slot 3", DESCR: "SCE8000-SIP"
PID: SCE8000-SIP       , VID: V01, SN: CAT1150G07F

NAME: "SCE8000 Link 0", DESCR: "SCE8000 Link"
PID: ""                , VID: "" , SN: ""

NAME: "SCE8000 Link 1", DESCR: "SCE8000 Link"
PID: ""                , VID: "" , SN: ""

NAME: "SCE8000 SIP bay 3/0", DESCR: "SCE8000 SIP bay"
PID: ""                , VID: "" , SN: ""

NAME: "SCE8000 SIP bay 3/1", DESCR: "SCE8000 SIP bay"
PID: ""                , VID: "" , SN: ""

NAME: "SCE8000 SIP bay 3/2", DESCR: "SCE8000 SIP bay"
PID: ""                , VID: "" , SN: ""

NAME: "SCE8000 SIP bay 3/3", DESCR: "SCE8000 SIP bay"
PID: ""                , VID: "" , SN: ""

NAME: "SCE8000 SPA module 3/0", DESCR: "SPA-1X10GE-L-V2"
PID: SPA-1X10GE-L-V2   , VID: V02, SN: JAE11517RMR

NAME: "SCE8000 SPA module 3/1", DESCR: "SPA-1X10GE-L-V2"
PID: SPA-1X10GE-L-V2   , VID: V02, SN: JAE11496E1P

NAME: "SCE8000 SPA module 3/2", DESCR: "SPA-1X10GE-L-V2"
PID: SPA-1X10GE-L-V2   , VID: V02, SN: JAE11517RIO

NAME: "SCE8000 SPA module 3/3", DESCR: "SPA-1X10GE-L-V2"
PID: SPA-1X10GE-L-V2   , VID: V02, SN: JAE115295HH

NAME: "TenGigabitEthernet3/0/0", DESCR: "SCE8000 SPA port"
PID: ""                , VID: "" , SN: ""

NAME: "TenGigabitEthernet3/1/0", DESCR: "SCE8000 SPA port"
PID: ""                , VID: "" , SN: ""

NAME: "TenGigabitEthernet3/2/0", DESCR: "SCE8000 SPA port"
PID: ""                , VID: "" , SN: ""


NAME: "TenGigabitEthernet3/3/0", DESCR: "SCE8000 SPA port"
PID: ""                , VID: "" , SN: ""

NAME: "SCE8000 FAN 1", DESCR: "FAN-MOD-4HS"
PID: FAN-MOD-4HS       , VID: V01, SN: DCH11013744

NAME: "SCE8000 AC power supply 0", DESCR: "PWR-2700-AC/4"
PID: PWR-2700-AC/4     , VID: V02, SN: APQ105000MV

NAME: "SCE8000 DC power supply 1", DESCR: "PWR-2700-DC/4"
PID: PWR-2700-DC/4     , VID: V03, SN: APQ1049000S

NAME: "SCE8000 optic 3/0/0", DESCR: "XFP-10GLR-OC192SR "
PID: XFP-10GLR-OC192SR , VID: V02, SN: AGA1142N4B7

NAME: "SCE8000 optic 3/1/0", DESCR: "XFP-10GLR-OC192SR "
PID: XFP-10GLR-OC192SR , VID: V02, SN: AGA1142N4AL
```

```
NAME: "SCE8000 optic 3/2/0", DESCR: "XFP-10GLR-OC192SR "
PID: XFP-10GLR-OC192SR , VID: V02, SN: AGA1141N43R

NAME: "SCE8000 optic 3/3/0", DESCR: "XFP-10GLR-OC192SR "
PID: XFP-10GLR-OC192SR , VID: V02, SN: AGA1143N4JN

NAME: "SCE8000 traffic processor 1", DESCR: "SCE8000 traffic processor"
PID: ""                , VID: "" , SN: ""

NAME: "SCE8000 traffic processor 2", DESCR: "SCE8000 traffic processor"
PID: ""                , VID: "" , SN: ""

NAME: "SCE8000 traffic processor 3", DESCR: "SCE8000 traffic processor"
PID: ""                , VID: "" , SN: ""

NAME: "SCE8000 traffic processor 4", DESCR: "SCE8000 traffic processor"
PID: ""                , VID: "" , SN: ""

NAME: "SCE8000 traffic processor 5", DESCR: "SCE8000 traffic processor"
PID: ""                , VID: "" , SN: ""

NAME: "SCE8000 traffic processor 6", DESCR: "SCE8000 traffic processor"
PID: ""                , VID: "" , SN: ""

NAME: "SCE8000 traffic processor 7", DESCR: "SCE8000 traffic processor"
PID: ""                , VID: "" , SN: ""

NAME: "SCE8000 traffic processor 8", DESCR: "SCE8000 traffic processor"
PID: ""                , VID: "" , SN: ""

NAME: "SCE8000 traffic processor 9", DESCR: "SCE8000 traffic processor"
PID: ""                , VID: "" , SN: ""

NAME: "SCE8000 traffic processor 10", DESCR: "SCE8000 traffic processor"
PID: ""                , VID: "" , SN: ""

NAME: "SCE8000 traffic processor 11", DESCR: "SCE8000 traffic processor"
PID: ""                , VID: "" , SN: ""

NAME: "SCE8000 traffic processor 12", DESCR: "SCE8000 traffic processor"
PID: ""                , VID: "" , SN: ""
```

How to Display the System Uptime

Use this command to see how long the system has been running since the last reboot.

At the SCE>prompt, type:

Command	Purpose
show system-uptime	Displays the system uptime.

Example for Displaying the System Uptime

The following example shows how to display the system uptime of the Cisco SCE platform.

```
SCE#show system-uptime
Cisco SCE8000 uptime is 21 minutes, 37 seconds
```

Configuring the System Mode

The Cisco SCE 8000 devices operates in one the following system modes:

- IPv4 only system mode—All traffic processors handle only IPv4 traffic.
- IPv6 only system mode—All traffic processors handle only IPv6 traffic.
- Dual stack system mode—Default system mode. All traffic processors handle both IPv4 and IPv6 traffic. Dual static subscribers can either be configured statically on the device or they can be configured through Gx. The number of subscriber mappings in the device is divided between IPv4 and IPv6 based on the const -db value.

The system mode is configured by modifying the value of the **seCommonConstDb.box.ipv6MappingsPercentage** const -db.

The default value of the **seCommonConstDb.box.ipv6MappingsPercentage** const-db is 20. Dual stack mode is active if the value is between 1 and 99. IPv4 mode is active when the value is 0 and IPv6 mode when the value is 100.

To configure the IPv4 only mode, follow these steps:

Step 1 Enter the Global Configuration mode:

SCE8000# **configure**

Step 2 Configure the IPv4 only mode:

From the Global Configuration prompt, enter **debug const-db name seCommonConstDb.box.ipv6MappingsPercentage value 0** and press **Enter**.

Step 3 Copy the running configuration to the startup configuration:

SCE8000#> **copy running-config startup-config**

Step 4 Reboot the Cisco SCE 8000 device.

To configure the IPv6 only mode, follow these steps:

Step 1 Enter the Global Configuration mode:

SCE8000# **configure**

Step 2 Configure the IPv6 only mode:

From the Global Configuration prompt, enter **debug const db seCommonConstDb.box.ipv6MappingsPercentage value 100** and press **Enter**.

Step 3 Copy the running configuration to the startup configuration:

SCE8000#> **copy running-config startup-config**

Reboot the Cisco SCE 8000 device.

To configure dual stack mode with a value 30 (other than the default value 20), follow these steps:

Step 1 Enter the Global Configuration mode:

SCE8000# **configure**

Step 2 Configure the dual stack mode with value 30:

From the Global Configuration prompt, enter **debug const db seCommonConstDb.box.ipv6MappingsPercentage value 30** and press **Enter**.

Step 3 Copy the running configuration to the startup configuration:

SCE8000#> **copy running-config startup-config**

Step 4 Reboot the Cisco SCE 8000 device.

Configuring the IPv6 Prefix Length

Cisco SCE 8000 devices identifies the IPv6 subscribers based on the MSB 64 bits of the subscriber IPv6 address. Cisco SCE 8000 devices support IPv6 subscribers with a range of /32 to /64 and not less than /32.

You can configure the IPv6 prefix length using the **const -db seCommonConstDb.box. ipv6SystemPrefixLength** command. The value of IPv6 prefix length can be between 32 and 64 and the default value is 64. The changes will be effective only after you copy the running configuration to the startup configuration and restart the device.

The Cisco SCE 8000 device restricts static subscriber configuration to use prefixes greater than 31. If the Gx CCA-I reaches the device with a Framed-IPV6-Prefix with prefix delegation less than 32, it is ignored and subscriber stays as anonymous.

To configure the IPv6 Prefix Length 48, follow these steps:

Step 1 Enter the Global Configuration mode:

SCE8000# **configure**

Step 2 Configure the IPv6 prefix length:

From the Global Configuration prompt, enter **debug const-db name seCommonConstDb.box.ipv6SystemPrefixLength value 48** and press **Enter**.

Step 3 Copy the running configuration to the startup configuration:

SCE8000#> **copy running-config startup-config**

Step 4 Reboot the Cisco SCE 8000 device.

You can configure the system mode and prefix length together. The configuration will be effective with a single reload.

Based on the IPv6 prefix length, corresponding number of MSB bits (from the total 128 bits) will be used to identify subscribers.

For example, if the system prefix length is 48 for a party mapping configuration **party mapping ipv6-address** *1234:abcd:2123:abbc:0:0:1e:0* **name** *test*, only MSB 48 bits *1234:abcd:2123* is considered for identifying subscriber *test*.

Monitoring Control Processor CPU Utilization

You can monitor the CPU utilization of the control processor by displaying the actual load on the control processor. This feature provides visibility into the performance envelop of the Control processor under different management schemes.

Information regarding CPU utilization is available using any of the following methods:

- SNMP—The Cisco Process MIB shows the CPU utilization of both the control processor and the traffic processors. An SNMP walk on the cpmCPUTotalTable provides information regarding overall CPU statistics.
- CLI commands—The following admin level CLI commands can be used to monitor CPU utilization:
  - **show processes cpu**
  - **show processes cpu sorted**
  - **show snmp MIB cisco-process**
- Cisco SCE support file—The entire measured CPU utilization of the control processor, as well as a number of specific internal tasks that were marked important to track, is written to the Cisco SCE log files , which are part of the Cisco SCE support file. This data can be used to monitor the CPU utilization trend of the control processor and the specific internal tasks over time or to view the CPU utilization required for a specific event.

CLI Commands for Monitoring Control Processor CPU Utilization

Use this command to display the CPU utilization of the control processor.

From the SCE> prompt, type:

Command	Purpose
show processes cpu [sorted]	Displays the CPU utilization of the control processor.



Tip To display the CPU history sorted by percentage of utilization, use the **sorted** keyword.

Example for Monitoring Control Processor CPU Utilization

The following example shows how to display the CPU utilization of the control processor.

```
SCE>show processes cpu

CORE_0 CPU utilization for five seconds: 100%/  0%; one minute:  81%; five minutes:  38%
PID   Runtime(ms)   Invoked      uSecs   5Sec   1Min   5Min TTY Process
    1       78790      6374          0  0.00%  0.00%  0.00%   0 (init)
    2          10         1          0  0.00%  0.00%  0.00%   0 (kthreadd)
    3        5010       501          0  0.00%  0.00%  0.00%   0 (migration/0)
    4          90         9          0  0.00%  0.00%  0.00%   0 (ksoftirqd/0)
    5       63130      6313          0  0.00%  0.00%  0.00%   0 (watchdog/0)
    6        4940       494          0  0.00%  0.00%  0.00%   0 (migration/1)
    7           0         0          0  0.00%  0.00%  0.00%   0 (ksoftirqd/1)
    8       10530      1053          0  0.00%  0.00%  0.00%   0 (watchdog/1)
    9     2606490    207337          0  0.00%  0.02%  0.03%   0 (events/0)
   10     1246730    123793          0  0.00%  0.02%  0.02%   0 (events/1)
   11           0         0          0  0.00%  0.00%  0.00%   0 (khelper)
   12      177810     17781          0  0.00%  0.00%  0.00%   0 (kblockd/0)
   13        8010       801          0  0.00%  0.00%  0.00%   0 (kblockd/1)
   16           0         0          0  0.00%  0.00%  0.00%   0 (kswapd0)
   17           0         0          0  0.00%  0.00%  0.00%   0 (aio/0)
   18           0         0          0  0.00%  0.00%  0.00%   0 (aio/1)
   19           0         0          0  0.00%  0.00%  0.00%   0 (nfsiod)
   20           0         0          0  0.00%  0.00%  0.00%   0 (mtdblockd)
   21     1198570    119326          0  0.00%  0.02%  0.02%   0 (skynet)
   22     7413850    741207          0  0.00%  0.11%  0.10%   0 (hw-mon-regs)
   23      556170     49614          0  0.00%  0.02%  0.01%   0 (scos-dump)
   24      527310     52718          0  0.00%  0.00%  0.01%   0 (wdog-kernel)
```

The following table lists and describes the fields in the **show processes cpu** output.

Table 3-1 show processes cpu Command Output Fields

Field	Description
CPU utilization for five seconds	CPU utilization for the last five seconds. The first number indicates the total, the second number indicates the percent of CPU time spent at the interrupt level.
one minute	CPU utilization for the last minute
five minutes	CPU utilization for the last five minutes
PID	The process ID
Runtime (msecs)	CPU time the process has used, expressed in msecs
Invoked	The number of times the process has been invoked (progress once every 5sec)
uSecs	Microseconds of CPU time for each process invocation
5Sec	CPU utilization by task in the last five seconds
1Min	CPU utilization by task in the last minute
5Min	CPU utilization by task in the last five minutes
TTY	Currently not relevant in the Cisco Service Control system.
Process	Name of the process. For more information, refer to The Processes section of this document. **Note** Linux tasks are also presented as processes.

Note When CPU utilization is higher than about 90%, the CPU utilization per task is not reliable and can sum to more than 100%. This is because high CPU utilization can influence the task that samples CPU utilization.

Rebooting and Shutting Down the Cisco SCE Platform

- Rebooting the Cisco SCE Platform, page 3-23
- How to Shut Down the Cisco SCE Platform, page 3-23

Rebooting the Cisco SCE Platform

Rebooting the Cisco SCE platform is required after installing a new package, in order for that package to take effect. There might be other occasions where rebooting the Cisco SCE platform is necessary.

Note When the Cisco SCE restarts, it loads the startup configuration, so all changes made in the running configuration will be lost. You are advised to save the running configuration before performing reload, as described in "How to Save or Change the Configuration Settings" section on page 3-6.

Step 1 At the SCE# prompt, enter **reload** and press **Enter**.

A confirmation message appears.

Step 2 Press **Enter** to confirm the reboot request (accept default, which is *Yes*).

Examples for Rebooting the Cisco SCE Platform

The following example shows the commands for system reboot.

```
SCE# reload

Are you sure? Y.
the system is about to reboot, this will end your CLI session
```

How to Shut Down the Cisco SCE Platform

Shutting down the Cisco SCE platform is required before turning the power off. This helps to ensure that non-volatile memory devices in the Cisco SCE platform are properly flushed in an orderly manner.

Note When the Cisco SCE platform restarts, it loads the startup configuration, so all changes made in the running configuration will be lost. You are advised to save the running configuration before performing reload, as described in "How to Save or Change the Configuration Settings" section on page 3-6.

Step 1 Connect to the serial console port (The CON connector on the front panel of the Service Control module in slot #1, 9600 baud).

The SCE# prompt appears.

Step 2 Type **reload shutdown**.

A confirmation message appears.

Step 3 Type **y to** confirm the shutdown request and press **Enter**.

Example for Shutting Down the Cisco SCE Platform

The following example shows the commands for system shutdown.

```
SCE# reload shutdown

You are about to shut down the system.
The only way to resume system operation after this
is to cycle the power off, and then back on.
Continue?
y
IT IS NOW SAFE TO TURN THE POWER OFF.
```

Note Since the Cisco SCE platform can recover from the power-down state only by being physically turned off (or cycling the power), this command can only be executed from the serial CLI console. This limitation helps prevent situations in which users issue this command from a Telnet session, and then realize that they have no physical access to the Cisco SCE platform.



Utilities

Revised: September 23, 2015

Introduction

This chapter describes the following utilities:

- Working with Cisco SCE Platform Files, page 4-2
- The User Log, page 4-7
- Managing Syslog, page 4-10
- Flow Capture, page 4-17

Working with Cisco SCE Platform Files

The CLI commands include a complete range of file management commands. These commands allow you to create, delete, copy, and display both files and directories

Note Regarding disk capacity: While performing disk operations, the user should take care that the addition of new files that are stored on the Cisco SCE disk do not cause the disk to exceed 70%.

- Working with Directories, page 4-2
- Working with Files, page 4-4

Working with Directories

- How to Create a Directory, page 4-2
- How to Delete a Directory, page 4-2
- How to Change Directories, page 4-3
- How to Display your Working Directory, page 4-3
- How to List the Files in a Directory, page 4-3

How to Create a Directory

mkdir

From the SCE# prompt, type:

Command	Purpose
mkdir *directory-name*	Creates a directory.

How to Delete a Directory

There are two different commands for deleting a directory, depending on whether the directory is empty or not.

- How to Delete a Directory and All its Files, page 4-2
- How to Delete an Empty Directory, page 4-3

How to Delete a Directory and All its Files

delete

From the SCE# prompt, type:

Command	Purpose
delete *directory-name* **/recursive**	The recursive flag deletes all files and sub-directories contained in the specified directory.

How to Delete an Empty Directory

rmdir

From the SCE# prompt, type:

Command	Purpose
rmdir *directory-name*	Use this command only for an empty directory.

How to Change Directories

Use this command to change the path of the current working directory. **cd**

From the SCE# prompt, type:

Command	Purpose
cd *new path*	Changes the path of the current working directory.

How to Display your Working Directory

pwd

From the SCE# prompt, type:

Command	Purpose
pwd	Displays the working directory.

How to List the Files in a Directory

You can display a listing of all files in the current working directory. This list may be filtered to include only application files. The listing may also be expanded to include all files in any sub-directories.

- How to List the Files in the Current Directory, page 4-3
- How to List the Applications in the Current Directory, page 4-4
- How to Include Files in Sub-Directories in the Directory Files List, page 4-4

How to List the Files in the Current Directory

From the SCE# prompt, type:

Command	Purpose
dir	Lists the files in the current directory.

How to List the Applications in the Current Directory

From the SCE# prompt, type:

Command	Purpose
dir applications	Lists the applications in the current directory.

How to Include Files in Sub-Directories in the Directory Files List

From the SCE# prompt, type:

Command	Purpose
dir -r	Includes files in the sub-directories in the directory files list.

Working with Files

- How to Rename a File, page 4-4
- How to Delete a File, page 4-4
- Copying Files, page 4-5
- How to Display File Contents, page 4-6
- How to Unzip a File, page 4-6

How to Rename a File

From the SCE# prompt, type:

Command	Purpose
rename *current-file-name new-file-name*	Renames a file.

How to Delete a File

From the SCE# prompt, type:

Command	Purpose
delete *file-name*	Deletes a file.

Copying Files

You can copy a file from the current directory to a different directory. You can also copy a file (upload/download) to or from an FTP site.

To copy a file using passive FTP, use the **copy-passive** command.

- How to Copy a File, page 4-5
- How to Download a File from an FTP Site, page 4-5
- How to Upload a File to a Passive FTP Site, page 4-5

How to Copy a File

From the SCE# prompt, type:

Command	Purpose
copy *source-file-name destination-file-name*	Copies a file.

Copying a File: Example

The following example copies the local analysis.sli file located in the root directory to the applications directory.

```
SCE#copy analysis.sli applications/analysis.sli
sce#
```

How to Download a File from an FTP Site

Use the copy command to upload and download commands from and FTP site. In this case, either the source or destination filename must begin with *ftp://*.

Step 1 From the SCE# prompt, type **copy** *ftp://username:password@ip-address/path:/source-file destination-file-name* and press **Enter**.

To upload a file to an FTP site, specify the FTP site as the destination (*ftp://username:password@ip-address/path:/destination-file*)

How to Upload a File to a Passive FTP Site

Step 1 From the SCE# prompt, type **copy-passive** *source-file-name ftp:/username:password@10.10.10.10/h:/destination-file* and press **Enter**.

To download a file from a passive FTP site, specify the FTP site as the source (*ftp://username:password@ip-address/path:/source-file*)

Uploading a File to a Passive FTP Site: Example

The following example uploads the analysis.sli file located on the local flash file system to the host 10.1.1.1, specifying Passive FTP.

```
SCE#copy-passive /appli/analysis.sli ftp://myname:mypw@10.1.1.1/p:/appli/analysis.sli
sce#
```

How to Display File Contents

From the SCE# prompt, type:

Command	Purpose
more *file-name*	Displays file contents.

How to Unzip a File

From the SCE# prompt, type:

Command	Purpose
unzip *file-name*	Unzips a file.

The User Log

The user log is an ASCII file that can be viewed in any editor. It contains a record of system events, including startup, shutdown and errors. You can use the Logger to view the user log to determine whether or not the system is functioning properly, as well as for technical support purposes.

- The Logging System, page 4-7
- Generating a File for Technical Support, page 4-9

The Logging System

- Copying the User Log, page 4-7
- Enabling and Disabling the User Log, page 4-8
- Viewing the User Log Counters, page 4-8
- Viewing the User Log, page 4-9
- Clearing the User Log, page 4-9

Events are logged to one of two log files. After a file reaches maximum capacity, the events logged in that file are then temporarily archived. New events are then automatically logged to the alternate log file. When the second log file reaches maximum capacity, the system then reverts to logging events to the first log file, thus overwriting the temporarily archived information stored in that file.

Basic operations include:

- Copying the User Log to an external location
- Viewing the User Log
- Clearing the User Log
- Viewing/clearing the User Log counters

Copying the User Log

You can view the log file by copying it to an external location or to disk. This command copies both log files to the local Cisco SCE platform disk or any external host running a FTP server.

- Copying the User Log to an External Location, page 4-7
- Copying the User Log to an Internal Location, page 4-8

Copying the User Log to an External Location

From the SCE# prompt, type:

Command	Purpose
logger get user-log file-name *ftp://username:password@ipaddress/path*	Copies user log to an external location.

Copying the User Log to an Internal Location

From the SCE# prompt, type:

Command	Purpose
logger get user-log file-name *target-filename*	Copies the user log to an internal location.

Enabling and Disabling the User Log

By default, the user log is enabled. You can disable the user log by configuring the status of the logger.

- Disabling the User Log, page 4-8
- Enabling the User Log, page 4-8

Disabling the User Log

Step 1 From the SCE# prompt, type **configure** and press **Enter**.

Step 2 From the SCE (config)# prompt, type **logger device User-File-Log disabled** and press **Enter**.

Enabling the User Log

Step 1 From the SCE# prompt, type **configure** and press **Enter**.

Step 2 From the SCE (config)# prompt, type **logger device User-File-Log enabled** and press **Enter**.

Viewing the User Log Counters

- Viewing the user log counters for the current session, page 4-8
- Viewing the non-volatile counter for the user-file-log, page 4-9

There are two types of log counters:

- User log counters—Count the number of system events logged from the Cisco SCE platform last reboot.
- Non-volatile counters—These are not cleared during boot time

Viewing the user log counters for the current session

From the SCE# prompt, type:

Command	Purpose
show logger device user-file-log counters	Displays the user log counters for the current session.

Viewing the non-volatile counter for the user-file-log

From the SCE# prompt, type:

Command	Purpose
show logger device user-file-log nv-counters	Displays the non-volatile counter for the user-file-log.

Viewing the User Log

Note This command is not recommended when the user log is large. Copy a large log to a file to view it (see Copying the User Log, page 4-7)

From the SCE# prompt, type:

Command	Purpose
more user-log	Displays the user log.

Clearing the User Log

Step 1 From the SCE# prompt, type **clear logger device user-file-log** and press **Enter**.

Step 2 The system asks **Are you sure?**

Step 3 Type **Y** and press **Enter**.

Generating a File for Technical Support

In order for technical support to be most effective, the user should provide them with the information contained in the system logs. Use the logger get support-file command to generate a support file via FTP for the use of Cisco technical support staff.

From the SCE# prompt, type:

Command	Purpose
logger get support-file *filename*	Creates a support file. The support information file is created using the specified filename. The specified file must be a file located on an FTP site, not on the local file system. This operation may take some time.

Generating a File for Technical Support: Example

```
SCE#logger get support-file ftp://user:1234@10.10.10.10/c:/support.zip
```

Managing Syslog

System messages are also written to the Syslog server. When enabled, all user-log messages are sent to the configured Syslog servers as well as to the Cisco SCE user logs.

You can configure the following options for syslog support:

- Up to five remote syslog hosts
- Port number
- Minimum severity level to be logged
- Logging rate limit
- Syslog facility (such as system daemon, local printer, or user process)
- Time stamp format

Transport protocol is not configurable, since the Cisco SCE platform supports Syslog over UDP, only.

Enabling and Disabling Syslog

By default, logging to the syslog server is disabled.

Enabling Syslog

Step 1 From the SCE# prompt, type **configure** and press **Enter**.

Step 2 From the SCE (config)# prompt, type **logging on** and press **Enter**.

Disabling Syslog

Step 1 From the SCE# prompt, type **configure** and press **Enter**.

Step 2 From the SCE (config)# prompt, type **no logging on** and press **Enter**.

Configuring Remote Syslog Hosts

You can configure up to five remote syslog hosts. You can also assign a UDP port to each host.

Guidelines

- You can define the host using either the hostname or the IP address.
- To assign a port, you must use the **transport udp** option. If you are not assigning a port, this is not required, since UDP is the only transport protocol supported for Syslog on the Cisco SCE platform.
- Each host requires a separate command.

Options

The following options are available:

- **hostname**—Logical name of the remote host
- **ip-address**—IP address of the remote host
- **port-number**—Number of the UDP port (1 – 65535)
  - default = 514

How to Add a Remote Syslog Host

Step 1 From the SCE# prompt, type **configure** and press **Enter**.

Step 2 From the SCE (config)# prompt, type **logging host (*hostname* | *ip-address*) [transport udp [port *port-number*]]** and press **Enter**.

How to Remove a Remote Syslog Host

Step 1 From the SCE# prompt, type **configure** and press **Enter**.

Step 2 From the SCE (config)# prompt, type **no logging host (*hostname* | *ip-address*)** and press **Enter**.

Configuring the Minimum Severity Level to be Logged to Syslog

By default, all messages are logged to the Syslog server when it is enabled, with the exception of debug messages. However, you can configure the minimum severity level of the messages to logged to Syslog.

Table 4-1 lists the syslog severity levels and the corresponding SCOS severity levels. Not all syslog severity levels are supported on the Cisco SCE platform.

Table 4-1 Syslog and SCOS Severity Levels

Syslog Severity	Level	SCOS Severity	SCOS Definition
emergency	0	Not defined	SEVERITY_EMERGENCY_LEVEL
alert	1	Not defined	SEVERITY_ALERT_LEVEL
critical	2	Fatal	SEVERITY_FATAL_LEVEL
error	3	Error	SEVERITY_ERROR_LEVEL
warning	4	Warning	SEVERITY_WARNING_LEVEL
notice	5	Not defined	SEVERITY_NOTICE_LEVEL
informational	6	Info	SEVERITY_INFORMATIONAL_LEVEL
debug	7	Not defined	SEVERITY_DEBUG_LEVEL

Options

The following option is available:

- **severity-level**—The name of the desired severity level at which messages should be logged. Messages at or lower than the specified level are logged. Severity levels supported on the Cisco SCE platform are as follows:
  - **fatal**
  - **error**
  - **warning**
  - **info**

  Default = info

How to Configure the Minimum Severity Level for Syslog Messages

Step 1 From the SCE# prompt, type **configure** and press **Enter**.

Step 2 From the SCE (config)# prompt, type **logging trap** ***severity-level*** and press **Enter**.

How to Restore the Default Minimum Severity Level for Syslog Messages

Step 1 From the SCE# prompt, type **configure** and press **Enter**.

Step 2 From the SCE (config)# prompt, type **no logging trap** and press **Enter**.

Configuring the Syslog Facility

You can assign Syslog messages to a specified facility.

Options

The following option is available:

- **facility-type**—Syslog facility. See Table 4-2
  - default = local7

Table 4-2 logging facility types

Facility-type keyword	Description
auth	Authorization system
cron	Cron facility
daemon	System daemon
kern	Kernel
local0-local7	Reserved for locally defined messages
lpr	Line printer system

Table 4-2 logging facility types (continued)

Facility-type keyword	Description
mail	Mail system
news	USENET news
sys9	System use
sys10	System use
sys11	System use
sys12	System use
sys13	System use
sys14	System use
syslog	System log
user	User process
uucp	UNIX-to-UNIX copy system

How to Configure the Syslog Facility

Step 1 From the SCE# prompt, type **configure** and press **Enter**.

Step 2 From the SCE (config)# prompt, type **logging facility** ***facility-type*** and press **Enter**.

How to Restore the Default Syslog Facility

Step 1 From the SCE# prompt, type **configure** and press **Enter**.

Step 2 From the SCE (config)# prompt, type **no logging facility** and press **Enter**.

Configuring the Syslog Logging Rate Limit

You can configure a maximum number of messages logged per second. In addition, you can specify a severity level above which the rate is unlimited. For example, you can configure a rate limit for all messages below the **fatal** severity level.

Options

The following options are available:

- **rate**—Number of messages to be logged per second (1 to 10000).
  - default = 10
- **severity-level**—Excludes messages of this severity level and higher. Severity levels supported on the Cisco SCE platform are as follows:
  - **fatal**
  - **error**
  - **warning**
  - **info**

How to Configure the Syslog Rate Limit

Step 1 From the SCE# prompt, type **configure** and press **Enter**.

Step 2 From the SCE (config)# prompt, type **logging rate-limit** *rate* **[except** *severity-level***]** and press **Enter**.

How to Restore the Default Syslog Rate Limit

Step 1 From the SCE# prompt, type **configure** and press **Enter**.

Step 2 From the SCE (config)# prompt, type **no logging rate-limit** and press **Enter**.

Configuring the Syslog Time Stamp Format

You can configure the format of the the time stamp on the messages on the Syslog server. You can use the **no** form of this command to specify the default Syslog time stamp format (uptime).

Options

The following time stamp format options are available:

- **uptime** (default)—Time stamp shows time since the system was last rebooted. For example "4w6d" (time since last reboot is 4 weeks and 6 days).
- **datetime**—Time stamp shows date and time.

  The following additional options are available for the **datetime option:**
  - **msec—**Include milliseconds in the date-time format.
  - **localtime—**Time stamp relative to the local time zone.
  - **show-timezone—**Include the time zone name in the date-time format.
  - **year—**Include the year in the date-time format.

  If the **datetime** keyword is used without additional keywords, time stamps will be shown using UTC, without the year, without milliseconds, and without a time zone name.

Tip The optional **msec**, **localtime**, **show-timezone**, and **year** keywords, if present, must be in the order shown in the command syntax. All keywords up to the last specified keyword must be present
Incorrect: **service timestamps log datetime msec year**
Correct: **service timestamps log datetime msec localtime show-timezone year**

How to Configure the Syslog Time Stamp Format

Step 1 From the SCE# prompt, type **configure** and press **Enter**.

Step 2 From the SCE (config)# prompt, type **service timestamps log (uptime | (datetime [msec] [localtime] [show-timezone] [year]))** and press **Enter**.

How to Restore the Default Syslog Time Stamp Format

Step 1 From the SCE# prompt, type **configure** and press **Enter**.

Step 2 From the SCE (config)# prompt, type **no service timestamps log** and press **Enter**.

Enabling and Disabling the Syslog Message Counter

By default, the syslog message counter is enabled. You can use this command to disable the syslog message counter. When it is disabled, no line count appears in the syslog messages.

Disabling the Syslog Message Counter

Step 1 From the SCE# prompt, type **configure** and press **Enter**.

Step 2 From the SCE (config)# prompt, type **no logging message-counter** and press **Enter**.

Enabling the Syslog Message Counter

Step 1 From the SCE# prompt, type **configure** and press **Enter**.

Step 2 From the SCE (config)# prompt, type **logging message-counter** and press **Enter**.

Monitoring Syslog

You can display the following Syslog information:

- Current Syslog server configuration.
- Syslog counters

How to Display the Syslog Configuration

From the SCE# prompt, type:

Command	Purpose
show logging	Displays the syslog configuration.

How to Display the Syslog Counters

From the SCE# prompt, type:

Command	Purpose
show logging counters	Displays the syslog counters.

Flow Capture

- Limitations, page 4-17
- The Flow Capture Process, page 4-18

The flow capture utility is a CLI-controlled utility that is used to capture traffic according to Layer 4 attributes.

The traffic flow captured by this utility is accumulated in a cap format file. Traffic that is identified by the capture mechanism is not available for traffic control or any other service for the duration of the capture. When flow capture is configured, the traffic that matches the rules are bypassed on Cisco SCE. After the completion of a capture, normal service to all the traffic is resumed.

The recorded data is sent online to a distant location using FTP. The data is sent in a standard format, with a maximum file size of 128 MB on the Cisco SCE 8000 platform (configurable by a const DB).

In Cisco SCE 8000 that has two SCM modules, a separate cap file is created by each SCM module, each with a maximum file size of 64 MB.

Note Flow capture is used for troubleshooting. Enabling the flow capture may impact the CPU performance. We recommend that you avoid enabling flow capture in a production environment.

Limitations

Note the following known limitations of the flow capture utility:

- The actual capture starts only for newly opened flows. Therefore, already opened flows cannot be captured by this utility.
- The termination of a capture flow is verified for every new relevant packet that is being captured. As long as no packets matching the capturing attributes arrives after the time is exceeded, the capturing is not stopped and must be stopped manually.
- File size limitation: the maximum captured file size is limited on Cisco SCE 8000 platform to 128MB (configurable by a const DB). In this platform the FTP connection is opened when the capturing starts, but the data is first kept on the system disk and is transferred to the FTP destination only when the capturing procedure is concluded.
- In a system with two Cisco SCE 8000-SCM modules, which creates two capture files, the maximum file size for each file is 64MB, for a total of 128MB.
- FTP timeout: As the recorded traffic might be sparse, and since on the Cisco SCE 8000 platform the data is sent only upon termination of the capture, it is required to disable any connection timeout on the destination FTP server.
- Capture may end prematurely due to a shortage event on the Cisco SCE platform.
- Capturing throughput is limited by the following:
  - system architectural limitations
  - line capacity to the remote FTP destination (for non-Linux platforms only, such as the Cisco SCE 2000 platform).

  The approximated throughput on a live setup is 2Mbps. When this throughput is exceeded, packets are absent from the cap file and the appropriate field in the consequent captured packet is updated to note the number of lost packets. The maximum allowed number of sequential lost packets is configurable by a const DB.

The Flow Capture Process

There are three main steps in the overall flow capture process:

1. Configure the traffic rules to define the traffic to be captured. ("Configuring a Flow Capture Traffic Rule" section on page 4-18)
2. Configure the flow capture settings. (Optional) ("Configuring the Flow Capture Settings" section on page 4-18)
3. Perform the actual flow capture. ("Performing the Flow Capture" section on page 4-20)

Configuring a Flow Capture Traffic Rule

The flow capture traffic rules define the traffic to be captured. You can configure a flow capture traffic rule by specifying the **flow-capture** action for the relevant flows.

For example, in order to capture all the traffic sent to or coming from subscribers whose IP addresses are in the range.2.3.0-1.2.3.255, define a traffic rule as follows:

```
SCE(config if)# traffic-rule name flowcapturerule IP-addresses subscriber-side 1.2.3.0/24
network-side all protocol all direction both traffic-counter none action flow-capture
```

Multiple rules can be configured, but note that all configured flow capture rules are in effect during the flow capture process. It is not possible to apply only a subset of the configured rules.

For more information regarding configuring traffic rules, see "Configuring Traffic Rules and Counters" section on page 7-26.

Configuring the Flow Capture Settings

The flow capture settings control aspects of the flow capture process, as opposed to defining the flow to be captured. These settings limit the scope of the process to maximize the recorded information while minimizing the effect on traffic.

- Maximum duration of the capture: By limiting the duration of the capture, you can limit the effect of the capture on live traffic.

  You can stop the capture at any time before the maximum duration has been reached.

- Maximum length of the L4 payload of each captured packet: If you want to capture mainly the L2-L4 headers, you need only a small portion of the payload of each packet. Setting a limit on the length of the payload makes the capture more efficient, as it allows more packets to be captured within a given time frame and for a given throughput.

  Guidelines and Information:

  – If maximum L4 payload length is not configured, all bytes of each captured packet are recorded.

  – If maximum L4 payload length is configured, each captured packet will contain the entire L2/L3/L4 headers and no more than the configured maximum bytes of L4 payload.

  – Only one maximum L4 payload length value can be configured. This value applies to all recorded packets.

- If the maximum L4 payload length value is changed while recording is performed, it will not take effect until the next recording session.
- The cap file contains marking for packets which had TCP or UDP checksum error when received in the Cisco SCE platform, since the validity of the TCP and UDP checksum cannot be checked for the captured packets due to missing bytes.
- The cap file contains the information to retrieve the original length of each packet that was truncated.

How to Configure the Maximum Flow Capture Duration

The following options are available:

- **duration—**The maximum duration of the flow capture in seconds.

  Default = 3600 seconds
- **unlimited—**There is no time limit to the flow capture, and it will continue until stopped by the operator.

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

Command	Purpose
flow-capture controllers time (*duration* | **unlimited**)	Configures the maximum flow capture duration.

How to Configure the Maximum Length of the L4 Payload

The following options are available:

- **length—**The maximum number of L4 payload bytes to capture from each packet.
- **unlimited**—There is no limit on the number of L4 payload bytes. (Default)

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

Command	Purpose
flow-capture controllers max-l4-payload-length (*length* | **unlimited**)	Configures the maximum length of the L4 payload.

How to Restore the Default Flow Capture Settings

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

Command	Purpose
default flow-capture controllers (time | max-l4-payload-length)	Restores the default flow capture settings.

Performing the Flow Capture

The flow capture begins when you execute the flow-capture command. You can stop the capture at any time. If the capture is not stopped, it continues for the configured maximum duration ("Configuring the Flow Capture Settings" section on page 4-18).

How to Start a Flow Capture

The following option is available:

- **filename**—Name and FTP location to which to record the flow capture data in the format *ftp://<username>:<password>@<IP_address>/<path>/<file_name>*. (Do not include the ".cap" file extension; it is appended automatically.)

  In a system with two Cisco SCE 8000-SCM modules, which creates two capture files, an indicator is appended to this prefix to indicate which Cisco SCE 8000-SCM module created the file. For example, if you assign the filename "myCapFile", the system creates *myCapFile1.cap* and *myCapFile2.cap*.

From the SCE# prompt, type:

Command	Purpose
flow-capture start format cap file-name-prefix *filename*	Starts a flow capture.

How to Stop a Flow Capture

From the SCE# prompt, type:

Command	Purpose
flow-capture stop	Stops a flow capture.

Monitoring the Flow Capture

Use the following command to monitor the flow capture process. It displays the following information:

- status of the recording process
- current target file size
- number of packets captured
- number of packets lost
- configured values of the different controllers

How to Monitor the Flow Capture

From the SCE> prompt, type:

Command	Purpose
show interface linecard 0 flow-capture	Monitors the flow capture.



Configuring the Management Interface and Security

Revised: September 23, 2015

Introduction

This chapter describes how to configure the physical management interfaces (ports) as well as the various management interface applications, such as SNMP, SSH, and TACACS+. It also explains how to configure users, passwords, IP configuration, clock and time zone, and domain name settings.

- Management Interface and Security, page 5-2
- Configuring the Management Ports, page 5-3
- Configuring Management Interface VLANs, page 5-11
- TACACS+ Authentication, Authorization, and Accounting, page 5-15
- Configuring Access Control Lists (ACLs), page 5-32
- Managing the Telnet Interface, page 5-35
- Configuring the SSH Server, page 5-37
- Configuring and Managing the SNMP Interface, page 5-41

Management Interface and Security

The Cisco SCE 8000 platform is equipped with two RJ-45 management ports (Port1 and Port2 on the Cisco SCE 8000-SCM-E module in slot 1). These ports provide access from a remote management console to the Cisco SCE platform via a LAN.

The two management ports support management interface redundancy, providing the possibility for a backup management link.

Note The second management port is reflected in all objects related to it in the SNMP interface.

Note Cisco SCE 8000 does not support IPv6 addresses or connectivity on the management interfaces.

Perform the following tasks to configure the management interface and management interface security:

- Configure the management port:
  - Physical parameters
  - Specify active port (if not redundant installation)
  - Redundancy (if redundant installation)
- Configure management interface security
  - Configure the permitted and not-permitted IP addresses

Configuring the Management Ports

- Entering the Management Interface Configuration Mode, page 5-3
- Configuring the Management Port Physical Parameters, page 5-4
- Management Interface Redundancy, page 5-8
- Monitoring the Management Interface, page 5-10

Perform the following tasks to configure the management ports:

- Configure the IP address and subnet mask (only one IP address for the management interface, not one IP address per port).
- Configure physical parameters:
  - Duplex
  - Speed
  - active port (optional)
- Configure redundant management interface behavior (optional):
  - Failover mode

Step 1 Cable the desired management port, connecting it to the remote management console via the LAN. If connecting both management ports for redundancy, connect the to the LAN using a switch.

Step 2 Configure the management port physical parameters. (See "Configuring the Management Port Physical Parameters" section on page 5-4.)

Step 3 (Optional) Configure the system with management interface redundancy. (see "Management Interface Redundancy" section on page 5-8.)

Entering the Management Interface Configuration Mode

When entering Management Interface Configuration Mode, you must indicate the number of the management port to be configured:

- 0/1—Port1
- 0/2—Port2

The following Management Interface commands are applied only to the port specified when entering Management Interface Configuration Mode. Therefore, each port must be configured separately:

- **speed**
- **duplex**

The following Management Interface commands are applied to both management ports, regardless of which port had been specified when entering Management Interface Configuration Mode. Therefore, both ports are configured with one command:

- **ip address**
- **auto-failover**

The GBE management interface is configured as follows:

- mode: MNG Interface configuration mode
- interface designation: 0/1 or 0/2

Step 1 Enter **configure** and press **Enter**.

Enters Global Configuration mode.

The command prompt changes to SCE(config)#.

Step 2 Enter **interface mng (0/1 | 0/2)** and press **Enter**.

Enters Management Interface Configuration mode.

The command prompt changes to SCE(config if)#

Configuring the Management Port Physical Parameters

This interface has a transmission rate of 10, 100, or 1000 Mbps and is used for management operations and for transmitting RDRs, which are the output of traffic analysis and management operations.

- Setting the IP Address and Subnet Mask of the Management Interface, page 5-4
- Configuring the Management Interface Speed and Duplex Parameters, page 5-5
- Specifying the Active Management Port, page 5-7

Setting the IP Address and Subnet Mask of the Management Interface

You must define the IP address of the management interface.

When both management ports are connected, providing a redundant management port, this IP address always acts as a virtual IP address for the currently active management port, regardless of which port is the active port.

Options

The following options are available:

- **IP address—**The IP address of the management interface.

  If both management ports are connected, so that a backup management link is available, this IP address will be act as a virtual IP address for the currently active management port, regardless of which physical port is currently active.

  The following IP addresses are used internally by the Cisco SCE 8000 platform and cannot be assigned to the management interface:

  - 192.168.207.241 to 192.168.207.255
  - 92.168.207.145 to 192.168.207.159

- **subnet mask—**Subnet mask of the management interface.

Step 1 On a physically connected local console, access the interface configuration mode for either management interface. The specified IP address is configured for both interfaces.

From the SCE(config)# prompt, type **interface Mng (0/1 | 0/2)** and press **Enter**.

Step 2 From the SCE(config if)# prompt, type **ip address** *ip-address subnet-mask* and press **Enter**.

The command might fail if there is a routing table entry that is not part of the new subnet defined by the new IP address and subnet mask.

Caution Changing the IP address of the management interface via telnet will result in loss of the telnet connection and inability to reconnect with the interface.

Note After changing the IP address, you must reload the Cisco SCE platform so that the change will take effect properly in all internal and external components of the Cisco SCE platform. (See "Rebooting and Shutting Down the Cisco SCE Platform" section on page 3-23.)

Setting the IP Address and Subnet Mask of the Management Interface: Example

The following example shows how to set the IP address of the Cisco SCE platform to 10.1.1.1 and the subnet mask to 255.255.0.0.

```
SCE#config
SCE(config)#interface mng 0/1
SCE(config if)#ip address 10.1.1.1 255.255.0.0
```

Configuring the Management Interface Speed and Duplex Parameters

This section presents sample procedures that describe how to configure the speed and the duplex of the management interface.

Both these parameters must be configured separately for each port.

- Interface State Relationship to Speed and Duplex, page 5-6
- How to Configure the Speed of the Management Interface, page 5-6
- How to Configure the Duplex Operation of the Management Interface, page 5-7

Interface State Relationship to Speed and Duplex

Table 5-1 summarizes the relationship between the interface state and speed and duplex.

Note It is not possible to set one parameter to "Auto" and to specify the other parameter. If either speed or duplex is set to "Auto", then both parameters will behave as if set to "Auto".

Table 5-1 Interface State Relationship to Speed and Duplex

Speed	Duplex	Actual GBE Interface State
Auto	Auto	Auto negotiation
Auto	Full/Half	Auto negotiation
10/100/1000	Auto	Auto negotiation
10	Full	10 Mbps and full duplex
10	Half	10 Mbps and half duplex
100	Full	100 Mbps and full duplex
100	Half	100 Mbps and half duplex

How to Configure the Speed of the Management Interface

Options

The following options are available:

- **speed—**Speed in Mbps of the currently selected management port (0/1 or 0/2):
  - **10**
  - **100**
  - **auto** (default)—Auto-negotiation (do not force speed on the link)

If the duplex parameter is configured to **auto**, changing the speed parameter has no effect.

Step 1 Access the interface configuration mode for the management interface you want to configure.

From the SCE(config)# prompt, type **interface Mng (0/1 | 0/2)** and press **Enter**.

Step 2 From the SCE(config if)# prompt, type **speed (10|100|auto)** and press **Enter**.

Specify the desired speed option.

Configuring the Speed of the Management Interface: Example

The following example shows how to use this command to configure the Management port to 100 Mbps speed.

```
SCE#config
SCE(config)#interface mng 0/1
SCE(config if)#speed 100
```

How to Configure the Duplex Operation of the Management Interface

Options

The following options are available:

- **duplex—**Duplex operation of the management port (0/1 or 0/2):
  - full
  - half
  - auto (default)—Auto-negotiation (do not force duplex on the link)

If the speed parameter is configured to **auto**, changing the duplex parameter has no effect.

Step 1 Access the interface configuration mode for the management interface you want to configure.

From the SCE(config)# prompt, type **interface Mng (0/1 | 0/2)** and press **Enter**.

Step 2 From the SCE(config if)# prompt, type **duplex auto|full|half** and press **Enter**.

Specify the desired duplex option.

Configuring the Duplex Operation of the Management Interface: Example

The following example shows how to use this command to configure the management port to half duplex mode.

```
SCE#config
SCE(config)#interface mng 0/2
SCE(config if)#duplex half
```

Specifying the Active Management Port

This command explicitly specifies which management port is currently active. Its use varies slightly, depending on whether the management interface is configured as a redundant interface (auto failover enabled) or not (auto failover disabled).

- auto failover enabled (automatic mode)—The specified port becomes the currently active port, in effect forcing a failover action even if a failure has not occurred.
- auto failover disabled (manual mode)—The specified port should correspond to the cabled Mng port, which is the only functional port and therefore must be and remain the active management port.

Step 1 Access the interface configuration mode for the management interface you want to configure as the active management port.

From the SCE(config)# prompt, type **interface Mng (0/1 | 0/2)** and press **Enter**.

Step 2 Type **active-port** and press **Enter**

Specifying the Active Management Port: Example

The following example shows how to use this command to configure Mng port 2 as the currently active management port.

```
SCE#config
SCE(config)#interface mng 0/2
SCE(config if)#active-port
```

Management Interface Redundancy

- Configuring the Management Ports for Redundancy, page 5-8
- Configuring the Fail-Over Mode, page 5-9

The Cisco SCE platform contains two RJ-45 management ports. The two management ports provide the possibility for a redundant management interface, thus ensuring management access to the Cisco SCE platform even if there is a failure in one of the management links. If a failure is detected in the active management link, the standby port automatically becomes the new active management port.

Note that both ports must be connected to the management console via a switch. In this way, the IP address of the MNG port is always the same, regardless of which physical port is currently active.

Important information:

- Only one port is active at any time.
- The same virtual IP address and MAC address are assigned to both ports.
- Default:
  - Port 1 = active
  - Port 2 = standby
- The standby port sends no packets to the network and packets from the network are discarded.
- When a problem in the active port is encountered, the standby port automatically becomes the new active port.
- Link problem, with switch to standby MNG port, is declared after the link is down for 300 msec.
- Service does not revert to the default active port if/when that link recovers. The currently active MNG port remains active until link failure causes a switch to the other MNG port.

Configuring the Management Ports for Redundancy

Step 1 Cable both management ports (Mng 1 and Mng 2), connecting them both to the remote management console via the LAN and via a switch.

Using the switch ensures that the IP address of the MNG port is always the same, regardless of which physical port is currently active

Step 2 Configure the automatic failover mode.

See Configuring the Fail-Over Mode, page 5-9.

Step 3 Configure the IP address for the management interface.

The same IP address will always be assigned to the active management port, regardless of which physical port is currently active.

See Setting the IP Address and Subnet Mask of the Management Interface, page 5-4.

Step 4 Configure the speed and duplex for both management ports.

See Configuring the Management Interface Speed and Duplex Parameters, page 5-5.

Configuring the Fail-Over Mode

- Options, page 5-9
- How to Enable the Automatic Fail-Over Mode, page 5-9
- How to Disable the Automatic Fail-Over Mode, page 5-9

Use the following command to enable automatic failover. The automatic mode must be enabled to support management interface redundancy. This mode automatically switches to the backup management link when a failure is detected in the currently active management link.

This parameter can be configured when in management interface configuration mode for either management port, and is applied to both ports with one command.

Options

The following options are available:

- **auto/ no auto—**Enable or disable automatic failover switching mode
  - Default—Auto (automatic mode)

How to Enable the Automatic Fail-Over Mode

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

Command	Purpose
auto-fail-over	Enables automatic failover mode.

How to Disable the Automatic Fail-Over Mode

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

Command	Purpose
no auto-fail-over	Disables automatic failover mode.

Monitoring the Management Interface

Use this command to display the following information for the management interface.

- speed
- duplex
- IP address
- auto-failover configuration

From the SCE# prompt, type:

Command	Purpose
show mng interface Mng (0/1 | 0/2) [speed | duplex | |ip address | auto-fail-over]	Displays the specified GBE management interface configuration for the specified interface. If no option is specified, all management interface information is displayed for the specified interface.

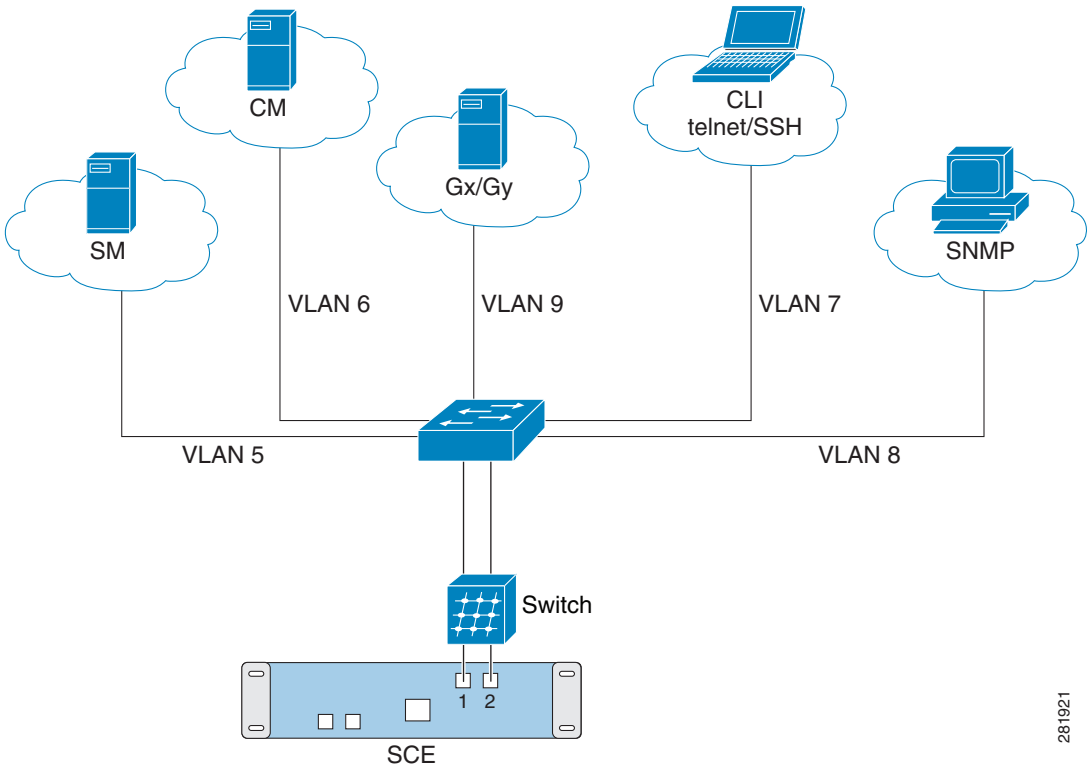
Configuring Management Interface VLANs

The Cisco SCE management network interface is used for various management services such as:

- Accessing the Cisco SCE shell through Telnet or SSH.
- SNMP

Management interface VLANS provide a way to distinguish between these management services (Telnet, SSH, SNMP) by using separate VLANs carried over same physical port (see Figure 5-1).

Figure 5-1 Management Interface VLANS



There are two steps in configuring management VLANs:

1. Create the VLAN and assign an IP address (**mng-vlan** command).
2. Assign the VLAN to a management or control service. Use one of the following commands:
   - **ip ssh mng-vlan**
   - **vty mng-vlan**
   - **snmp-server mng-vlan**

When a new management VLAN is configured, the device adds a new routing entry in the routing table. So, each configured IP address should be routable from the Cisco SCE.

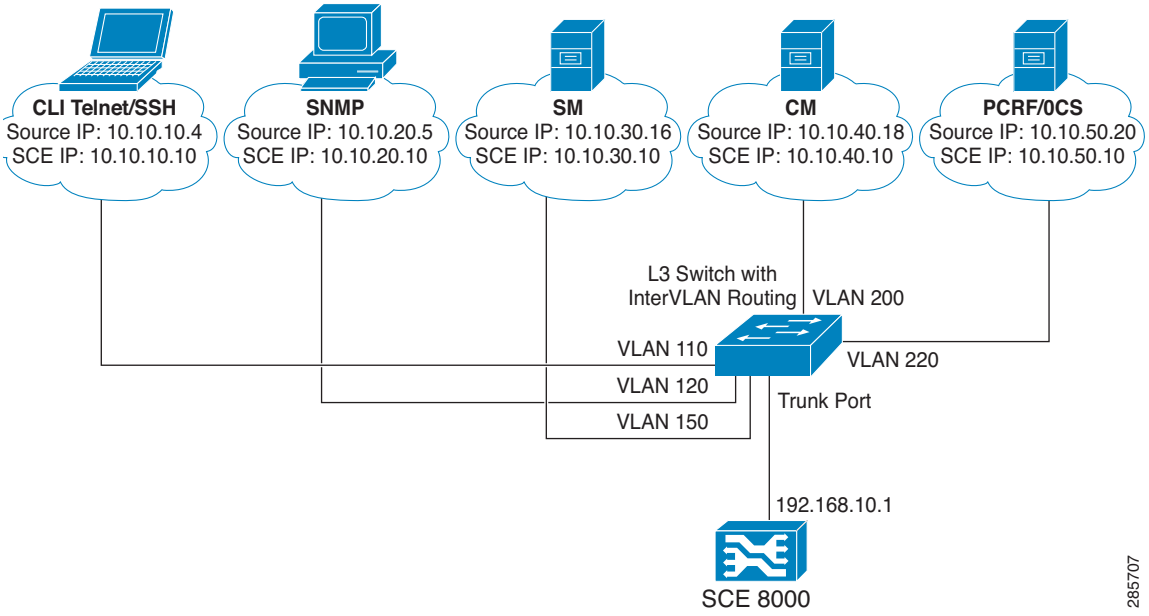
Here is an example of a sample configuration.

The following are some of the valid IP entries:

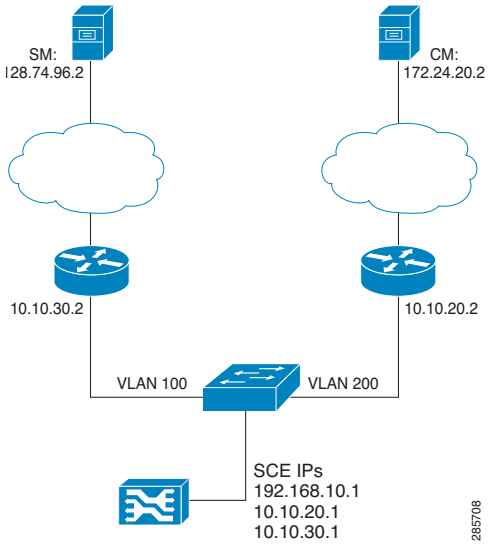
- 10.10.10.10/255.255.255.192 - VLAN ID 100
- 10.10.20.10/255.255.255.192 - VLAN ID 120
- 10.10.30.10/255.255.255.192 - VLAN ID 150

If the IP addresses listed above are added, the device creates the following entries in the routing table:

Destination	Gateway	Genmask	Flags	Metric	Ref	Use	Iface
10.10.20.0	*	255.255.255.192	U	0	0		0bond0.120
10.10.30.0	*	255.255.255.192	U	0	0		0bond0.150
10.10.10.0	*	255.255.255.192	U	0	0		0bond0.100



The following diagram provides another view of the configured management VLAN:



SUMMARY STEPS

1. **enable**
2. **configure**
3. **mng-vlan** *vlan-id* **address** *ip-address* **mask** *mask*
4. **<service> mng-vlan** *vlan-id*

DETAILED STEPS

	Command	Purpose
Step 1	**enable** **Example:** `SCE> enable`	Enables privileged EXEC mode. Enter your password when prompted.
Step 2	**configure** **Example:** `SCE#> configure`	Enters global configuration mode.

	Command	Purpose
Step 3	**mng-vlan** *vlan-id* **address** *ip-address* **mask** *mask* **Example:** `SCE(config)#> mng-vlan 10 address 10.10.10.20 mask 255.255.255.0`	Creates the VLAN and assigns the IP address to the VLAN. Range for VLAN tag: 1-4094
Step 4	**<service> mng-vlan** *vlan-id* **Example:** `SCE(config)#> snmp-server mng-vlan 10`	Assigns the VLAN to the specified service. Service options are: • ip ssh • vty • snmp-server **Note** The **snmp-server mng-vlan** command, in either the positive or negative form, restarts the SNMP process in order for the changes to take effect. This generates a cold-start trap.

Monitoring Management VLANs

To monitor managment VLANs, use one or more of the following commands.

These commands are in viewer mode.

Command	Purpose
show mng-vlan [*vlan-id* | **all**]	Displays the IP configuration and configured management service for the specified management VLAN.
show vlan *vlan-id* **service-bind**	Displays the management service assigned to the specified management VLAN.
show vlan *vlan-id* **statistics**	Displays the traffic statistics for the specified VLAN.
show vty mng-vlan	Displays the management interface VLAN configured for Telnet services.
show ip ssh mng-vlan	Displays the management interface VLAN configured for SSH services.

TACACS+ Authentication, Authorization, and Accounting

- Information About TACACS+ Authentication, Authorization, and Accounting, page 5-15
- Configuring the Cisco SCE Platform TACACS+ Client, page 5-19
- Managing the User Database, page 5-22
- Configuring AAA Login Authentication, page 5-26
- Configuring AAA Privilege-Level Authorization Methods, page 5-28
- Configuring AAA Command-Level Authorization Methods, page 5-28
- Configuring AAA Accounting, page 5-29
- Monitoring TACACS+, page 5-30
- Monitoring TACACS+ Users, page 5-31

Information About TACACS+ Authentication, Authorization, and Accounting

- Login Authentication, page 5-15
- Accounting, page 5-16
- Privilege-Level Authorization, page 5-16
- General AAA Fallback and Recovery Mechanism, page 5-17
- About Configuring TACACS+, page 5-18

TACACS+ is a security application that provides centralized authentication of users attempting to gain access to a network element. The implementation of TACACS+ protocol allows customers to configure one or more authentication servers for the Cisco SCE platform, providing a secure means of managing the Cisco SCE platform, as the authentication server will authenticate each user. This then centralizes the authentication database, making it easier for the customers to manage the Cisco SCE platform.

TACACS+ services are maintained in a database on a TACACS+ server running, typically, on a UNIX or Windows NT workstation. You must have access to and must configure a TACACS+ server before the configured TACACS+ features on your network element are available.

The TACACS+ protocol provides authentication between the network element and the TACACS+ ACS, and it can also ensure confidentiality, if a key is configured, by encrypting all protocol exchanges between a network element and a TACACS+ server.

The TACACS+ protocol provides the following three features:

- Login authentication
- Privilege level authorization
- Accounting

Login Authentication

The Cisco SCE platform uses the TACACS+ ASCII authentication message for CLI, Telnet and SSH access.

TACACS+ allows an arbitrary conversation to be held between the server and the user until the server receives enough information to authenticate the user. This is usually done by prompting for a username and password combination.

The login and password prompts may be provided by the TACACS+ server, or if the TACACS+ server does not provide the prompts, then the local prompts will be used.

The user log in information (user name and password) is transmitted to the TACACS+ server for authentication. If the TACACS+ server indicates that the user is not authenticated, the user will be re-prompted for the user name and password. The user is re-prompted a user-configurable number of times, after which the failed login attempt is recorded in the Cisco SCE platform user log and the telnet session is terminated (unless the user is connected to the console port.)

The Cisco SCE platform will eventually receive one of the following responses from the TACACS+ server:

- ACCEPT – The user is authenticated and service may begin.
- REJECT – The user has failed to authenticate. The user may be denied further access, or will be prompted to retry the login sequence depending on the TACACS+ server.
- ERROR – An error occurred at some time during authentication. This can be either at the server or in the network connection between the server and the Cisco SCE platform. If an ERROR response is received, the Cisco SCE platform will try to use an alternative method\server for authenticating the user.
- CONTINUE – The user is prompted for additional authentication information.

If the server is unavailable, the next authentication method is attempted, as explained in "General AAA Fallback and Recovery Mechanism" section on page 5-17.

Accounting

The TACACS+ accounting supports the following functionality:

- Each executed command (the command must be a valid one) will be logged using the TACACS+ accounting mechanism (including login and exit commands).
- The command is logged both before and after it is successfully executed.
- Each accounting message contains the following:
  - User name
  - Current time
  - Action performed
  - Command privilege level

TACACS+ accounting is in addition to normal local accounting using the Cisco SCE platform dbg log.

Privilege-Level Authorization

After a successful login the user is granted a default privilege level of 0, giving the user the ability to execute a limited number of commands. Changing privilege level is done by executing the "enable" command. This command initiates the privilege level authorization mechanism.

Privilege level authorization in the Cisco SCE platform is accomplished by the use of an "enable" command authentication request. When a user requests an authorization for a specified privilege level, by using the "enable" command, the Cisco SCE platform sends an authentication request to the TACACS+ server specifying the requested privilege level. The Cisco SCE platform grants the requested privilege level only after the TACACS+ server does the following:

- Authenticates the " **enable** " command password

- Verifies that the user has sufficient privileges to enter the requested privilege level.

Once the user privilege level has been determined, the user is granted access to a specified set of commands according to the level granted.

As with login authentication, if the server is unavailable, the next authentication method is attempted, as explained in "General AAA Fallback and Recovery Mechanism" section on page 5-17.

Command-Level Authorization

When command level authorization is enabled, each CLI command that is issued must be authorized by the external TACACS server before the system actually executes the command. You can configure the authorization level at which command level authorization is required. For example, you can require command level authorization only at root level.

As with login and privilege level authentication, if the TACACS+ server is unavailable, the regular fall back mechanism will be used.

General AAA Fallback and Recovery Mechanism

The Cisco SCE platform uses a fall-back mechanism to maintain service availability in case of an error.

The AAA methods available are:

- **TACACS+** – AAA is performed by the use of a TACACS+ server, allows authentication, authorization and accounting.
- **Local** – AAA is performed by the use of a local database, allows authentication and authorization.
- **Enable** – AAA is performed by the use of user configured passwords, allows authentication and authorization.
- **None** – no authentication\authorization\accounting is performed.

In the current implementation the order of the methods used isn't configurable but the customer can choose which of the methods are used. The current order is:

- **TACACS+**
- **Local**
- **Enable**
- **None**

Caution If the server goes to AAA fault, the Cisco SCE platform will not be accessible until one of the AAA methods is restored. In order to prevent this, it is advisable to use the "none" method as the last AAA method. If the Cisco SCE platform becomes inaccessible, the shell function "AAA_MethodsReset" will allow you to delete the current AAA method settings and set the AAA method used to "enable".
To run the "AAA_MethodsReset" shell function, complete the following steps:
1. Connect to AUX with username "root"
2. Run the debug shell: **scos_xinetd --service debug-shell --on**
3. Use Telnet to access the shell: **telnet localhost 2301**
4. Run the shell function: **AAA_MethodsReset**

About Configuring TACACS+

The following is a summary of the procedure for configuring TACACS+. All steps are explained in detail in the remainder of this section.

1. Configure the remote TACACS+ servers.

   Configure the remote servers for the protocols. Keep in mind the following guidelines

   - Configure the encryption key that the server and client will use.
   - The maximal user privilege level and enable password (password used when executing the enable command) should be provided.
   - The configuration should always include the root user, giving it the privilege level of 15.
   - Viewer (privilege level 5) and superuser (privilege level 10) user IDs should be established at this time also.

1. For complete details on server configuration, refer to the appropriate configuration guide for the particular TACACS+ server that you will be using.
2. Configure the Cisco SCE client to work with TACACS+ server:
   - hostname of the server
   - port number
   - shared encryption key (the configured encryption key must match the encryption key configured on the server in order for the client and server to communicate.)
3. (Optional) Configure the local database, if used.
   - add new users

   If the local database and TACACS+ are both configured, it is recommended to configure the same user names in both TACACS+ and the local database. This will allow the users to access the Cisco SCE platform in case of TACACS+ server failure.



Note If TACACS+ is used as the login method, the TACACS+ username is used automatically in the enable command. Therefore, it is important to configure the same usernames in both TACACS+ and the local database so that the enable command can recognize this username.

   - specify the password
   - define the privilege level
4. Configure the authentication methods on the Cisco SCE platform.
   - login authentication methods
   - privilege level authorization methods
   - command level authorization methods
5. Review the configuration.

   Use the " **show running-config** " command to view the configuration.

Configuring the Cisco SCE Platform TACACS+ Client

- Adding a New TACACS+ Server Host, page 5-19
- Removing a TACACS+ Server Host, page 5-20
- Configuring the Global Default Key, page 5-20
- Configuring the Global Default Timeout, page 5-21

The user must configure the remote servers for the TACACS+ protocol. Then the Cisco SCE platform TACACS+ client must be configured to work with the TACACS+ servers. The following information must be configured:

- TACACS+ server hosts definition—A maximum of three servers is supported.

  For each sever host, the following information can be configured:

  - hostname (required)
  - port
  - encryption key
  - timeout interval
- Default encryption key (optional)—A global default encryption key may be defined. This key is defined as the key for any server host for which a key is not explicitly configured when the server host is defined.

  If the default encryption key is not configured, a default of no key is assigned to any server for which a key is not explicitly configured.
- Default timeout interval (optional)—A global default timeout interval may be defined. This timeout interval is defined as the timeout interval for any server host for which a timeout interval is not explicitly configured when the server host is defined.

  If the default timeout interval is not configured, a default of five seconds is assigned to any server for which a timeout interval is not explicitly configured.

The procedures for configuring the Cisco SCE platform TACACS+ client are explained in the following sections:

- Adding a New TACACS+ Server Host, page 5-19
- Removing a TACACS+ Server Host, page 5-20
- Configuring the Global Default Key, page 5-20
- Configuring the Global Default Timeout, page 5-21

Adding a New TACACS+ Server Host

Use this command to define a new TACACS+ server host that is available to the Cisco SCE platform TACACS+ client.

The Service Control solution supports a maximum of three TACACS+ server hosts.

Options

The following options are available:

- **host-name—**Name of the server
- **port number—**TACACS+ port number
  - Default = 49
- **timeout interval—**Time in seconds that the server waits for a reply from the server host before timing out
  - Default = 5 seconds or user-configured global default timeout interval (see "To define the global default timeout, do the following:" section on page 5-21.)
- **key-string—**Encryption key that the server and client will use when communicating with each other. Make sure that the specified key is actually configured on the TACACS+ server host.
  - Default = no key or user-configured global default key (see "To define a global default key, do the following:" section on page 5-21.)

From the SCE(config)# prompt, type:

Command	Purpose
tacacs-server host *host-name* **[port** *portnumber***] [timeout** *timeout-interval***] [key** *key-string***]**	Define a new TACACS+ server host that is available to the Cisco SCE platform TACACS+ client.

Removing a TACACS+ Server Host

Options

The following options are available:

- **host-name—**Name of the server to be deleted

From the SCE(config)# prompt, type:

Command	Purpose
no tacacs-server host *host-name*	Removes a TACACS+ server host.

Configuring the Global Default Key

Use this command to define the global default key for the TACACS+ server hosts. This default key can be overridden for a specific TACACS+ server host by explicitly configuring a different key for that TACACS+ server host.

Options

The following options are available:

- **key-string—**Default encryption key that all TACACS+ servers and clients will use when communicating with each other. Make sure that the specified key is actually configured on the TACACS+ server hosts.
  - Default = no encryption

To define a global default key, do the following:

From the SCE(config)# prompt, type:

Command	Purpose
tacacs-server key *key-string*	Defines the global default key for the TACACS+ server hosts.

To clear a global default key, do the following:

From the SCE(config)# prompt, type:

Command	Purpose
no tacacs-server key	No global default key is defined. Each TACACS+ server host may still have a specific key defined. However, any server host that does not have a key explicitly defined (uses the global default key) is now configured to use no key.

Configuring the Global Default Timeout

Use this command to define the global default timeout interval for the TACACS+ server hosts. This default timeout interval can be overridden for a specific TACACS+ server host by explicitly configuring a different timeout interval for that TACACS+ server host.

Options

The following options are available:

- **timeout interval—**Default time in seconds that the server waits for a reply from the server host before timing out.
  - Default = 5 seconds

To define the global default timeout, do the following:

From the SCE(config)# prompt, type:

Command	Purpose
tacacs-server timeout *timeout-interval*	Defines global default timeout.

To clear the global default timeout, do the following:

From the SCE(config)# prompt, type:

Command	Purpose
no tacacs-server timeout	No global default timeout interval is defined. Each TACACS+ server host may still have a specific timeout interval defined. However, any server host that does not have a timeout interval explicitly defined (uses the global default timeout interval) is now configured to a five second timeout interval.

Managing the User Database

TACACS+ maintains a local user database. Up to 100 users can be configured in this local database, which includes the following information for all users:

- Username
- Password—May be configured as encrypted or unencrypted
- Privilege level

The procedures for managing the local user database are explained in the following sections:

- Adding a New User to the Local Database, page 5-22
- Defining the User Privilege Level, page 5-24
- Adding a New User with Privilege Level and Password, page 5-24
- Deleting a User, page 5-26

Adding a New User to the Local Database

Use these commands to add a new user to the local database. Up to 100 users may be defined.

- How to Add a User with a Clear Text Password, page 5-23
- How to Add a User with No Password, page 5-23
- How to Add a User with an MD5 Encrypted Password Entered in Clear Text, page 5-23
- How to Add a User with an MD5 Encrypted Password Entered as an MD5 Encrypted String, page 5-24

Options

The password is defined with the username. There are several password options:

- No password—Use the **nopassword** keyword.
- Password—Password is saved in clear text format in the local list.

  Use the *password* parameter.

- Encrypted password—Password is saved in encrypted (MD5) form in the local list. Use the secret keyword.

  Password may be defined by either of the following methods:

  - Specify a clear text password, which is saved in MD5 encrypted form
  - Specify an MD5 encryption string, which is saved as the user MD5-encrypted secret password

The following options are available:

- **name—**Name of the user to be added
- **password—**A clear text password. May be saved in the local list in either of two formats:
  - as clear text
  - in MD5 encrypted form if the secret keyword is used
- **encrypted-secret—**An MD5 encryption string password

The following keywords are available:

- **nopassword—**There is no password associated with this user
- **secret—**The password is saved in MD5 encrypted form. Use with either of the following keywords to indicate the format of the password as entered in the command:
  - **0—**Use with the **password** option to specify a clear text password that will be saved in MD5 encrypted form
  - **5—**Use with the **encrypted-secret** option to specify an MD5 encryption string that will be saved as the user MD5-encrypted secret password

How to Add a User with a Clear Text Password

From the SCE(config)# prompt, type:

Command	Purpose
username *name* **password** *password*	Adds a user with a clear text password.

How to Add a User with No Password

From the SCE(config)# prompt, type:

Command	Purpose
username *name* **nopassword**	Adds a user with no password.

How to Add a User with an MD5 Encrypted Password Entered in Clear Text

From the SCE(config)# prompt, type:

Command	Purpose
username *name* **secret 0** *password*	Adds a user with an MD5 encrypted password entered in clear text.

How to Add a User with an MD5 Encrypted Password Entered as an MD5 Encrypted String

From the SCE(config)# prompt, type:

Command	Purpose
username *name* **secret 5** *encrypted-secret*	Adds a user with an MD5 encrypted password entered as an MD5 encrypted string.

Defining the User Privilege Level

Privilege level authorization in the Cisco SCE platform is accomplished by the use of an " **enable** " command authentication request. When a user requests an authorization for a specified privilege level, by using the " **enable** " command, the Cisco SCE platform sends an authentication request to the TACACS+ server specifying the requested privilege level. The Cisco SCE platform grants the requested privilege level only after the TACACS+ server authenticates the " **enable** " command password and verifies that the user has sufficient privileges the enter the requested privilege level.

Options

The following options are available:

- **name—**Name of the user whose privilege level is set
- **level—**The privilege level permitted to the specified user. These levels correspond to the CLI authorization levels, which are entered via the **enable** command:
  - 0—User
  - 10—Admin
  - 15 (default)—Root

From the SCE(config)# prompt, type:

Command	Purpose
username *name* **privilege** *level*	Defines user privilege level.

Adding a New User with Privilege Level and Password

Use these commands to define a new user, including password and privilege level, in a single command.

Note In the config files (**running config** and **startup config**), this command will appear as two separate commands.

- How to Add a User with a Privilege Level and a Clear Text Password, page 5-25
- How to Add a User with a Privilege Level and an MD5 Encrypted Password Entered in Clear Text, page 5-25
- How to Add a User with a Privilege Level and an MD5 Encrypted Password Entered as an MD5 Encrypted String, page 5-26

Options

The following options are available:

- **name—**Name of the user whose privilege level is set
- **level—**The privilege level permitted to the specified user. These levels correspond to the CLI authorization levels, which are entered via the **enable** command:
  - 0—User
  - 10—Admin
  - 15 (default)—Root
- **password—**A clear text password. May be saved in the local list in either of two formats:
  - as clear text I
  - n MD5 encrypted form if the secret keyword is used
- **encrypted-secret—**An MD5 encryption string password

The following keywords are available:

- **secret—**The password is saved in MD5 encrypted form. Use with either of the following keywords to indicate the format of the password as entered in the command:
  - **0—**Use with the **password** option to specify a clear text password that will be saved in MD5 encrypted form
  - **5** = use with the **encrypted-secret** option to specify an MD5 encryption string that will be saved as the user MD5-encrypted secret password

How to Add a User with a Privilege Level and a Clear Text Password

From the SCE(config)# prompt, type:

Command	Purpose
username *name* **privilege** *level* **password** *password*	Adds a user with a privilege level and a clear text password.

How to Add a User with a Privilege Level and an MD5 Encrypted Password Entered in Clear Text

From the SCE(config)# prompt, type:

Command	Purpose
username *name* **privilege** *level* **secret 0** *password*	Adds a user with a privilege level and an MD5 encrypted password entered in clear text.

How to Add a User with a Privilege Level and an MD5 Encrypted Password Entered as an MD5 Encrypted String

From the SCE(config)# prompt, type:

Command	Purpose
username *name* **privilege** *level* **secret 5** *encrypted-secret*	Adds a user with a privilege level and an MD5 encrypted password entered as an MD5 encrypted string.

Deleting a User

Options

The following options are available:

- **name—**Name of the user to be deleted

From the SCE(config)# prompt, type:

Command	Purpose
no username *name*	Deletes a user.

Configuring AAA Login Authentication

There are two features to be configured for login authentication:

- Maximum number of permitted Telnet login attempts
- The authentication methods used at login (see "General AAA Fallback and Recovery Mechanism" section on page 5-17.)

The procedures for configuring login authentication are explained in the following sections:

- Configuring Maximum Login Attempts, page 5-26
- Configuring the AAA Login Authentication Methods, page 5-27

Configuring Maximum Login Attempts

Use this command to set the maximum number of login attempts that will be permitted before the session is terminated.

Options

The following options are available:

- **number-of-attempts—**The maximum number of login attempts that will be permitted before the telnet session is terminated.

  This is relevant only for Telnet sessions. From the local console, the number of re-tries is unlimited.

  - Default = three

From the SCE(config)# prompt, type:

Command	Purpose
aaa authentication attempts login *number-of-attempts*	Configures maximum login attempts.

Configuring the AAA Login Authentication Methods

You can configure "backup" login authentication methods to be used if failure of the primary login authentication method (see "General AAA Fallback and Recovery Mechanism" section on page 5-17).

Use this command to specify which login authentication methods are to be used, and in what order of preference.

- How to Specify the Login Authentication Methods, page 5-27
- How to Delete the Login Authentication Methods List, page 5-27

Options

The following options are available:

- **method—**The login authentication methods to be used. You may specify up to four different methods, in the order in which they are to be used.
  - **group TACACS+—**Use TACACS+ authentication.
  - **local—**Use the local username database for authentication
  - **enable** (default)—Use the " **enable** " password for authentication
  - **none—**Use no authentication.

How to Specify the Login Authentication Methods

From the SCE(config)# prompt, type:

Command	Purpose
aaa authentication login default method1 [method2...]	Specifies login authentication methods. You may list a maximum of four methods; all four methods explained above. List them in the order of priority.

How to Delete the Login Authentication Methods List

From the SCE(config)# prompt, type:

Command	Purpose
no aaa authentication login default	Deletes login authentication methods list. If the login authentication methods list is deleted, the default login authentication method only (enable password) will be used. TACACS+ authentication will not be used.

Configuring AAA Privilege-Level Authorization Methods

- How to Specify AAA Privilege-Level Authorization Methods, page 5-28
- How to Delete the AAA Privilege Level Authorization Methods List, page 5-28

Options

The following options are available:

- **method—**The login authorization methods to be used. You may specify up to four different methods, in the order in which they are to be used.
  - **group TACACS+—**Use TACACS+ authorization.
  - **local—**Use the local username database for authorization
  - **enable** (default)—Use the " **enable** " password for authorization
  - **none—**Use no authorization.

How to Specify AAA Privilege-Level Authorization Methods

From the SCE(config)# prompt, type:

Command	Purpose
aaa authorization enable default method1 [method2...]	Specifies AAA privilege level authorization methods. You may list a maximum of four methods; all four methods explained above. List them in the order of priority.

How to Delete the AAA Privilege Level Authorization Methods List

From the SCE(config)# prompt, type:

Command	Purpose
no aaa authorization enable default	Deletes AAA privilege level authorization methods list. If the privilege level authorization methods list is deleted, the default login authentication method only (enable password) will be used. TACACS+ authentication will not be used.

Configuring AAA Command-Level Authorization Methods

- How to Specify AAA Command-Level Authorization Methods, page 5-29
- How to Delete the AAA Command-Level Authorization Methods List, page 5-29

Options

The following options are available:

- **level—**The privilege level for which to enable the TACACS+ command level authorizarion (0, 5, 10, 15)
- **method—**The command authorization methods to be used. You may specify up to two methods, in the order in which they are to be used.
  - **group TACACS+—**Use TACACS+ authorization.
  - **none—**Use no authorization.

How to Specify AAA Command-Level Authorization Methods

From the SCE(config)# prompt, type:

Command	Purpose
aaa authorization command *level* **default** *method1 [method2]***]**	Specifies AAA command level authorization methods. You may list a maximum of two methods. List them in the order of priority.

How to Delete the AAA Command-Level Authorization Methods List

From the SCE(config)# prompt, type:

Command	Purpose
no aaa authorization command *level* **default**	Deletes AAA command level authorization methods list. If the command level authorization methods list is deleted, the default login authentication method only (enable password) will be used. TACACS+ authentication will not be used.

Configuring AAA Accounting

Use this command to enable or disable TACACS+ accounting.

- How to Enable AAA Accounting, page 5-30
- How to Disable AAA Accounting, page 5-30

If TACACS+ accounting is enabled, the Cisco SCE platform sends an accounting message to the TACACS+ server after every command execution. The accounting message is logged in the TACACS+ server for the use of the network administrator.

By default, TACACS+ accounting is disabled.

Options

The following options are available:

- **level—**The privilege level for which to enable the TACACS+ accounting

How to Enable AAA Accounting

From the SCE(config)# prompt, type:

Command	Purpose
aaa authentication accounting commands *level* **default stop-start group tacacs+**	Enables AAA accounting. The **start-stop** keyword (required) indicates that the accounting message is sent at the beginning and the end (if the command was successfully executed) of the execution of a CLI command.

How to Disable AAA Accounting

From the SCE(config)# prompt, type:

Command	Purpose
aaa authentication accounting commands *level* **default**	Disables AAA accounting.

Monitoring TACACS+

- Displaying Statistics for TACACS+ Servers, page 5-30
- Displaying Statistics, Keys, and Timeouts for TACACS+ Servers, page 5-30
- Monitoring TACACS+ Users, page 5-31

Displaying Statistics for TACACS+ Servers

From the SCE# prompt, type:

Command	Purpose
show tacacs	Displays statistics for TACACS+ servers.

Displaying Statistics, Keys, and Timeouts for TACACS+ Servers

From the SCE# prompt, type:

Command	Purpose
show tacacs all	Displays statistics, keys, and timeouts for TACACS+ servers. Note that, although most show commands are accessible to viewer level users, the ' **all** ' option is available only at the admin level. Use the command ' **enable 10** ' to access the admin level.

Monitoring TACACS+ Users

Use this command to display the users in the local database, including passwords.

From the SCE# prompt, type:

Command	Purpose
show users	Displays the users in the local database, including passwords. Note that, although most show commands are accessible to viewer level users, this command is available only at the admin level. Use the command ' **enable 10** ' to access the admin level.

Configuring Access Control Lists (ACLs)

- Adding Entries to an ACL, page 5-33
- Removing an ACL, page 5-33
- Defining a Global ACL, page 5-34

The Cisco SCE platform can be configured with Access Control Lists (ACLs), which are used to permit or deny incoming connections on the management interface. An access list is an ordered list of entries, each consisting of an IP address and an optional wildcard "mask" defining an IP address range, and a permit/deny field.

The order of the entries in the list is important. The default action of the first entry that matches the connection is used. If no entry in the Access List matches the connection, or if the Access List is empty, the default action is deny.

Configuration of system access is done in two stages:

1. Creating an access list. ("Adding Entries to an ACL" section on page 5-33).
2. Associating the ACL with a management service. (See "Defining a Global ACL" section on page 5-34.)

Creating an access list is done entry by entry, from the first to the last.

When the system checks for an IP address on an access list, the system checks each line in the ACL for the IP address, starting at the first entry and moving towards the last entry. The first match that is detected (that is, the IP address being checked is found within the IP address range defined by the entry) determines the result, according to the permit/deny flag in the matched entry. If no matching entry is found in the ACL, access is denied.

You can create up to 99 ACLs. ACLs can be associated with system access on the following levels:

- Global (IP) level: If a global list is defined using the **ip access-class** command, when a request comes in, the Cisco SCE platform first checks if there is permission for access from that IP address. If not, the Cisco SCE does not respond to the request. Configuring the Cisco SCE platform to deny a certain IP address would preclude the option of communicating with that address using any IP-based protocol including Telnet, FTP, ICMP, RPC, SSH, and SNMP. The basic IP interface is low-level, blocking the IP packets before they reach the interfaces.
- Service level: Access to each management service (Telnet, SNMP, and SSH) can be restricted to an ACL. Interface-level lists are, by definition, a subset of the global list defined. If access is denied at the global level, the IP will not be allowed to access using one of the interfaces. Once an ACL is associated with a specific management service, that service checks the ACL to find out if there is permission for a specific external IP address trying to access the management interface.

  Use the following CLI commands to assign an ACL to the specified management service:

  - Telnet–**access-class in**
  - SSH—**ip ssh access-class**
  - SNMP—**snmp-server community**

It is possible to configure several management services to the same ACL, if this is the desired behavior of the Cisco SCE platform.

If no ACL is associated to a management service or to the global IP level, access is permitted from all IP addresses.

Note The Cisco SCE Platform will respond to **ping** commands only from IP addresses that are allowed access. Pings from a non-authorized address will not receive a response from the Cisco SCE platform, as ping uses ICMP protocol.

Options

The following options are available:

- **number—**The ID number assigned to the Access Control List
- **ip-address—**The IP address of the interface to be permitted or denied. Enter in x.x.x.x format.
- **ip-address/mask—**Configures a range of addresses in the format x.x.x.x y.y.y.y where x.x.x.x specifies the prefix bits common to all IP addresses in the range, and y.y.y.y is a wildcard-bits mask specifying the bits that are ignored. In this notation, '0' means bits to ignore.

The following keywords are available:

- **permit—**The specified IP addresses have permission to access the Cisco SCE platform.
- **deny—**The specified IP addresses are denied access to the Cisco SCE platform.

Adding Entries to an ACL

Step 1 Type **configure** and press **Enter**.

Enables Global Configuration mode.

Step 2 Enter the desired IP address or addresses.

- To configure one IP address type:

  access-list *number* **permit|deny** *ip-address* and press **Enter**.

- To configure more than one IP address type:

  access-list *number* **permit|deny** *ip-address/mask* and press **Enter**.

When you add a new entry to an ACL, it is always added to the end of the list.

Adding Entries to an ACL: Example

The following example adds an entry to the access list number 1, that permits access only to IP addresses in the range of 10.1.1.0–10.1.1.255.

```
SCE(config)#access-list 1 permit 10.1.1.0 0.0.0.255
```

Removing an ACL

Use this command to remove an ACL with all its entries.

From the SCE(config)# prompt, type:

Command	Purpose
no access-list *number*	Removes the specified ACL with all its entries.

Defining a Global ACL

A global ACL for permits or denies all traffic to the Cisco SCE platform.

From the SCE(config)# prompt, type:

Command	Purpose
ip access-class *number*	Applies the specified ACL to all traffic attempting to access the Cisco SCE platform, rather than to a specific type of traffic, such as Telnet traffic.

Managing the Telnet Interface

- Preventing Telnet Access, page 5-35
- Assigning an ACL to the Telnet Interface, page 5-35
- Configuring Telnet Timeout, page 5-36

This section discusses the Telnet interface of the Cisco SCE platform. A Telnet session is the most common way to connect to the Cisco SCE platform CLI interface.

You can set the following parameters for the Telnet interface:

- Enable/disable the interface
- Assign an ACL to permit or deny incoming connections.
- Timeout for Telnet sessions, that is, if there is no activity on the session, how long the Cisco SCE platform waits before automatically cutting off the Telnet connection.

The following commands are relevant to Telnet interface:

- **access-class in**
- **line vty**
- **[no] access list**
- **[no] service telnetd**
- **[no] timeout**
- **show line vty access-class in**
- **show line vty timeout**

Preventing Telnet Access

Use this command to disable access by Telnet altogether.

From the SCE(config)# prompt, type:

Command	Purpose
no service telnetd	Disables access by Telnet. Current Telnet sessions are not disconnected, but no new Telnet sessions are allowed.

Assigning an ACL to the Telnet Interface

From the SCE(config-line)# prompt, type:

Command	Purpose
access-class *acl-number* **in**	The specified ACL controls access to the Telnet interface. *acl-number* is the ID number of an existing ACL.

Assigning an ACL to the Telnet Interface: Example

The following example shows how to assign ACL #1 to the Telnet interface.

```
SCE#configure
SCE(config)#line vty 0
SCE(config-line)#access-class 1 in
```

Removing ACL Assignment from the Telnet Interface

From the SCE(config-line)# prompt, type:

Command	Purpose
no access-class in	Removes the ACL assignment from the Telnet interface, so that any IP address may now access the Telnet interface.

Configuring Telnet Timeout

The Cisco SCE platform supports timeout of inactive Telnet sessions.

Options

The following options are available:

- **timeout—**The length of time in minutes before an inactive Telnet session will be timed-out.
  - Default—30 minutes

From the SCE(config-line)# prompt, type:

Command	Purpose
timeout *timeout*	Configures Telnet timeout.

Configuring the SSH Server

- The SSH Server, page 5-37
- Key Management, page 5-37
- Managing the SSH Server, page 5-38
- Monitoring the Status of the SSH Server, page 5-40

The SSH Server

A shortcoming of the standard telnet protocol is that it transfers password and data over the net unencrypted, thus compromising security. Where security is a concern, using a Secure Shell (SSH) server rather than telnet is recommended.

An SSH server is similar to a telnet server, but it uses cryptographic techniques that allow it to communicate with any SSH client over an insecure network in a manner which ensures the privacy of the communication. CLI commands are executed over SSH in exactly the same manner as over telnet.

The SSH server supports both the SSHv1 and SSHv2 protocols. You can disable SSHv1, so that only SSHv2 is running.

The SSH server supports the following encryption ciphers:

- aes256-ctr, aes192-ctr, aes128-ctr (RFC-4344, section 4).
- 3des-cbc, blowfish-cbc, aes256-cbc, aes192-cbc, aes128-cbc, arcfour, cast128-cbc (RFC-4253, section 6.3)
- arcfour128, arcfour256 (RFC-4345, section 4).
- rijndael-cbc@lysator.liu.se (as provided by OpenSSH 4.7p1).

An ACL can be configured for SSH as for any other management protocol, limiting SSH access to a specific set of IP addresses (see "Configuring Access Control Lists (ACLs)" section on page 5-32)

Key Management

Each SSH server should define a set of keys (DSA2, RSA2 and RSA1) to be used when communicating with various clients. The key sets are pairs of public and private keys. The server publishes the public key while keeping the private key in non-volatile memory, never transmitting it to SSH clients. Note that the keys are kept on the tffs0 file system, which means that a person with knowledge of the 'enable' password can access both the private and public keys. The SSH server implementation provides protection against eavesdroppers who can monitor the management communication channels of the Cisco SCE platform, but it does not provide protection against a user with knowledge of the 'enable' password.

Key management is performed by the user via a special CLI command. A set of keys must be generated at least once before enabling the SSH server.

Size of the encryption key is always 2048 bits.

Managing the SSH Server

Use these commands to manage the SSH server. These commands do the following:

- Generate an SSH key set
- Enable/disable the SSH server
- Enable/disable SSHv1. (Disabling SSHv1 allows you to run SSHv2 only.)
- Delete existing SSH keys

Generating a Set of SSH Keys

Remember that you must generate a set of SSH keys before you enable the SSH server.

From the SCE(config)# prompt, type:

Command	Purpose
ip ssh key generate	Generates a new SSH key set and immediately saves it to non-volatile memory. (Key set is not part of the configuration file). Key size is always 2048 bits.

Enabling the SSH Server

SSH allows you to login only when the user password and AAA authentication are configured.

- Configure at least one user name and password.

  SCE8000(config)# username <username> password <password>

- Configure AAA authentication for login.

  SCE8000(config)# aaa authentication login default none

From the SCE(config)# prompt, type:

Command	Purpose
ip ssh	Enables SSH server.

Disabling the SSH Server

From the SCE(config)# prompt, type:

Command	Purpose
no ip ssh	Disables SSH server.

Running Only SSHv2

Step 1 From the SCE(config)# prompt, type **ip ssh** and press **Enter**.

Step 2 From the SCE(config)# prompt, type **no ip ssh sshv1** and press **Enter**

To re-enable SSHv1, use the command **ip ssh SSHv1.**

Assigning an ACL to the SSH Server

From the SCE(config)# prompt, type:

Command	Purpose
ip ssh access-class *acl-number*	The specified ACL controls access to the SSH server. *acl-number* is the ID number of an existing ACL

Removing the ACL Assignment from the SSH Server

rom the SCE(config)# prompt, type:

Command	Purpose
no ip ssh access-class	Removes the ACL assignment from the SSH server, so that any IP address may now access the SSH server.

Deleting the Existing SSH Keys

From the SCE(config)# prompt, type:

Command	Purpose
ip ssh key remove	Removes the existing SSH key set from non-volatile memory.

If the SSH server is currently enabled, it will continue to run, since it only reads the keys from non-volatile memory when it is started. However, if the startup-configuration specifies that the SSH server is enabled, the Cisco SCE platform will not be able to start the SSH server on startup if the keys have been deleted. To avoid this situation, after executing this command, always do one of the following before the Cisco SCE platform is restarted (using **reload**):

- Generate a new set of keys.
- Disable the SSH server and save the configuration.

Monitoring the Status of the SSH Server

Use this command to monitor the status of the SSH sever, including current SSH sessions.

From the SCE> prompt, type:

Command	Purpose
show ip ssh	Monitors the status of SSH server.

Configuring and Managing the SNMP Interface

- About the SNMP Interface, page 5-41
- Enabling the SNMP Interface, page 5-45
- Configuring SNMP Community Strings, page 5-45
- Configuring SNMP Notifications, page 5-47

About the SNMP Interface

This section explains how to configure the SNMP agent parameters. It also provides a brief overview of SNMP notifications and relevant CLI commands.

- SNMP Protocol, page 5-41
- Security Considerations, page 5-42
- About CLI, page 5-43
- About MIBs, page 5-44
- Configuration via SNMP, page 5-44

SNMP Protocol

SNMP (Simple Network Management Protocol) is a set of protocols for managing complex networks. SNMP works by sending messages, called protocol data units (PDUs), to different parts of a network. SNMP-compliant devices, called agents, store data about themselves in Management Information Bases (MIBs) and return this data to the SNMP requesters.

Cisco SCE platform supports the original SNMP protocol (also known as SNMPv1), Community-based SNMPv2 (also known as SNMPv2C), and SNMPv3.

- **SNMPv1**—Is the first version of the Simple Network Management Protocol, as defined in RFCs 1155 and 1157, and is a full Internet standard. SNMPv1 uses a community-based form of security.
- **SNMPv2c**—The revised version of SNMPv1, includes improvements to SNMPv1 in the areas of protocol packet types, transport mappings, and MIB structure elements but using the existing SNMPv1 administration structure. It is defined in RFC 1901, RFC 1905, and RFC 1906.
- **SNMPv3**—This version mainly addresses the deficiencies related to security and administration. SNMPv3 indroduced advanced security features that split the authentication and authorization into two separate functions. The User-based Security Model (USM) is the default security model. USM and its attributes are described in RFC 2574.

Cisco SCE platform implementation of SNMP supports all MIB II variables, as described in RFC 1213, and defines the SNMP traps using the guidelines described in RFC 1215.

The SNMPv1, SNMPv2C, and SNMPv3 specifications define the following basic operations that are supported by Cisco SCE platform. Table 5-2 Table 5-2 lists the request types and descriptions.

Table 5-2 Request Types

Request Type	Description	Remarks
Set Request	Writes new data to one or more of the objects managed by an agent.	Set operations immediately affect the Cisco SCE platform running-config but do not affect the startup config
Get Request	Requests the value of one or more of the objects managed by an agent.	
Get Next Request	Requests the Object Identifier(s) and value(s) of the next object(s) managed by an agent.	
Get Response	Contains the data returned by an agent.	
Trap	Sends an unsolicited notification from an agent to a manager, indicating that an event or error has occurred on the agent system	Cisco SCE platform may be configured to send either SNMPv1 or SNMPv2 style traps.
Get Bulk Request	Retrieves large amounts of object information in a single Request / response transaction. GetBulk behaves as if many iterations of GetNext request/responses were issued, except that they are all performed in a single request/response.	This is an SNMPv2c message.

Security Considerations

By default, the SNMP agent is disabled for both read and write operations. When enabled, SNMP is supported over the management port only (in-band management is not supported).

In addition, the Cisco SCE platform supports the option to configure community of managers for read-write accessibility or for read-only accessibility. Furthermore, an ACL may be associated with the SNMP agent by assigning it to one of the community strings to allow SNMP management to a restricted set of manager IP addresses. If different ACLs are assigned to different community strings, access by all community strings is controlled by all assigned ACLs. Assigning different ACLs to different community strings is not supported.

When using SNMPv3, Each user is assigned a name (called a securityName), an authentication type (authProtocol), and a privacy type (privProtocol) as well as associated keys for each of these (authKey and privKey). Authentication is performed by using the authKey of the user to sign the message being sent. The authProtocol can be either in MD5 or SHA format. authKeys (and privKeys) are generated from a passphrase that is at least 8 characters in length. Encryption is performed by using the privKey of a user to encrypt the data portion of the message being sent. The privacy Protocol can be AES or DES. Messages can be sent unauthenticated and unencrypted (noAuthNoPriv), authenticated but unencrypted (authNoPriv), or authenticated and encrypted (authPriv) by setting the securityLevel.

You can create VIEW and GROUPS for access control. A VIEW is a set of MIBs or OIDs, defined by including or excluding different MIBs and OIDs. A view defines the range of OIDs that are exposed in a specific view for a user. A GROUP is defined with a particular security level, read-view, and

write-view. All users in a group have common access rights. A read-view defines a view, containing the OIDs available for the user to read (SNMPGET/GETNEXT), and similarly a write-view contains the OIDs available for update (SNMPSET)

Note SNMPv3 configurations are available only through Command Line Interface.

About CLI

- CLI Commands for Configuring SNMP, page 5-43
- CLI Commands for Monitoring SNMP, page 5-44

The Cisco SCE platform supports the CLI commands that control the operation of the SNMP agent. All the SNMP commands are available in Admin authorization level. The SNMP agent is disabled by default and any SNMP configuration command enables the SNMP agent (except where there is an explicit disable command).

CLI Commands for Configuring SNMP

Following is a list of CLI commands available for configuring SNMPv1 and SNMPv2c. These are Global Configuration mode commands.

- **snmp-server enable**
- **no snmp-server**
- **[no] snmp-server community [all]**
- **[no | default] snmp-server enable traps**
- **[no] snmp-server host [all]**
- **[no] snmp-server contact**
- **[no] snmp-server location**

Following is a list of CLI commands available for configuring SNMPv3. These are Global Configuration mode commands.

- **snmp-server enable**
- **no snmp-server**
- **snmp-server view <view-name> OID <oidOrMIBName> operation {included | excluded}**
- **no snmp-server view all**
- **no snmp-server view view-name OID <oidOrMIBName> operation {included | excluded}**
- **snmp-server group <group-name> 3 { noAuthNoPriv | authNoPriv | authPriv } [read-view <view-name>] [write-view <view-name>]**
- **no snmp-server group all**
- **no snmp-server group <group-name>**
- **snmp-server user <user-name> group <group-name> [auth {MD5-auth | SHA-auth} auth-pass-phrase <authPassPhrase>] [priv {AES-priv | DES-priv} priv-pass-phrase <privPassPhrase>]**
- **no snmp-server user all**
- **no snmp-server user <user-name>**

- **snmp-server host {<hostname> |< ip-address>} [traps] version 3 {noAuthNoPriv | authNoPriv | authPriv} user <user-name>**

CLI Commands for Monitoring SNMP

Following is a list of CLI commands available for monitoring SNMP. These are Viewer mode commands, and are available when the SNMP agent is enabled:

- **show snmp** (available when SNMP agent is disabled)
- **show snmp community**
- **show snmp contact**
- **show snmp enabled**
- **show snmp host**
- **show snmp location**
- **show snmp MIB** (available when SNMP agent enabled and community was set)
- **show snmp traps**

Following is a list of CLI commands available for monitoring SNMPv3. These are Viewer mode commands, and are available when the SNMP agent is enabled:

- **show snmp view {all-Views | { view-name <view-name>}}**
- **show snmp group {all-Groups | { group-name < group-name >}}**
- **show snmp user {all-Users | { user-name < user-name >}}**
- **show snmp engine-id**

About MIBs

MIBs (Management Information Bases) are databases of objects that can be monitored by a network management system (NMS). SNMP uses standardized MIB formats that allow any SNMP tools to monitor any device defined by a MIB.

For further information concerning MIBs used by the Cisco SCE 8000 platform, see the "Cisco Service Control MIBs" section on page A-1

Configuration via SNMP

Cisco SCE platform supports a limited set of variables that may be configured via SNMP (read-write variables). Setting a variable via SNMP (as via the CLI) takes effect immediately and affects only the running-configuration. To make this configuration stored for next reboots (startup-configuration) the user must specify it explicitly via CLI or via SNMP using the Cisco enterprise MIB objects.

The Cisco SCE platform takes the approach of a single configuration database with multiple interfaces that may change this database. Therefore, executing the **copy running-config startup-config** command via CLI or SNMP makes permanent all the changes made by either SNMP or CLI.

The SNMPv3 user configurations are persistent across reboots. The configurations are not part of the running-configuration, hence these configurations are not part of the startup-configuration.

Enabling the SNMP Interface

Use this command to explicitly enable the SNMP interface.

The SNMP interface is implicitly enabled when any **snmp-server** command is executed to configure any SNMP parameter. For more information on configuring and managing the SNMP parameters, including hosts, communities, contact, location, and trap destination hosts, see "Configuring and Managing the SNMP Interface" section on page 5-41.

- How to Enable the SNMP Interface, page 5-45
- How to Disable the SNMP Interface, page 5-45

How to Enable the SNMP Interface

You must define at least one community string to allow SNMP access. For complete information on community strings see "Configuring SNMP Community Strings" section on page 5-45.

From the SCE(config)# prompt, type:

Command	Purpose
snmp-server enable	Enables SNMP interface.

How to Disable the SNMP Interface

From the SCE(config)# prompt, type:

Command	Purpose
no snmp-server	Disables SNMP interface.

Configuring SNMP Community Strings

- Defining a Community String, page 5-46
- Removing a Community String, page 5-46
- Displaying the Configured Community Strings, page 5-47

To enable SNMP management, you must configure SNMP community strings to define the relationship between the SNMP manager and the agent.

After receiving an SNMP request, the SNMP agent compares the community string in the request to the community strings that are configured for the agent. The requests are valid under the following circumstances:

- SNMP *Get* , *Get-next* , and *Get-bulk* requests are valid if the community string in the request matches the read-only community.
- SNMP *Get* , *Get-next* , *Get-bulk* and *Set* requests are valid if the community string in the request matches the agent's read-write community.

Defining a Community String

Options

The following options are available:

- **community-string—**A security string that identifies a community of managers who are permitted to access the SNMP server
- **acl-number**—(Optional) ID number of the ACL to be assigned to the SNMP interface. It should list the IP addresses of all SNMP managers permitted to access the SNMP server.

Note Assigning different ACLs to different community strings is not supported. If you specify an ACL in this command, it is assigned to the SNMP server globally, not just to the specified community string. For example, if you configure two community strings and assign a different ACL to each, access to the SNMP agent for both communities is controlled by both ACLs.

If no ACL is specified, all IP addresses can access the SNMP service. For more information about ACLs, see "Configuring Access Control Lists (ACLs)" section on page 5-32.

The following keywords are available:

- **ro—**Read only (default accessibility)
- **rw**—Read and write

From the SCE(config)# prompt, type:

Command	Purpose
snmp-server community *community-string* **ro|rw** [*acl-number*]	Defines a community string. If you specify ACLs for any communities, all assigned ACLs in conjunction will control access for all communities. Repeat the command as necessary to define all community strings.

Defining a Community String: Example

This example shows how to configure a community string called "mycommunity" with read-only rights. ACL "1" will be assigned to the SNMP server and all configured community strings, not just "mycommunity".

Since read-only is the default, it does not need to be defined explicitly.

```
SCE(config)#snmp-server community mycommunity 1
```

Removing a Community String

From the SCE(config)# prompt, type:

Command	Purpose
no snmp-server community *community-string*	Removes a community string.

Removing a Community String: Example

The following example shows how to remove a community string called “mycommunity”. If an ACL was assigned to this community string, it is also removed.

```
SCE(config)#no snmp-server community mycommunity
```

Displaying the Configured Community Strings

From the SCE> prompt, type:

Command	Purpose
show snmp-server community *community-string*	Displays configured community strings

Displaying the Configured Community Strings: Example

The following example shows how to display the configured SNMP communities.

```
SCE>show snmp community
Community: public, Access Authorization: RO, Access List Index: 1
SCE>
```

Configuring SNMP Notifications

Use these commands to configure:

- The destinations that will receive SNMP notifications (hosts)
- Which types of notifications will be sent (traps)

Notifications are unsolicited messages that are generated by the SNMP agent that resides inside the Cisco SCE platform when an event occurs. When the Network Management System receives the notification message, it can take suitable actions, such as logging the occurrence or ignoring the signal.

By default, the Cisco SCE platform is not configured to send any SNMP notifications. You must define the Network Management System to which the Cisco SCE platform should send notifications. (See the table below, Configurable Notifications, for a list of configurable notifications). Whenever one of the events that trigger notifications occurs in the Cisco SCE platform, an SNMP notification is sent from the Cisco SCE platform to the list of IP addresses that you define.

The Cisco SCE platform supports two general categories of notifications:

- Standard SNMP notifications - As defined in RFC1157 and using the conventions defined in RFC1215.
- Proprietary service control enterprise notifications - As defined in the service control proprietary MIB (See Table A-20 on page A-21).

After a host or hosts are configured to receive notifications, by default, the Cisco SCE platform sends to the host or hosts all the notifications supported by the Cisco SCE platform except for the AuthenticationFailure notification. The Cisco SCE platform provides the option to enable or disable the sending of this notification, as well as some of the Cisco SCE enterprise notifications, explicitly.

Cisco SCE platform can be configured to generate either SNMPv1 style, SNMPv2c style, or SNMPv3 style notifications. By default, the Cisco SCE platforms sends SNMPv1 notifications.

Following are some sample procedures illustrating how to do the following:

- (SNMPv3 only) Configure SNMP Server View
- (SNMPv3 only) Configure SNMP Server Group
- (SNMPv3 only) Configuration SNMP Server User
- Configure hosts (NMS) to which the SNMP agent should send notifications
- Remove/disable a host (NMS) from receiving notifications
- Enable the SNMP agent to send authentication-failure notifications
- Enable the SNMP agent to send enterprise notifications
- Reset all notifications to the default setting

Configuring SNMP Server Group

Use this command to configure the SNMPv3 server group on the Cisco SCE platform.

At the SCE(config)# prompt, type:

Command	Purpose
snmp-server group *groupname* **3** { **noAuthNoPriv**| **authNoPriv** | **authPriv**}} [**read-view** *view-name*] [**write-view** *write-view*]	Configure the SNMPv3 server group on the Cisco SCE platform. The server group details are included in the running-configuration. By default an SNMP group is read-only.

Options

The following options are available:

- **groupname**—Name of the server group.
- **noAuthNoPriv**—Sets the security level to no authentication and no privacy.Group does not require any authentication or encryption security levels for access.
- **authNoPriv**—Sets the security level with authentication and no privacy. The users of this group has to authenticate themselves to get access.
- **authPriv**—Sets the security level with authentication and privacy. The users of this group requires both authentication and encryption security levels for access.
- **read-view**—(Optional)Specifies that the view is a read-view. This is the view used for SNMPGET/GET-NEXT/WALK.
- **write-view**—(Optional)Specifies that the view is a write-view. This is the view used for SNMPSET.

Configuring SNMP Server View

Use this command to configure the SNMP v3 server view on the Cisco SCE platform.

At the SCE(config)# prompt, type:

Command	Purpose
snmp-server view *view-name* **3** OID *OID or MIB-name* **Operation {Included | Excluded}**	Configure the SNMPv3 server view on the Cisco SCE platform.

Options

The following options are available:

- **view-name**—Name of the view.
- **OID**—Specify the OID or MIB name.
- **Operation**—Sets the the type of the entry for view.
- **Included**—Specifies to include all entries from the MIB/OID name to the view.
- **Excluded**—Specifies to exclude the MIB/OID name from the view.

Note The OID or the MIB name is read and is considerd as plain text.

Configuring SNMP Server User

Use this command to configure the SNMP v3 server user on the Cisco SCE platform.

To configure large number of SNMPv3 users, disable SNMP agent before configuring the users. Enable the SNMP agent after configuring all users.

To configure less number of SNMPv3 user or few adhoc configurations, you need not disable the SNMP agent. You may notice a delay of few seconds for the configuration to take effect.

At the SCE(config)# prompt, type:

Command	Purpose
snmp-server user *user-name* **group** *group-name* **[auth {MD5-auth | SHA-auth} auth-pass-phrase** *authPassPhrase***] [priv {***AES-priv* **|** *DES-priv***} priv-pass-phrase** *privPassPhrase***]**	Configure the SNMP v3 server user on the Cisco SCE platform.

Options

The following options are available:

- **user-name**—Sets the SNMP user name.
- **group-name**—Specifies the groups to add the user.
- **auth**—Sets the type of authentication. MD5 and SHA authentication types are available.
- **auth-pass-phrase**—Sets the authentication pass phrase.
- **priv**—Sets the type of privacy. AES and DES privacy types are available.
- **priv-pass-phrase**—Sets the privacy pass phrase.

SNMPv3 users are not shown in the running-configuration. To see a list of configured users, use the **show snmp user** command.

Defining SNMP Hosts

Use this command to define the hosts that will receive notifications from the Cisco SCE platform.

- How to Configure the Cisco SCE Platform to Send Notifications to a Host (NMS), page 5-50
- How to Configure the Cisco SCE Platform to Stop Sending Notifications to a Host, page 5-50
- Configuring SNMP Traps, page 5-51

Options

The following options are available:

- **hostname**—Name of the SNMP notification host
- **ip-address**—IP address of the SNMP notification host
- **community-string**—Security string that identifies a community of managers who are permitted to access the SNMP server
- **version**—SNMP version running in the system. Can be set to 1, 2c, or 3. The default is 1 (SNMPv1).
- **noAuthNoPriv**—(SNMPv3 only)Sets the security level to no authentication and no privacy. Trap does not require any authentication or encryption security levels for access.
- **authNoPriv**—(SNMPv3 only)Sets the security level with authentication and no privacy. The users of this trap has to authenticate themselves to get access.
- **authPriv**—(SNMPv3 only)Sets the security level with authentication and privacy. The users of this trap requires both authentication and encryption security levels for access.

How to Configure the Cisco SCE Platform to Send Notifications to a Host (NMS)

At the SCE(config)# prompt, type:

Command	Purpose
snmp-server host { *hostname* | *ip-address* } [**traps**] [**version** *version*] {*community-string* | { **noAuthNoPriv**| **authNoPriv** | **authPriv**}} **user** *user-name*	Configures Cisco SCE platform to send notifications to a host (NMS). If the version is not specified, SNMPv1 is assumed. Only one host can be specified per command. To define multiple hosts, execute one command for each host.

Configuring the Cisco SCE Platform to Send Notifications to Multiple Hosts: Example

The following example shows how to configure the Cisco SCE platform to send SNMPv1 notifications to one host, SNMPv2c notifications to another, and SNMPv3 notification to a third host:

```
SCE(config)# snmp-server host myhost.cisco.com traps public
SCE(config)# snmp-server host 20.20.20.20 version 2c mycommunity
SCE(config)# snmp-server host host 10.1.1.205 version 3 authPriv user user1
```

How to Configure the Cisco SCE Platform to Stop Sending Notifications to a Host

At the SCE(config)# prompt, type:

Command	Purpose
no snmp-server host *ip-address*	Configures Cisco SCE platform to stop sending notifications to a host.

Configuring the Cisco SCE Platform to Stop Sending Notifications to a Host: Example

The following example shows how to remove the host with the IP Address: “192.168.0.83”.

```
SCE(config)#no snmp-server host 192.168.0.83
```

Configuring SNMP Traps

Use this command to configure the notifications that will be sent to the defined host.

- How to Enable the SNMP Server to Send Authentication Failure Notifications, page 5-51
- How to Enable the SNMP Server to Send All Enterprise Notifications, page 5-52
- How to Enable the SNMP Server to Send a specific Enterprise Notification, page 5-52
- How to Restore All Notifications to the Default Status, page 5-52

Options

The following options are available:

- **snmp—**Optional parameter that specifies that all or specific snmp traps should be enabled or disabled.

  By default, snmp traps are disabled.

  snmp trap name—Optional parameter that specifies a specific snmp trap that should be enabled or disabled.

  Currently the only accepted value for this parameter is **Authentication** .

- **enterprise—**Optional parameter that specifies that all or specific enterprise traps should be enabled or disabled.

  By default, enterprise traps are enabled.

- **enterprise trap name—**Optional parameter that specifies a specific snmp trap that should be enabled or disabled.

  Values: attack, chassis, link-bypass, logger, operational-status, port-operational-status, pull-request-failure, RDR-formatter, session, SNTP, subscriber, system-reset, telnet, vas-traffic-forwarding

Use these parameters as follows:

- To enable/disable all traps of one type: Specify only **snmp** or **enterprise** .
- To enable/disable only one specific trap: Specify **snmp** or **enterprise** with the additional trap name parameter naming the desired trap.
- To enable/disable all traps: Do not specify either **snmp** or **enterprise** .

How to Enable the SNMP Server to Send Authentication Failure Notifications

At the SCE(config)# prompt, type:

Command	Purpose
snmp-server enable traps snmp authentication	Enables SNMP server to send authentication failure notifications.

How to Enable the SNMP Server to Send All Enterprise Notifications

At the SCE(config)# prompt, type:

Command	Purpose
snmp-server enable traps enterprise	Enables SNMP server to send all enterprise notifications.

How to Enable the SNMP Server to Send a specific Enterprise Notification

At the SCE(config)# prompt, type:

Command	Purpose
snmp-server enable traps enterprise *[attack|chassis|link-bypass|logger|operational-status|port-operational-status|pull-request-failure|RDR-formatter|session|SNTP|subscriber|system-reset|telnet|vas-traffic-forwarding]*	Specify the desired enterprise trap type.

Enabling the SNMP Server to Send a Specific Enterprise Notification: Example

The following example shows how to configure the SNMP server to send the logger enterprise notification only.

```
SCE(config)#snmp-server enable traps enterprise logger
```

How to Restore All Notifications to the Default Status

At the SCE(config)# prompt, type:

Command	Purpose
default snmp-server enable traps	Resets all notifications supported by the Cisco SCE platform to their default status.

SNMP Walk Acceleration for linkServiceUsage Queries

The time taken for the SNMP walk on any of the linkServiceUsage queries is reduced considerably.

The SNMP walk acceleration enables Cisco SCE 8000 device to perform SNMP queries for LinkUsage MIB queries in background and cache the results. This may result in more CPU utilization.

The following queries are accelerated:

- snmpbulkwalk on linkServiceUsageUpVolume
- snmpbulkwalk on linkServiceUsageDownVolume
- snmpbulkwalk on linkServiceUsageNumSessions
- snmpbulkwalk on linkServiceUsageDuration
- snmpbulkwalk on linkServiceUsageConcurrentSessions
- snmpbulkwalk on linkServiceUsageActiveSubscribers
- snmpwalk on complete linkServiceUsageTable

Use the following commands to enable, disable, and show the SNMP query acceleration for linkServiceUsage queries:

- **snmp-server accelerate-query**
- **show snmp-server accelerate-query**

How to Enable SNMP Query Acceleration

From the SCE(config)#> prompt, type:

Command	Purpose
snmp-server accelerate-query <*linkServiceUsage query*>	Enables the SNMP query acceleration.

SNMPv3 Configuration Example

```
SCE8000#> configure
SCE8000(config)#> snmp-server view IPView OID 1.3.6.1.2.1.4 operation include
SCE8000(config)#> snmp-server view IfView OID 1.3.6.1.2.1.2.2 operation include
SCE8000(config)#> snmp-server group ipGroup 3 authNoPriv read-view IPView
SCE8000(config)#> snmp-server group ifGroup 3 authPriv read-view IfView write-view IfView
SCE8000(config)#> snmp-server group ifGroupReadOnly 3 noAuthNoPriv read-view IfView
SCE8000(config)#> snmp-server user ipUser group ipGroup auth MD5-auth auth-pass-phrase
password
Note: Do not forget to enable the SNMP agent after all the user configuration, to make
them active.
SCE8000(config)#> snmp-server user ipUser01 group ipGroup auth SHA-auth auth-pass-phrase
passwordsha
Note: Do not forget to enable the SNMP agent after all the user configuration, to make
them active.
SCE8000(config)#> snmp-server user ifUserNoAuth group ifGroupReadOnly
Note: Do not forget to enable the SNMP agent after all the user configuration, to make
them active.
SCE8000(config)#> snmp-server user ifUserRW group ifGroup auth SHA-auth auth-pass-phras
passwordsha priv AES-priv priv-pass-phrase passwordaes
Note: Do not forget to enable the SNMP agent after all the user configuration, to make
them active.
```

```
SCE8000(config)#> snmp-server enable

SCE8000(config)# >end

SCE8000#> show snmp view all-Views

List of configured Views:-
  View: full-view - include OID: 1
  View: none-view - exclude OID: 1
  View: IPView - include OID: 1.3.6.1.2.1.4
  View: IfView - include OID: 1.3.6.1.2.1.2.2

SCE8000#> show snmp view IPView
Error -  invalid input     ^

SCE8000#> show snmp view view-name IPView
View: IPView -
  include OID: 1.3.6.1.2.1.4

SCE8000#> show snmp group all-Groups

List of configured Groups:-
  Group Name: ipGroup :
    Security Level: authNoPriv
    read view: IPView
    write view: none-view

  Group Name: ifGroup :
    Security Level: authPriv
    read view: IfView
    write view: IfView

  Group Name: ifGroupReadOnly :
    Security Level: noAuthNoPriv
    read view: IfView
    write view: none-view

SCE8000#> show snmp group group-name ifGroup

Group Name: ifGroup :
  Security Level: authPriv
  read view: IfView
  write view: IfView
  Users belonging to this group:
        ifUserRW

SCE8000#>sh snmp user all-Users
User: ipUser :
  Group Name: ipGroup
  Authentication Protocol: MD5
  Privacy protocol: NONE

User: ipUser01 :
  Group Name: ipGroup
  Authentication Protocol: SHA
  Privacy protocol: NONE

User: ifUserNoAuth :
  Group Name: ifGroupReadOnly
  Authentication Protocol: NONE
  Privacy protocol: NONE

User: ifUserRW :
  Group Name: ifGroup
```

```
  Authentication Protocol: SHA
  Privacy protocol: AES

SCE8000#> show snmp user user-name ipUser
User: ipUser :
  Group Name: ipGroup
  Authentication Protocol: MD5
  Privacy protocol: NONE
```



Global Configuration

Revised: September 23, 2015,

Introduction

This chapter explains how to perform global configuration tasks, including IP routing and clock and time zone settings.

- IP Routing Configuration, page 6-2
- Configuring Time Clocks and Time Zone, page 6-6
- Configuring SNTP, page 6-13
- Domain Name Server (DNS) Settings, page 6-17
- Configuring Cisco Discovery Protocol, page 6-20
- Enabling the CLI Interface Warning Banner, page 6-29
- OS Fingerprinting and NAT Detection, page 6-30
- Using the Bursty Traffic Convergence, page 6-33
- DNS Assisted Classification/Sampling, page 6-36

IP Routing Configuration

- Configuring the IP Routing Table, page 6-2
- IP Advertising, page 6-4

Configuring the IP Routing Table

- How to Configure the Default Gateway, page 6-2
- How to Add an Entry to the IP Routing Table, page 6-3
- How to Display the IP Routing Table, page 6-3

For handling IP packets on the out-of-band MNG port, the Cisco SCE platform maintains a static routing table. When a packet is sent, the system checks the routing table for proper routing, and forwards the packet accordingly. In cases where the Cisco SCE platform cannot determine where to route a packet, it sends the packet to the default gateway.

The Cisco SCE platform supports the configuration of the default gateway as the default next hop router, as well as the configuration of the routing table to provide different next hop routers for different subnets (for maximum configuration of 100 subnets).

The following sections illustrate how to use CLI commands to configure various parameters.

The following commands are relevant to IP routing tables:

- **ip default-gateway**
- **ip route prefix mask next-hop**
- **no ip route all**
- **no ip route prefix mask**
- **show ip route**
- **show ip route prefix**
- **show ip route prefix mask**

How to Configure the Default Gateway

Options

The following option is available:

- ip-address—the IP address of the default gateway.

From the SCE(config)# prompt, type:

Command	Purpose
ip default-gateway *ip-address*	Configures the default gateway.

Configuring the Default Gateway: Example

The following example shows how to set the default gateway IP of the Cisco SCE platform to 10.1.1.1.

```
SCE(config)#ip default-gateway 10.1.1.1
```

How to Add an Entry to the IP Routing Table

Options

The following options are available:

- **prefix—**IP address of the routing entry, in dotted notation.
- **mask—**The relevant subnet mask, in dotted notation
- **next-hop—**The IP address of the next hop in the route, in dotted notation.

  Must be within the MNG interface subnet.

From the SCE(config)# prompt, type:

Command	Purpose
ip route *prefix mask next-hop*	Adds the specified IP routing entry to the routing table.

How to Add an Entry to the IP Routing Table: Example

The following example shows how to set the router 10.1.1.250 as the next hop to subnet 10.2.0.0.

```
SCE(config)#ip route 10.2.0.0 255.255.0.0 10.1.1.250
```

How to Display the IP Routing Table

- How to Display the Entire IP Routing Table, page 6-3
- How to Display the IP Routing Table for a Specified Subnet, page 6-4

How to Display the Entire IP Routing Table

From the SCE# prompt, type:

Command	Purpose
show ip route	Displays the entire routing table and the destination of last resort (default-gateway).

Displaying the Entire IP Routing Table: Example

This example shows how to display the routing table.

```
SCE#show ip route
gateway of last resort is        10.1.1.1
|    prefix       |     mask         |    next hop     |
|-----------------|------------------|-----------------|
|        10.2.0.0 |     255.255.0.0  |      10.1.1.250 |
|        10.3.0.0 |     255.255.0.0  |      10.1.1.253 |
|        198.0.0.0 |       255.0.0.0 |      10.1.1.251 |
|        10.1.60.0 |   255.255.255.0 |      10.1.1.5   |
```

How to Display the IP Routing Table for a Specified Subnet

Options

The following options are available:

- **prefix—**IP address of the routing entry, in dotted notation.
- **mask—**The relevant subnet mask, in dotted notation

From the SCE# prompt, type:

Command	Purpose
show ip route *prefix mask*	Displays the routing table for the specified subnet (prefix/mask).

Displaying the IP Routing Table for a Specified Subnet: Example

This example shows how to display the routing table for a specified subnet.

```
SCE#show ip route 10.1.60.0 255.255.255.0
|    prefix       |     mask        |    next hop     |
|-----------------|-----------------|-----------------|
|       10.1.60.0 |   255.255.255.0 |        10.1.1.5 |
sce#
```

IP Advertising

- Configuring IP Advertising, page 6-4
- How to Display the Current IP Advertising Configuration, page 6-5

IP advertising is the act of periodically sending ping requests to a configured address at configured intervals. This maintains the Cisco SCE platform IP/MAC addresses in the memory of adaptive network elements, such as switches, even during a long period of inactivity.

The following commands are relevant to IP advertising:

- **[no] ip advertising**
- **ip advertising destination**
- **ip advertising interval**
- **default ip advertising destination**
- **default ip advertising interval**
- **show ip advertising**
- **show ip advertising destination**
- **show ip advertising interval**

Configuring IP Advertising

To configure IP advertising, you must first enable IP advertising. You may then specify a destination address to which the ping request is to be sent and/or the frequency of the ping requests (interval). If no destination or interval is explicitly configured, the default values are assumed.

Options

The following options are available in the IP advertising commands:

- **interval—**The time interval between pings in seconds.

  default interval = 300 seconds
- **destination—**The IP address of the destination for the ping requests

  default destination = 127.0.0.1

How to Enable IP Advertising

From the SCE(config)# prompt, type:

Command	Purpose
ip advertising	Enables IP advertising.

How to Configure the IP Advertising Destination

From the SCE(config)# prompt, type:

Command	Purpose
ip advertising destination *destination*	Configures the destination for the IP advertising pings.

How to Configure the IP Advertising Interval

From the SCE(config)# prompt, type:

Command	Purpose
ip advertising interval *interval*	Configures the frequency of the IP advertising pings.

Configuring IP Advertising: Example

The following example shows how to configure IP advertising, specifying 10.1.1.1 as the destination and an interval of 240 seconds.

```
SCE(config)#ip advertising destination 10.1.1.1
SCE(config)#ip advertising interval 240
```

How to Display the Current IP Advertising Configuration

From the SCE# prompt, type:

Command	Purpose
show ip advertising	Displays the status of IP advertising (enabled or disabled), the configured destination, and the configured interval.

Configuring Time Clocks and Time Zone

- Displaying the System Time, page 6-6
- Displaying the Calendar Time, page 6-7
- Setting the System Clock, page 6-7
- Setting the Calendar, page 6-7
- Setting the Time Zone, page 6-8
- Removing the Current Time Zone Setting, page 6-9
- Configuring Daylight Saving Time, page 6-9

The Cisco SCE platform has three types of time settings, which can be configured: the clock, the calendar, and the time zone. It is important to synchronize the clock and calendar to the local time, and to set the time zone properly. The Cisco SCE platform does not track Daylight Saving Time automatically, so you must update the time zone when the time changes bi-annually.

The Cisco SCE platform has the following two time sources:

- A real-time clock, called the calendar, that continuously keeps track of the time, even when the Cisco SCE platform is not powered up. When the Cisco SCE platform reboots, the calendar time is used to set the system clock. The calendar is not used for time tracking during system operation.
- A system clock, which creates all the time stamps during normal operation. This clock clears if the system shuts down. During a system boot, the clock is initialized to show the time indicated by the calendar.

It does not matter which clock you set first, as long as you use the clock and calendar read commands to ensure they are synchronized.

The time zone settings are important because they allow the system to communicate properly with other systems in other time zones. The system is configured based on Coordinated Universal Time (UTC), which is standard in the industry for coordination with other manufacturers' hardware and software. For example, Pacific Standard Time would be written as PST-10, meaning that the name of the time zone is PST, which is 10 hours behind Universal Time.

When setting and showing the time, the time is always typed or displayed according to the local time zone configured.

Displaying the System Time

From the SCE(config)# prompt, type:

Command	Purpose
show clock	Displays system time.

Displaying the System Time: Example

The following example shows the current system clock.

```
SCE#show clock
12:50:03  UTC  MON  November 13 2001
sce#
```

Displaying the Calendar Time

From the SCE(config)# prompt, type:

Command	Purpose
show calendar	Displays calendar time.

Displaying the Calendar Time: Example

The following example shows the current system calendar.

```
SCE#show calendar
12:50:03  UTC  MON  May 11 2007
sce#
```

Setting the System Clock

Options

The following option is available:

- **time-date** the time and date you want to set, in the following format:

  hh:mm:ss day month year

From the SCE# prompt, type:

Command	Purpose
clock set *time-date*	Sets the system clock to the specified time and date.

Setting the System Clock: Example

The following example shows how to set the clock to 20 minutes past 10 AM, May 13, 2007, updates the calendar and then displays the time.

```
SCE#clock set 10:20:00 13 may 2007
SCE#clock update-calendar
SCE#show clock
10:21:10  UTC  THU  May  13  2007
```

Setting the Calendar

The calendar is a system clock that continues functioning even when the system shuts down.

Options

The following option is available:

- **time-date** —the time and date you want to set, in the following format:

  hh:mm:ss day month year

Step 1 From the SCE# prompt, type **calendar set** *time-date* and press **Enter**.

Sets the system calendar to the specified time and date.

The time specified in this command is relative to the configured time zone.

Step 2 From the SCE# prompt, type **clock read-calendar** and press **Enter**.

Synchronizes the system clock with the calendar time you just set .

Setting the Calendar: Example

The following example shows that the calendar is set to 10:20 AM, May 13, 2007. The clock is then synchronized with the calendar setting.

```
SCE#calendar set 10:20:00 13 may 20017
SCE#clock read-calendar
SCE#show calendar
10:21:06  UTC  THU  May  13  2007
```

Setting the Time Zone

Options

The following options are available:

- **zone—**The name of the time zone to be displayed.

  default = GMT

- **hours—**The hours offset from UTC. This must be an integer in the range –23 to 23.

  default = 0

- **minutes—**The minutes offset from UTC. This must be an integer in the range of 0 to 59. Use this parameter to specify an additional offset in minutes when the offset is not measured in whole hours.

  default = 0

From the SCE(config)# prompt, type:

Command	Purpose
clock timezone *zone hours minutes*	Sets the timezone to the specified timezone name with the configured offset in hours and minutes.

Setting the Time Zone: Example

The following example shows how to set the time zone to Pacific Standard Time with an offset of 10 hours behind UTC.

```
SCE(config)#clock timezone PST -10
SCE(config)#
```

Removing the Current Time Zone Setting

From the SCE(config)# prompt, type:

Command	Purpose
no clock timezone	Removes the timezone configuration and resets the timezone to the default value (UTC).

Configuring Daylight Saving Time

The Cisco SCE platform can be configured to automatically switch to daylight saving time on a specified date, and also to switch back to standard time. In addition, the time zone code can be configured to vary with daylight saving time if required. (For instance, in the eastern United States, standard time is designated EST, and daylight saving time is designated EDT).

- Options, page 6-9
- Guidelines, page 6-10
- How to Define Recurring Daylight Saving Time Transitions, page 6-11
- How to Define Non-Recurring Daylight Saving Time Transitions, page 6-11
- How to Cancel the Daylight Saving Time Configuration, page 6-11
- How to Display the Current Daylight Saving Time Configuration, page 6-12

Options

The transition times into and out of daylight saving time may be configured in one of two ways, depending on how the dates for the beginning and end of daylight saving time are determined for the particular location:

- recurring—If daylight saving time always begins and ends on the same day every year, (as in the United States), the **clock summer-time recurring** command is used. The beginning and ending days for daylight saving time can be configured once, and the system will automatically perform the switch every year.
- not recurring—If the start and end of daylight saving time is different every year, (as in Israel), the **clock summer-time** command is used. In this case, the transitions must be configured every year for that particular year. (Note that "year" is not necessarily a calendar year. If the transition days are determined in the fall, the transitions for that fall and the next spring may be configured.)

The day on which the transition takes place may be defined in several ways:

- Specific date—For example, March 29, 2004. A specific date, including the year, is defined for a not recurring configuration.
- First/last occurrence of a day of the week in a specified month—For example, the last Sunday in March. This is used for a recurring configuration.
- Day of the week in a specific week in a specified month—For example, Sunday of the fourth week of March. (This would be different from the last Sunday of the month whenever there were five Sundays in the month). This is used for a recurring configuration.

The following options are available:

- **zone—**The time zone code for daylight saving time
- **week** (recurring only)— the week of the month on which daylight saving begins (week1) and ends (week2)
- **day** (recurring only)—The day of the week on which daylight savings begin (day1) and ends (day2)
- **date** (non-recurring only)—The date of the month on which daylight saving begins (date1) and ends (date2)
- **month—**The month in which daylight saving begins (month1) and ends (month2)
- **year** (non-recurring only)—The year in which daylight saving begins (year1) and ends (year2)
- **offset—**The difference in minutes between standard time and daylight saving time.

  Default = 60 minutes

Guidelines

General guidelines for configuring daylight saving time transitions:

- Specify the time zone code for daylight saving time.
- recurring—Specify a day of the month (week#|first|last/day of the week/month).
- not recurring—Specify a date (month/day of the month/year).
- Define two days:
  - Day1 = beginning of daylight saving time.
  - Day2 = end of daylight saving time.
- In the Southern hemisphere, month2 must be before month1, as daylight saving time begins in the fall and ends in the spring.
- Specify the exact time that the transition should occur (24 hour clock).
  - Time of transition into daylight saving time—According to local standard time.
  - Time of transition out of daylight saving time—According to local daylight savings time.
- For the **clock summer-time recurring** command, the default values are the United States transition rules:
  - Daylight saving time begins: 2:00 (AM) on the second Sunday of March.
  - Daylight saving time ends: 2:00 (AM) on the first Sunday of November.

How to Define Recurring Daylight Saving Time Transitions

From the SCE(config)# prompt, type:

Command	Purpose
clock summer-time *zone* **recurring** *[week1 day1 month1 time1 week2 day2 month2 time2* [*offset*]]	Configures daylight saving time to start and stop on the specified days every year.

Defining Recurring Daylight Saving Time Transitions: Example

The following example shows how to configure recurring daylight saving time for a time zone designated "DST" as follows:

- Daylight saving time begins—0:00 on the last Sunday of March.
- Daylight saving time ends—23:59 on the Saturday of fourth week of November.
- Offset = 1 hour (default.

```
SCE(config)# clock summer-time DST recurring last Sunday March 00:00  4 Saturday November
23:59
```

How to Define Non-Recurring Daylight Saving Time Transitions

From the SCE(config)# prompt, type:

Command	Purpose
clock summer-time *zone* [*date1 month1 year1 time1 date2 month2 year2 time2* [*offset*]]	Defines non-recurring daylight saving time transitions.

Defining Non-Recurring Daylight Saving Time Transitions: Example

The following example shows how to configure non-recurring daylight saving time for a time zone designated "DST" as follows:

- Daylight saving time begins—0:00 on April 16, 2004.
- Daylight saving time ends—23:59 October 23, 2004.
- Offset = 1 hour (default)

```
SCE(config)# clock summer-time DST April 16 2004 00:00 October 23 2004 23:59
```

How to Cancel the Daylight Saving Time Configuration

From the SCE(config)# prompt, type:

Command	Purpose
no clock summer-time	Removes all daylight saving configuration.

How to Display the Current Daylight Saving Time Configuration

From the SCE# prompt, type:

Command	Purpose
show timezone	Displays the current time zone and daylight saving time configuration.

Configuring SNTP

- How to Enable the SNTP Multicast Client, page 6-13
- How to Disable the SNTP Multicast Client, page 6-14
- How to Enable the SNTP Unicast Client, page 6-14
- Disabling the SNTP Unicast Client, page 6-14
- How to Define the SNTP Unicast Update Interval, page 6-15
- How to Display SNTP Information, page 6-15

The Simple Network Timing Protocol (SNTP) is a simple solution to the problem of synchronizing the clocks in the various elements of the network. SNTP provides access to a time source via the network. The system clock and calendar are then set in accordance with this external source.

There are two options for the SNTP client. These functions are independent, and the system employ either one or both.

- Multicast SNTP client—Listens to SNTP broadcasts and updates the system clock accordingly.
- Unicast SNTP client—Sends a periodic request to a configured SNTP server, and updates the system clock according to the server response.

Note It is recommended that an IP access control list be configured to prevent access from unauthorized SNTP or NTP multicast servers (see "Configuring Access Control Lists (ACLs)" section on page 5-32).

The following commands are relevant to SNTP configuration:

- **[no] sntp broadcast client**
- **[no] sntp server address**
- **no sntp server all**
- **sntp update-interval**
- **show sntp**

How to Enable the SNTP Multicast Client

From the SCE(config)# prompt, type:

Command	Purpose
sntp broadcast client	Enables the SNTP multicast client. It will accept time updates from any broadcast server.

How to Disable the SNTP Multicast Client

From the SCE(config)# prompt, type:

Command	Purpose
no sntp broadcast client	Disables the SNTP multicast client. It will not accept any broadcast time updates.

How to Enable the SNTP Unicast Client

Options

The following option is available:

- **ip-address—**The IP address of the SNTP unicast server.

From the SCE(config)# prompt, type:

Command	Purpose
sntp server ip-address	Defines the SNTP unicast server so that SNTP client is able to query that server.

Enabling SNTP Unicast Client: Example

The following example shows how to enable an SNTP server at IP address 128.182.58.100.

```
SCE(config)# sntp server 128.182.58.100
```

Disabling the SNTP Unicast Client

How to Disable the SNTP Unicast Client and Remove All Servers

From the SCE(config)# prompt, type:

Command	Purpose
no sntp server all	Removes all SNTP unicast servers, preventing unicast SNTP query.

How to Remove One SNTP Server

Options

The following option is available:

- **ip-address—**The IP address of the SNTP unicast server.

From the SCE(config)# prompt, type:

s

Command	Purpose
no sntp server *ip-address*	Removes the specified SNTP unicast server.

How to Define the SNTP Unicast Update Interval

Options

The following option is available:

- **interval—**The time in seconds between updates (64 through 1024)

  default interval = 64 seconds

From the SCE(config)# prompt, type;

Command	Purpose
sntp update-interval *interval*	Configures the SNTP unicast client to query the server at the defined intervals.

Example

The following example shows how to set the SNTP update interval for 100 seconds.

```
SCE(config)# sntp update-interval 100
```

How to Display SNTP Information

From the SCE> prompt, type:

Command	Purpose
show sntp	Displays the configuration of both the SNTP unicast client and the SNTP multicast client.

Example

This example illustrates how to use this command.

```
SCE# show sntp
SNTP broadcast client: disabled
last update time: not available
SNTP unicast client: enabled
SNTP unicast server: 128.182.58.100
last update time: Feb 10 2002, 14:06:41
update interval: 100 seconds
```

Domain Name Server (DNS) Settings

- Configuring DNS Lookup, page 6-17
- Configuring Name Servers, page 6-18
- How to Add a Host to the Host Table, page 6-19
- How to Display Current DNS Settings, page 6-19

When a name of a host is given as a parameter to a CLI command that expects a host name or an IP address, the system translates the name to an IP address according to the following:

1. If the name is in a dotted decimal notation (that is, in the format x.x.x.x), it is directly translated to an IP address it represents.
2. If the name does not contain the dot character (.), the system looks it up in the IP Host table. If the name is found on the table, it is mapped to the corresponding IP address. The IP host table can be configured using the command **ip host.**
3. If the name does not contain the dot (.) character, and the domain name function is enabled (See the **ip domain-lookup** command), and a default domain name is specified (See the **ip domain-name** command), the default domain name is appended to the given name to form a fully qualified host name. This, in turn, is used to perform a DNS query translating the name to an IP address.
4. Otherwise, if the domain name function is enabled, the name is considered to be fully qualified, and is used to perform a DNS query translating the name to an IP address.

The following commands are relevant to DNS settings:

- **ip name-server**
- **ip domain-name**
- **no ip domain-name**
- **ip domain-lookup**
- **show hosts**

Configuring DNS Lookup

How to Enable DNS Lookup

From the SCE(config)# prompt, type:

Command	Purpose
ip domain-lookup	Enables DNS lookup.

How to Disable DNS Lookup

From the SCE(config)# prompt, type:

Command	Purpose
no ip domain-lookup	Disables DNS lookup.

Configuring Name Servers

- Options, page 6-18
- How to Define Domain Name Servers, page 6-18
- How to Remove a Domain Name Server, page 6-18
- How to Remove All Domain Name Servers, page 6-19

Options

The following options are available:

- **server-ip-address—**The IP address of the domain name server. You can define more than one DNS server (server-ip-address1, server-ip-address2, server-ip-address3)

How to Define Domain Name Servers

Use this command to specify the address of one or more name servers to use for name and address resolution.

From the SCE(config)# prompt, type:

Command	Purpose
ip name-server *server-address1 [server-address2 [server-address3]]*	Defines the servers at the specified addresses as domain name servers.

Defining Domain Name Servers: Example

The following example shows how to configure the two name server (DNS) IP addresses.

```
SCE(config)#ip name-server 10.1.1.60 10.1.1.61
```

How to Remove a Domain Name Server

From the SCE(config)# prompt, type:

Command	Purpose
no ip name-server *server-address1 [server-address2 [server-address3]]*	Removes the specified server from the DNS list.

Removing a Domain Name Server: Example

The following example shows how to remove name server (DNS) IP addresses.

```
SCE(config)#no ip name-server 10.1.1.60 10.1.1.61
```

How to Remove All Domain Name Servers

From the SCE(config)# prompt, type:

Command	Purpose
no ip name-server	Removes all configured DNS servers.

How to Add a Host to the Host Table

Options

The following options are available:

- **hostname—**The name of the host.
- **ip-address—**The IP address of the host

From the SCE(config)# prompt, type:

Command	Purpose
ip host *hostname ip-address*	Adds the specified host to the host table.

Adding Hosts to Removing them from the Host Table: Example

The following example shows how to add a host to the host table.

```
SCE(config)#ip host PC85 10.1.1.61
```

The following example shows how to remove a hostname together with all its IP mappings.

```
SCE(config)#no ip host PC85
```

How to Display Current DNS Settings

From the SCE# prompt, type:

Command	Purpose
show hosts	Displays current DNS settings.

Displaying Current DNS Settings: Example

The following example shows how to display current DNS information.

```
SCE#show hosts
Default domain is Cisco.com
Name/address lookup uses domain service
Name servers are 10.1.1.60, 10.1.1.61
Host                 Address
----                 -------
PC85                 10.1.1.61
sce#
```

Configuring Cisco Discovery Protocol

Cisco Discovery Protocol (CDP) is a device discovery protocol that runs on Cisco manufactured equipment, and is now supported on the Cisco SCE 8000 platform.

- Cisco Discovery Protocol, page 6-20
- Cisco Discovery Protocol on the Cisco SCE 8000 Platform, page 6-21
- Configuring CDP on the Cisco SCE 8000 Platform, page 6-22
- Monitoring and Maintaining CDP, page 6-25
- CDP Configuration Examples, page 6-27

Cisco Discovery Protocol

CDP is primarily used to obtain protocol addresses of neighboring devices and discover the platform of those devices. It is media- and protocol-independent, and runs on all equipment manufactured by Cisco, including routers, bridges, access servers, and switches.

CDP runs on all media that support Subnetwork Access Protocol (SNAP), including LAN, Frame Relay, and ATM physical media. CDP runs over the data link layer only. Therefore, two systems that support different network-layer protocols can learn about each other.

Each device configured for CDP sends periodic messages, known as advertisements, to a multicast address. Each device advertises at least one address where it can receive SNMP messages. The advertisements also contain time-to-live, or holdtime, information, which indicates the length of time a receiving device should hold CDP information before discarding it. Each device also listens to the periodic CDP messages sent by others in order to learn about neighboring devices and determine when their interfaces to the media go up or down.

CDP Version-2 (CDPv2) is the most recent release of the protocol and provides more intelligent device tracking features. These features include a reporting mechanism that allows for more rapid error tracking, thereby reducing costly downtime. Reported error messages can be sent to the console or to a logging server, and include instances of native VLAN IDs (IEEE 802.1Q) on connecting ports that do not match, and port duplex states between connecting devices that do not match.

Type-Length-Value fields (TLVs) are blocks of information embedded in CDP advertisements. Table 6-1 summarizes the TLV definitions for CDP advertisements.

Table 6-1 Type-Length-Value Definitions for CDPv2

TLV	Definition
Device-ID TLV	Identifies the device name in the form of a character string.
Address TLV	Contains a list of network addresses of both receiving and sending devices.
Port-ID TLV	Identifies the port on which the CDP packet is sent.
Capabilities TLV	Describes the functional capability for the device in the form of a device type, for example, a switch.
Version TLV	Contains information about the software release version on which the device is running.
Platform TLV	Describes the hardware platform name of the device, for example, Cisco 4500.

Table 6-1 Type-Length-Value Definitions for CDPv2 (continued)

TLV	Definition
IP Network Prefix TLV	Contains a list of network prefixes to which the sending device can forward IP packets. This information is in the form of the interface protocol and port number, for example, Eth 1/0.
VTP Management Domain TLV	Advertises the system's configured VTP management domain name-string. Used by network operators to verify VTP domain configuration in adjacent network nodes.
Native VLAN TLV	Indicates, per interface, the assumed VLAN for untagged packets on the interface. CDP learns the native VLAN for an interface. This feature is implemented only for interfaces that support the IEEE 802.1Q protocol.
Full/Half Duplex TLV	Indicates status (duplex configuration) of CDP broadcast interface. Used by network operators to diagnose connectivity problems between adjacent network elements.

Cisco Discovery Protocol on the Cisco SCE 8000 Platform

Because the Cisco SCE 8000 platform functions differently from a router or a switch, there are several unique features of CDP as supported on this device.

CDP Operational Modes on the Cisco SCE 8000

With a typical Cisco device, CDP is either enabled or disabled. When enabled, CDP packets are received and transmitted. When disabled, CDP packets are discarded and no packets are transmitted.

The Cisco SCE 8000 is not a typical Cisco device. It is usually installed as a bump-in-the-wire device, and transparently forwards packets from one interface to the corresponding interface. This behavior conflicts with typical Cisco CDP packet processing; a typical Cisco device never forwards CDP packets from one interface to another interface. To accommodate this behavior, the Cisco SCE 8000 extends the enabled state with three different CDP modes:

- **Standard mode**: Standard CDP operation. CDP packets are received and processed, as well as generated.

  In this mode CDP functions as it does on a typical Cisco device. This mode should be used in most cases, even though it is not the default mode.

- **Bypass mode** (default): CDP packets are received and transmitted unchanged. Received packets are not processed. No packets are generated.

  In this mode, "bump-in-the-wire" behavior is applied to CDP packets. This is the backward-compatible mode, equivalent to not having CDP support.

- **Monitor mode**: CDP packets are received, processed, and transmitted unchanged. CDP packets are analyzed and CDP neighbor information is available. No packets are generated.

  In this mode "bump-in-the-wire" behavior is applied to CDP packets. This mode may be confusing to operators and network management tools, because it is contrary to the concept of CDP as a physical link protocol.

Table 6-2 summarizes the CDP state and modes behavior in the Cisco SCE 8000.

Note When CDP is either not running or disabled at the interface level, CDP packets are discarded and CDP packets are not generated, regardless of the CDP mode.

Table 6-2 CDP Modes in the Cisco SCE 8000

CDP Mode	"cdp run" AND "cdp enable"	"no cdp run" OR "no cdp enable"
Standard	Received CDP packets processed CDP packets generated	Received CDP packets discarded CDP packets not generated
Bypass (Default)	Received CDP packets bypassed (not processed) CDP packets not generated	Received CDP packets discarded CDP packets not generated
Monitor	Received CDP packets processed and bypassed CDP packets not generated	Received CDP packets discarded CDP packets not generated

CDP Limitations on the Cisco SCE 8000

CDP as currently supported on the Cisco SCE 8000 has the following limitations:

- CDP is supported on traffic interfaces only (including cascade ports).
- CDP is currently managed by CLI only. There is currently no SNMP support for CDP on the Cisco SCE 8000.
- CDP always sends version 2 CDP packets. However it may receive v1 or v2 packets

Configuring CDP on the Cisco SCE 8000 Platform

To configure CDP, perform the tasks in the following sections:

- Enabling CDP Globally, page 6-22
- Setting CDP Mode, page 6-23
- Enabling CDP on a Specific Traffic Interface, page 6-23
- Setting the Hold Time, page 6-24
- Setting the Timer, page 6-24

Enabling CDP Globally

By default, CDP is enabled on the Cisco SCE 8000. If you prefer not to use the CDP device discovery capability, use the following command to disable it.

From the SCE(config)# prompt, type:

Command	Purpose
no cdp run	Disables CDP.

To reenable CDP after disabling it, use the following command.

From the SCE(config)# prompt, type:

Command	Purpose
cdp run	Enables CDP.

Note By default, when you enable CDP, it is set to bypass mode. To change the mode, see "Setting CDP Mode" section on page 6-23.

Setting CDP Mode

The Cisco SCE 8000 is usually installed as a bump-in-the-wire device, and therefore forwards packets (including CDP packets) from one interface to the corresponding interface, whereas a typical Cisco device never forwards CDP packets from one interface to another interface. Therefore, the Cisco SCE 8000 extends the enabled state with the following three CDP modes:

- standard—function as a typical CDP device
- monitor—monitor the CDP packets
- bypass—bypass the CDP packets

(See "CDP Operational Modes on the Cisco SCE 8000" section on page 6-21 for a description of the different CDP modes.)

Caution In cascade topologies, both Cisco SCE 8000 platforms must be configured to the same CDP mode.

By default, the CDP mode is set to bypass.

To reset the CDP mode to the default mode (bypass) use the **default cdp mode** command.

To change the CDP mode, use the following command in global configuration mode.

From the SCE(config)# prompt, type:

Command	Purpose
cdp mode (standard | monitor | bypass)	Changes the CDP mode.

Enabling CDP on a Specific Traffic Interface

By default, CDP is enabled on all traffic interfaces (see "CDP Limitations on the Cisco SCE 8000" section on page 6-22).

To disable CDP on a specific interface, use the **no cdp enable** command in the appropriate interface configuration mode.

To reenable CDP on a specific interface after disabling it, use the following command in the appropriate interface configuration mode. CDP must be enabled globally on the Cisco SCE 8000 platform (**cdp run** command) in order to enable a specific interface.

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

Command	Purpose
cdp enable	Enables CDP on a specific interface.



Tip For consistent CDP operation, it is recommended that both ports of any one traffic link be either enabled or disabled.

Setting the Hold Time

Use this command to set the mount of time the receiving device should hold a CDP packet from your router before discarding it. Use either the **no** or the **default** form of the command to restore the holdtime to the default value.

Options

The following option is available:

- **seconds**— Hold time value to be sent in the CDP update packets in seconds.

  default = 180 seconds

From the SCE(config)# prompt, type:

Command	Purpose
cdp holdtime *seconds*	Sets hold time.

Setting the Timer

Use this command to configure how often the Cisco SCE 8000 platform sends CDP updates. Use either the **no** or the **default** form of the command to restore the timer to the default value.

Options

The following option is available:

- **seconds**— How often the Cisco SCE 8000 platform sends CDP updates, in seconds.

  default = 60 seconds

From the SCE(config)# prompt, type:

Command	Purpose
cdp timer *seconds*	Sets the timer.

Monitoring and Maintaining CDP

To monitor and maintain CDP on the Cisco SCE 8000, use one or more of the following commands. The **clear** commands are in privileged EXEC mode. The **show** commands are in viewer mode.

Command	Purpose
clear cdp counters	Resets CDP traffic counters to zero
clear cdp table	Clears the table that contains CDP information about neighbors
show cdp	Displays the following information: • Interval between transmissions of CDP advertisements (transmission timer) • Number of seconds the CDP advertisement is valid for a given port (hold time) • Version of the advertisement • CDP mode
show cdp entry {*|*device-name[*]* **[protocol | version]}**	Displays protocol and version information about a specific neighboring device discovered using CDP. • Use "*" to display all devices. • Use *device-name** to display all devices beginning with device-name. • Use the **protocol** keyword to display only protocol information • Use the **version** keyword to display only version information.

Command	Purpose
show cdp neighbors [*type number*] [**detail**]	Displays the following information: • Type of device that was discovered • Name of the device • Number and type of the local interface (port) • Number of seconds the CDP advertisement is valid for the port • Device type • Device product number • Port ID If you use the **detail** keyword, the following additional information is displayed: • Entry address(es) • [Network protocol] address • Version • Advertisement version • Native VLAN ID • Duplex mode • VTP domain name associated with neighbor devices.
show cdp traffic	Displays the following information: • Total CDP packets output • Total CDP packets input • Number of CDP advertisements with bad headers • Number of times the checksum operation failed • Number of times CDP failed to send advertisements • Number of times the local device did not have enough memory to store the CDP advertisements • Number of invalid CDP advertisements • Number of times fragments of CDP advertisement were received • CDP version 1 advertisements output • CDP version 1 advertisements input • CDP version 2 advertisements output • CDP version 2 advertisements input

CDP Configuration Examples

Example: Setting the CDP Mode

The following example illustrates how to configure CDP mode to ‘standard’.

Caution In cascade topologies, both Cisco SCE 8000 platforms must be configured to the same CDP mode.

The **show** command verifies that the CDP configuration has been correctly updated.

```
sce(config)# cdp mode standard
sce(config)# do show cdp
Global CDP information:
        Sending CDP packets every 60 seconds
        Sending a holdtime value of 180 seconds
        Sending CDPv2 advertisements is enabled
        standard mode - CDP packets are received and processed. CDP packets are generated.
```

Example: Monitoring and Maintaining CDP

The following example shows a typical series of steps for viewing information about CDP neighbors.

Table 6-3 describes the significant fields shown in the output of the **show cdp neighbors** command.

```
sce> show cdp
Global CDP information:
        Sending CDP packets every 60 seconds
        Sending a holdtime value of 180 seconds
        Sending CDPv2 advertisements is enabled
        standard mode - CDP packets are received and processed. CDP packets are generated.
sce> show cdp neighbors
Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater, P - Phone

Device ID        Local Intrfce     Holdtme    Capability  Platform  Port ID
Lab-Router       Ten 3/0/0          169         R S I     CISCO7604 Ten 3/1
Lab-Router       Ten 3/1/0          169         R S I     CISCO7604 Ten 3/2
Lab-Router       Ten 3/2/0          169         R S I     CISCO7604 Ten 3/3
Lab-Router       Ten 3/3/0          169         R S I     CISCO7604 Ten 3/4
sce>
```

Table 6-3 show cdp neighbors Field Description

Field	Definition
Device ID	The name of the neighbor device and either the MAC address or the serial number of this device.
Local Intrfce	The protocol being used by the connectivity media.
Holdtme	The remaining amount of time (in seconds) the current device will hold the CDP advertisement from a sending router before discarding it.

Table 6-3 show cdp neighbors Field Description (continued)

Field	Definition
Capability (Capability Codes)	Capability (type of routing device) of the listed neighboring device. The capability types that can be discovered are: R—Router T—Transparent bridge B—Source-routing bridge S—Switch H—Host I— device is using IGMP r—Repeater **Note** The capability of the Cisco SCE 8000 is 'r' (Repeater), since it is installed as a bump-in-the-wire device.
Platform	The product number of the device.
Port ID	The protocol and port number of the device.

Enabling the CLI Interface Warning Banner

A warning banner is a message displayed when the user connects to the Cisco SCE using either Telnet or the console connection. It serves as a security warning for unauthorized users trying to connect to Cisco SCE platform. It can also provide device details, as well as information about the service and application.

By default the banner is disabled. You do not have to shutdown the Cisco SCE platform in order to enable or disable the banner.

From the SCE(config)# prompt, type:

Command	Purpose
banner login "*banner-text"*	Enables the display of the specified text as the warning banner when the CLI interface is accessed. Banner text should be enclosed in quotation marks or other delimiting characters.

OS Fingerprinting and NAT Detection

OS fingerprinting is the process of determining the identity of a remote host operating system by analyzing packets from that host. It detects the operating system used by the subscriber and whether the subscriber is present in a NAT environment by analyzing subscriber traffic. NAT detection is based on whether the same subscriber is connecting using multiple operating systems.

An encrypted fingerprint file that has the list of OS signatures is packaged with each SCOS release. Signature files are updated as needed, and the updated signature files are available on cisco.com.

The detected OS type is reported using the following mechanisms:

- RDRs—The subscriber OS type is reported in the Real-time Subscriber Usage RDR (SUR). These RDRs can be stored by the CM and interpreted using Insight.
- CLI—The subscriber OS type is available through OS fingerprinting and party info commands.
- VSA—Over mobile interfaces, the OS type is sent as a VSA in CCR-U over Gx.
- SCA BB Console—The OS type is available through an API that displays the OS type on the SCA BB console as part of the status of a subscriber.

Restrictions and Limitations

Due to the nature of the Cisco SCE platform, there are certain limitations to the scope of the OS fingerprinting and NAT detection feature:

- OS information is available only for logged-in and active subscribers.
- OS fingerprinting is not done continuously for any subscriber. If a subscriber changes OS or moves to a NAT environment during the time when they are not sampled, OS type or NAT environment cannot be detected.
- OS fingerprinting depends mainly on the parameters in the TCP-SYN packets. The signature database is built based on the default settings used by various operating systems. If the subscriber changes default parameters, such as TCP window size, through registries, it may lead to misclassification of the OS.
- The OS type will not be detected in any of the following situations:
  - If the subscriber connects to the internet using an http-proxy, or if there is a proxy or gateway that changes L3/L4 packets of the subscriber.
  - If the subscriber has only one flow.
  - If the subscriber has only UDP flows
- In case of multiple IP or IP range subscribers, OS fingerprinting is done only for a limited number IP addresses (default is 5).
- NAT detection is based on whether the same subscriber is connecting using multiple operating systems. Therefore, if all the users behind a NAT use the same OS, it is not possible to detect the NAT.
- When a subscriber runs multiple operating systems using vmware, it may be detected as a NAT even though the subscriber is not in NAT environment.

Configuring OS Fingerprinting

By default, the OS fingerprinting feature is disabled. When OS fingerprinting is enabled, you can also configure the following OS fingerprinting parameters:

- Sampling window—How long flows from a subscriber are fingerprinted
- Sampling interval—Interval between OS fingerprinting sampling windows

  OS fingerprinting is done for "sampling window" seconds every "sampling interval" minutes.
- NAT detection window—Time period within which detecting multiple operating systems for the same subscriber or IP address triggers NAT identification
- OS flush time–Time interval after which OS information is flushed from the system
- Signature file—Name of OS fingerprint signature file
- Scan port—Port used for opening OS fingerprinting flows
- GX reporting—Enable sending subscriber OS information in Gx messages

SUMMARY STEPS

1. **enable**
2. **configure**
3. interface linecard 0
4. os-fingerprinting
5. (Optional) **os-fingerprinting sampling window** *window* **interval** *interval*
6. (Optional) **os-fingerprinting NAT-detection-window** *time*
7. (Optional) **os-fingerprinting os-flush-time** *time*
8. (Optional) **os-fingerprinting signature-file** *filename*
9. (Optional) **os-fingerprinting scan-port** *port#*
10. (Optional) **os-fingerprinting gx-report**

DETAILED STEPS

	Command	Purpose
Step 1	**enable** **Example:** `SCE> enable`	Enables privileged EXEC mode. Enter your password when prompted.
Step 2	**configure** **Example:** `SCE# configure`	Enters global configuration mode.
Step 3	**interface linecard** **Example:** `SCE(config)# interface linecard 0`	Enters interface linecard configuration mode.

	Command	Purpose
Step 4	**os**-fingerprinting **Example:** `SCE(config if)# os-fingerprinting`	Enables OS fingerprinting and loads the *default.fp* signature file.
Step 5	**os-fingerprinting sampling window** *window* **interval** *interval* **Example:** `SCE(config if)# os-fingerprinting sampling window 60 interval 5`	(Optional) Configures the following: • Length of the OS sampling window, in seconds (10-300) • Interval between sampling windows, in minutes (10-1440)
Step 6	**os-fingerprinting NAT-detection-window** *time* **Example:** `SCE(config if)# os-fingerprinting NAT-detection-window 600`	(Optional) Enables NAT detection and configures the time period, in seconds, within which detecting multiple operating systems for one subscriber will trigger NAT identification. (10-300)
Step 7	**os-fingerprinting os-flush-time** *time* **Example:** `SCE(config if)# os-fingerprinting os-flush-time 3`	(Optional) Enables flushing the OS fingerprinting information and configures the time interval, in days, after which OS fingerprinting information is flushed from the system. (1-5)
Step 8	**os-fingerprinting signature-file** *filename* **Example:** `SCE(config if)# os-fingerprinting signature-file new-signature-file`	(Optional) Specifies the signature file used for OS fingerprinting.
Step 9	**os-fingerprinting scan-port** *port#* **Example:** `SCE(config if)# os-fingerprinting scan-port 50`	(Optional) Configures the port used for opening OS fingerprinting flows. The port numbers can be in the range of 0 - 65535. However, the following port numbers are blocked, and cannot be used for OS fingerprinting: 20, 21, 194, 554, 651, 654, 1720, 1755, 2000, 2948, 2949, 4374, 5060, 5061. For more information on this command, see the *Cisco SCE 8000 CLI Command Reference, Release 3.7.x*.
Step 10	**os**-fingerprinting gx-report **Example:** `SCE(config if)# os-fingerprinting gx-report`	(Optional) Enables sending subscriber OS information in Gx messages.

Monitoring OS Fingerprinting

To monitor OS fingerprinting, use one or more of the following commands.

These commands are in viewer mode.

Command	Purpose
show os-fingerprinting config	Displays the current OS fingerprinting configuration.The following information is displayed: • State of OS fingerprinting (enabled or disabled) • Sampling period • Sampling interval • NAT detection window • OS flush time • OS fingerprinting port • Signature file
show os-finger-printing signature-file	Displays the unencrypted contents of the signature file.
show interface linecard *slot-number* **subscriber name** *name* **[os-info]**	Displays information about a specified subscriber, including detected OS. To display only the OS fingerprinting information, use the **os-info** option**.**
show os-finger-printing subscriber-name *name*	Displays the OS fingerprinting information for the specified subscriber. This command displays the same information as the **show interface linecard** *slot-number* **subscriber name** *name* command with the **os-info** option.

Using the Bursty Traffic Convergence

During the peak hours, file sharing applications, such as P2P Bittorrent and other similar types of traffic, could go beyond the configured PIR level in the aggregative global controller mode. This is due to the bursty traffic nature of those applications. So, in order to have better convergence **bursty traffic convergence** enhancement is provided. The following steps are required to enable this **bursty traffic convergence** support to a particular AGC.

Step 1 Choose AGC mode in Cisco SCABB.

Step 2 It is recommended to create both upstream and downstream dedicated unique service AGC like P2P service AGC and apply SCABB policy.

Step 3 Configure the new cli “bursty-traffic-convergence” for such unique service AGC like P2P directional AGC in both directions.

```
aggregative-global-controller network-side <agc-index> bursty-traffic-convergence
aggregative-global-controller subscriber-side <agc-index> bursty-traffic-convergence
```

Step 4 Verify whether the bursty convergence is enabled in the running configuration.

Step 5 Verify whether the traffic is mapped to a unique service AGC such as P2P service dedicated AGC.

Step 6 Verify whether the corresponding traffic, such as P2P traffic, is mapped to the above AGC and is controlled in the configured PIR using insight graph reports.

The following steps are required to enable this **bursty traffic convergence** support to a particular GC.

Step 1 Choose GC mode in Cisco SCABB.

Step 2 It is recommended to create both upstream and downstream dedicated unique services GC, such as P2P service GC, and apply the SCABB policy.

Step 3 Configure the new cli "bursty-traffic-convergence" for such unique service GC, like P2P directional GC, in both directions.

The enable/disable configuration cli command to be carried in each Ten Gig/One Gig interface port is given below:

```
global-controller <gc-id from scabb> bursty-traffic-convergence <flag >
```

flag -> 0 disable

flag -> 1 enable

The show command to be carried in each Ten Gig/ One Gig interface port is given below:

```
Show interface <TenGig/One Gig> global-controller < gc-id from scabb >
bursty-traffic-convergence
```

Example:

For downstream:

```
SCE8000(config)#>int TenGigabitEthernet 3/0/0
SCE8000(config)#>do sh running-config | I "DS GBWC P2P"
global-controller 1 name "DS GBWC P2P"
```

Attach the bursty cli to the above GC as shown below:

```
SCE8000(config)#>global-controller 1 bursty-traffic-convergence 1
```

Similarly for upstream:

```
SCE8000(config)#>int TenGigabitEthernet 3/1/0
SCE8000(config)#>do sh running-config | include "US GBWC P2P"
global-controller 1 name "US GBWC P2P"
```

Attach the bursty cli to above GC as shown below:

```
SCE8000(config)#>global-controller 1 bursty-traffic-convergence 1
```

Step 4 Verify whether the bursty convergence is enabled in the running configuration.

Step 5 Verify whether the traffic is mapped to a unique service GC such as P2P service dedicated GC.

Step 6 Verify whether the corresponding traffic, such as P2P traffic, is mapped to the above GC and is controlled in the configured PIR using insight graph reports.

For more information on **bursty traffic convergence** CLI, see *CLI Command Reference* chapter of the *Cisco SCE 8000 CLI Command Reference, Release 5.1.0.*

Restrictions and Known Limitations in Bursty Traffic Convergence

- Bursty traffic convergence is found to be effective for controlled environment traffic.
- For controlling bursty traffic such as P2P Bittorrent in AGC and GC mode, it is recommended to have a unique service under dedicated AGC and GC that needs bursty traffic convergence.

DNS Assisted Classification/Sampling

DNS Assisted Classification supports existing Traffic-classification techniques to achieve more granular classification of encrypted subscriber traffic using the DNS traffic.

Configuring DNS Assisted Classification

By default DNS Assisted Classification is disabled. Once we enable it, the following parameters can be configured:

- DNS sampling time - Sampling time frame for DNS packets
- Maximum DNS packets - Maximum DNS packets to be sampled within the sampling period
- DNS refresh time - Time in days to flush old DNS-assistance entries

SUMMARY STEPS

1. **enable**
2. **configure**
3. interface linecard 0
4. Disable DNS bypass flow-filter rules and enable DNS classification on first packet via Service configuration
5. dns-flow sampling
6. (Optional) **dns-flow sampling** sample-time <*input in seconds*>
7. (Optional) **dns-flow sampling** max-dns-packets <*input as integer*>
8. (Optional) **dns-flow refresh-time** <*input is days*>
9. **do copy running-config startup-config**

DETAILED STEPS

	Command	Purpose
Step 1	**enable** **Example:** `SCE8000> enable`	Enables privileged EXEC mode. Enter your password when prompted.
Step 2	**configure** **Example:** `SCE8000# configure`	Enters global configuration mode.
Step 3	**interface linecard** **Example:** `SCE8000(config)# interface linecard 0`	Enters interface linecard configuration mode.

	Command	Purpose
Step 4	Disable DNS Bypass flow-filter rules and enable DNS First packet classification	By default, Flow-filter rules are configured to bypass dns traffic, for the DNS Assisted Classification to take effect, we need to configure the Service Configuration policy to allow SCE to process the dns traffic For more information on configure policy to process DNS traffic, see Cisco Service Control Application for Broadband User Guide
Step 5	**dns-flow sampling** **Example:** `SCE8000(config if)# dns-flow sampling`	Enables sampling of DNS packets
Step 6	**dns-flow sampling** sample-time <*input in seconds*> **Example:** `SCE8000(config if)# dns-flow sampling sample-time 25`	(Optional) Configures the length of the DNS sampling window, in seconds (1-1000)
Step 7	**dns-flow sampling** max-dns-packets <*input as integer*> **Example:** `SCE8000(config if)# dns-flow sampling max-dns-packets 2000`	(Optional) Configures the Maximum number of DNS Packets to be processed per sample-time per traffic processor
Step 8	**dns-flow refresh-time** <*input is days*> **Example:** `SCE8000(config if)# dns-flow refresh-time 2`	(Optional) Configures the Time Interval, in days, after which DNS Entries that exceed the interval are flushed from each of the traffic processor of the system. (1-1000).
Step 9	**do copy running-config startup-config**	Copies the current configured settings to the startup configuration.

For more information on dns classification related CLI, see “CLI Command Reference” chapter in Cisco SCE 8000 CLI Command Reference, Release 5.1.x.



Configuring Line Interfaces

Revised: September 23, 2015,

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 7-4
- Configuring Traffic Rules and Counters, page 7-26
- DSCP Marking, page 7-36
- Counting Dropped Packets, page 7-37

Line Interfaces

- Information About Line Interfaces, page 7-2
- How to Configure the Ten Gigabit Ethernet Line Interfaces, page 7-2
- Changing the Traffic Direction on the Ten Gigabit Ethernet Line Interfaces, page 7-3

The Line Interfaces (Subscriber and Network) are used to connect the Cisco SCE platform to the network. See the description of network topologies in the "Cisco SCE8000 Topology and Topology-Related Parameters" chapter of the *Cisco SCE8000 10GBE Installation and Configuration Guide*.

Information About Line Interfaces

The Cisco SCE 8000 10GBE line interfaces are found on the 1-port Ten Gigabit Ethernet SPAs installed in subslots 0 through 3 of slot 3. Each 1-port Ten Gigabit Ethernet SPA provides one 10GBE port, which interfaces 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 Cisco SCE platform reacts to Ethernet flow control and does not activate it. Therefore, a situation could arise in which flow control stalls the Cisco SCE platform by overflowing the Cisco 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 Cisco SCE platform, which may in turn trigger a reload of the Cisco SCE platform in an attempt to recover. However, the technical support file may not be generated when the Cisco 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.

How to Configure the Ten Gigabit Ethernet Line Interfaces

The 10GBE line interfaces are configured in TenGigabitEthernet mode.

Step 1 At the SCE# prompt, enter **configure** and press **Enter**.

Enters Global Configuration mode.

Step 2 At the SCE(config)# prompt, enter **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, enter **exit** and press **Enter**.

Exits to global configuration mode, from which you can access a different Ten Gigabit Ethernet interface.

Changing the Traffic Direction on the Ten Gigabit Ethernet Line Interfaces

The hardware design of the Cisco SCE 8000 10G platform is such that the traffic coming in and out of SPAs 0 and 2 is limited to a total of 16Gbps in each direction, as is the traffic coming in and out of SPAs 1 and 3. Since by default, in 10GBE installations, SPAs 0 and 2 are subscriber-side ports, and SPAs 1 and 3 are network-side ports, this effectively limits total traffic in each direction (upstream and downstream) to 16Gbps. Therefore, a site that has a total of more than 16Gbps of traffic in one direction (upstream or downstream), will exceed this limit and suffer packet loss.

You can use this command to swap the roles of the ports on link 1, which will switch part of the high-volume traffic to the opposite pair of SPAs and prevent either pair of SPAs from exceeding the 16Gbps limit.

Limitations

- This command is supported only Link 1 (3/2/0 and 3/3/0). It is not supported on Link 0.

  Only one interface on the link is explicitly configured. The corresponding interface is automatically set to the opposite traffic side.
- The connection mode must be either inline or receive-only. This command is not supported for cascade modes.
- This command can only be executed when no application is loaded, or in shutdown mode.

How to Specify the Traffic Direction on Link 1

Step 1 At the SCE# prompt, enter **configure** and press **Enter**.

Enters Global Configuration mode.

Step 2 At the SCE(config)# prompt, enter **interface TenGigabitEthernet 3/2/0** or **interface TenGigabitEthernet 3/3/0** and press **Enter**.

Enters Interface Configuration mode for the selected 10 GBE interface.

Step 3 At the SCE(config if)# prompt, enter **traffic-side (subscriber | network**) and press **Enter**.

The traffic side is set to the specified direction for the interface you are configuring and also automatically to the opposite direction for the corresponding interface.

Tunneling Protocols

- Selecting the Tunneling Mode, page 7-7
- Asymmetric L2 Support, page 7-18
- Monitoring VPN Support, page 7-23
- Displaying the Tunneling Configuration, page 7-19

Tunneling technology is used across various telecommunications segments to solve a wide variety of networking problems. The Cisco SCE platform is designed to recognize and process various tunneling protocols in several ways. The Cisco 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/ 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 7-18

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 Cisco 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, in environments 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 mangement.

  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 7-8
- Configuring DS-Lite Tunnels, page 7-9
- Configuring L2TP Tunnels, page 7-10
- Configuring GRE Tunneling, page 7-11
- Configuring IPinIP Tunneling, page 7-13
- Configuring DSCP Marking, page 7-14
- Configuring the 6to4 Environment, page 7-15
- Configuring the VLAN Environment, page 7-16
- Configuring the MPLS Environment, page 7-17
- 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 also 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 7-15.

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 SCE(config if)#> prompt, enter **shutdown** and press **Enter**.

Step 2 Enable 6to4 tunneling.

From the SCE(config if)#> prompt, enter **ip-tunnel 6to4** and press **Enter**.

Step 3 Restart the linecard.

From the SCE(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 SCE(config if)#> prompt, enter **shutdown** and press **Enter**.

Step 2 Disable 6to4 tunneling.

From the SCE(config if)#>prompt, enter **no ip-tunnel 6to4** and press **Enter**.

Step 3 Restart the linecard.

From the SCE(config if)#> prompt, enter **no shutdown** and press **Enter**.

Configuring 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 SCE(config if)#> prompt, enter **shutdown** and press **Enter**.

Step 2 Enable DS-Lite tunneling.

From the SCE(config if)#> prompt, enter **ip-tunnel DS-Lite** and press **Enter**.

Step 3 Restart the linecard.

From the SCE(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 SCE(config if)#> prompt, enter **shutdown** and press **Enter**.

Step 2 Disable DS-Lite tunneling.

From the SCE(config if)#> prompt, enter **no ip-tunnel DS-Lite** and press **Enter**.

Step 3 Restart the linecard.

From the SCE(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 SCE(config if)#> prompt, enter **shutdown** and press **Enter**.

Step 2 Enable DS-Lite tunneling.

From the SCE(config if)#> prompt, enter **ip-tunnel DS-Lite** and press **Enter**.

Step 3 Enable DS-Lite extension header support.

From the SCE(config if)#> prompt, enter **ip-tunnel DS-Lite Extention-Header-Support** and press **Enter**.

Step 4 Restart the linecard.

From the SCE(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 SCE(config if)#> prompt, enter **shutdown** and press **Enter**.

Step 2 Disable DS-Lite extension header support.

From the SCE(config if)#> prompt, enter **no ip-tunnel DS-Lite Extention-Header-Support** and press **Enter**.

Step 3 Restart the linecard.

From the SCE(config if)#> prompt, enter **no shutdown** and press **Enter**.

Configuring L2TP Tunnels

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

Enabling the 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 SCE(config if)#> prompt, enter **shutdown** and press **Enter**.

Step 2 Enable L2TP tunneling.

From the SCE(config if)#> prompt, enter **ip-tunnel l2tp skip** and press **Enter**.

Step 3 Restart the linecard.

From the SCE(config if)#> prompt, enter **no shutdown** and press **Enter**.

Disabling the L2TP Tunneling

Disables all IP tunnels except IPinIP and GRE.

Step 1 Shut down the linecard. (This is a root level command.)

From the SCE(config if)#> prompt, enter **shutdown** and press **Enter**.

Step 2 Disable L2TP tunneling.

From the SCE(config if)#> prompt, enter **no ip-tunnel l2tp skip** and press **Enter**.

Step 3 Restart the linecard.

From the SCE(config if)#> prompt, enter **no shutdown** and press **Enter**.

Configuring GRE Tunneling

- Enabling GRE Tunneling, page 7-11
- Disabling GRE Tunneling, page 7-12

GRE tunneling is an IP-based tunneling protocol; therefore the system must be specifically configured to recognize the flows inside the tunnel. The Cisco 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 Cisco 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 Cisco 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 SCE(config if)#> prompt, enter **shutdown** and press **Enter**.

Step 2 Enable GRE tunneling.

From the SCE(config if)#> prompt, enter **ip-tunnel gre skip** and press **Enter**.

Step 3 Restart the linecard.

From the SCE(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 SCE(config if)#> prompt, enter **shutdown** and press **Enter**.

Step 2 Disable GRE tunneling.

From the SCE(config if)#> prompt, enter **no ip-tunnel gre skip** and press **Enter**.

Step 3 Restart the linecard.

From the SCE(config if)#> prompt, enter **no shutdown** and press **Enter**.

Configuring IPinIP Tunneling

- Enabling IPinIP Tunneling, page 7-13
- Disabling IPinIP Tunneling, page 7-13

IPinIP is an IP-based tunneling protocol; therefore the system must be specifically configured to recognize the flows inside the tunnel. The Cisco 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 Cisco 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 IPinIP 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 Cisco 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 SCE(config if)#> prompt, enter **shutdown** and press **Enter**.

Step 2 Enable IPinIP tunneling.

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

Step 3 Restart the linecard.

From the SCE(config if)#> prompt, enter **no shutdown** and press **Enter**.

Disabling IPinIP Tunneling

Step 1 Shut down the linecard. (This is a root level command.)

From the SCE(config if)#> prompt, enter **shutdown** and press **Enter**.

Step 2 Disable IPinIP tunneling.

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

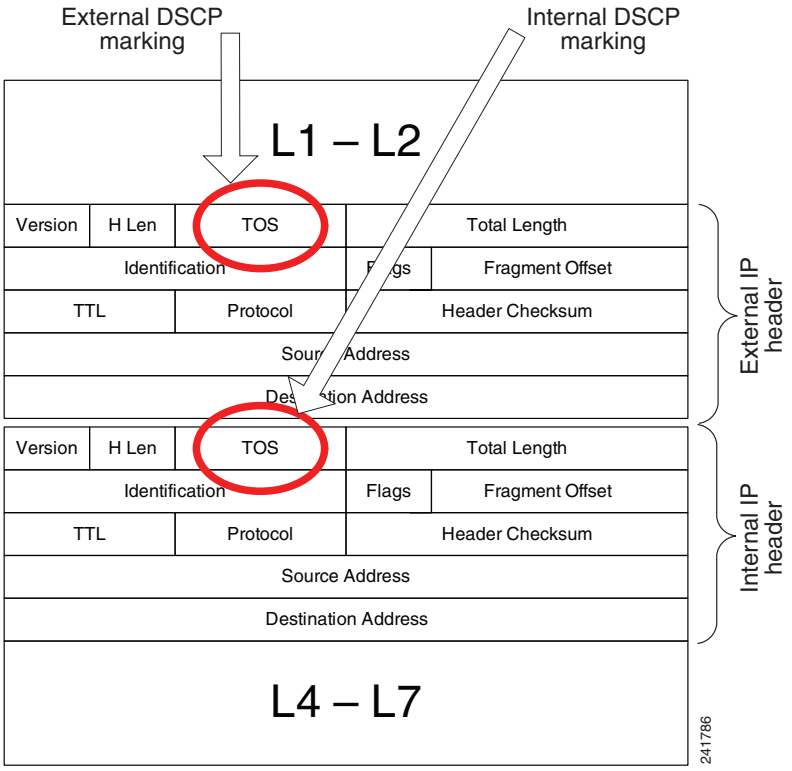
Step 3 Restart the linecard.

From the SCE(config if)#> prompt, enter **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 (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 Cisco 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 SCE(config if)#> prompt, type **shutdown** and press **Enter**.

Step 2 Configure the DSCP marking (.

From the SCE(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 SCE(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 SCE(config if)#> prompt, type **shutdown** and press **Enter**.

Step 2 Configure the DSCP marking.

From the SCE(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 SCE(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**.

Step 3 Configure the 6to4-IPv6 mode.

From the Global Configuration prompt, enter **debug const-db name seCommonConstDb.box.isIPv6to4IPv6Hash value true** and press **Enter**.

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 Cisco 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 Cisco SCE platform is configured by default to work in symmetric environments. A specific command should be used to allow correct operation of the Cisco 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 SCE(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.

```
SCE(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 SCE(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 7-26) 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 Cisco 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.

  Default port# = 1701

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

Command	Purpose
L2TP identify-by port-number *portnumber*	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 Cisco 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 Cisco 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 Cisco 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 Cisco 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 SCE(config if)# prompt, type:

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

Displaying the Tunneling Configuration

From the SCE# 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 SCE# 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 SCE# prompt, type:

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

How to Display the IPinIP Configuration

From the SCE# prompt, type:

Command	Purpose
show interface linecard 0 ip-tunnel IPinIP	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 SCE> prompt, type:

Command	Purpose
show interface linecard 0 VPN {name *vpn-name* **| all-names}**	Displays the logged-in VPNs.

How to Display the Asymmetric L2 Support Mode

From the SCE# prompt, type:

Command	Purpose
show interface linecard 0 asymmetric-L2-support	Displays asymmetric L2 support mode.

Managed VPNs

- Private IP Addresses, page 7-22
- Capacity, page 7-22
- Limitations for VPN mode, page 7-22

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 Cisco 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 Cisco 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 7-23
- Displaying Mappings for a Specified VPN: Examples, page 7-23

Options

The following option is available:

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

From the SCE> prompt, type:

Command	Purpose
show interface linecard 0 VPN name *vpn-name*	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.

```
SCE> 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.

```
SCE> 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 SCE> 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

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

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

- Options, page 7-24
- Displaying Subscribers Mapped to a IP 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.

From the SCE> prompt, type:

Command	Purpose
show interface linecard 0 subscriber mapping included-in IP *ip-range* **VPN** *vpn-name*	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

```
SCE> 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 7-24
- Displaying the Number of Subscribers Mapped to range on a Specified VPN: Example, page 7-25

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 SCE> prompt, type:

Command	Purpose
show interface linecard 0 subscriber amount mapping included-in IP *ip-range* **VPN** *vpn-name*	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

```
SCE> 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 7-26
- Configuring Traffic Counters, page 7-28
- Configuring Traffic Rules, page 7-29
- Managing Traffic Rules and Counters, page 7-34

Traffic Rules and Counters

- What are Traffic Rules and Counters?, page 7-26
- Traffic Rules, page 7-27
- Traffic Counters, page 7-27

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 Cisco 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 Cisco SCE platform, and thus is preserved when the application being run by the Cisco 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 Cisco SCE platform without being processed. This is useful when a particular type of traffic should be ignored by the Cisco 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 Cisco 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 Cisco SCE platform performance.

Traffic Rules

A traffic rule specifies that a defined action should be taken on packets processed by the Cisco 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 Cisco 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 that 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 7-28
- How to Delete a Traffic Counter, page 7-28
- How to Delete all Existing Traffic Counters, page 7-28

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 SCE(config if)# prompt, type:

Command	Purpose
traffic-counter name *name* **count-bytes|count-packets**	Adds a traffic counter with the specified name and counting mode.

How to Delete a Traffic Counter

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

Command	Purpose
no traffic-counter name *name*	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 SCE(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 7-29
- How to Create a Traffic Rule for IPv6 Addresses, page 7-32
- How to Delete a Traffic Rule, page 7-33
- How to Delete All Traffic Rules, page 7-33
- How to Delete All Flow Control Traffic Rules, page 7-34

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 7-16, 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 SCE(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 7-30
- Example 2, page 7-31
- Example 3, page 7-31
- Example 4, page 7-31

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
```

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.

```
SCE>enable 10
Password:<cisco>
SCE#configure
SCE(config)#interface linecard 0
SCE(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
SCE(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 the 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:

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:

Any of the following:

`upstream/downstream/both`

traffic-counter:

Either of the following:

- **name** <*name of an existing traffic counter*>— Packets that meet the traffice 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 SCE(config if)# prompt, type:

Command	Purpose
traffic-rule name *name* **ipv6 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.

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.

```
SCE> enable 10
Password:<cisco>
SCE # config
SCE(config)# interface linecard 0
SCE(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
SCE(config if)
```

How to Delete a Traffic Rule

From the SCE(config if)# prompt, enter:

Command	Purpose
no traffic-rule name *name*	Removes the specified traffic rule.

How to Delete All Traffic Rules

From the SCE(config if)# prompt, enter:

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

How to Delete All Flow Control Traffic Rules

From the SCE(config if)# prompt, enter**:**

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 7-34
- How to View All Traffic Rules, page 7-34
- How to View a Specified Traffic Counter, page 7-34
- How to View All Traffic Counters, page 7-35
- How to Reset a Specified Traffic Counter, page 7-35
- How to Reset All Traffic Counters, page 7-35

How to View a Specified Traffic Rule

From the SCE# prompt, enter:

Command	Purpose
show interface linecard 0 traffic-rule name *rule-name*	Displays the configuration of the specified traffic rule.

How to View All Traffic Rules

From the SCE# prompt, enter:

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 SCE# prompt, type:

Command	Purpose
show interface linecard 0 traffic-counter name *counter-name*	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

From the SCE# prompt, enter:

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.

```
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

From the SCE# prompt, enter:

Command	Purpose
clear interface linecard 0 traffic-counter name *counter-name*	Resets the specified traffic counter.

How to Reset All Traffic Counters

From the SCE# prompt, enter:

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 Cisco 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 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 SCE> prompt, enter:

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

Counting Dropped Packets

- About Counting Dropped Packets, page 7-37
- Disabling the Hardware Packet Drop, page 7-37

About Counting Dropped Packets

By default, the Cisco 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 SCE(config if)# prompt, type:

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



Configuring the Connection

Revised: September 23, 2015

Introduction

- Configuring the Connection Mode, page 8-2
- Monitoring the Connection Mode and Related Parameters, page 8-4
- Configuring the Link Mode, page 8-6
- External Optical Bypass, page 8-7
- Hardware Bypass, page 8-9
- Configuring the Static Subscribers, page 8-13
- Link Failure Reflection, page 8-15
- Asymmetric Routing Topology, page 8-18
- Configuring a Forced Failure, page 8-20
- Configuring the Failure Recovery Mode, page 8-21
- Configuring the Cisco SCE Platform/SM Connection, page 8-22

Note For more information regarding the physical installation of the Cisco SCE 8000 platform and cabling the connections, see the *Cisco SCE8000 10GBE Installation and Configuration Guide* and in particular, the following sections:

- Information About the SCE Platform (illustrations of modules and information regarding LEDs and interfaces)
- The Cisco SCE8000 Optical Bypass (information regarding the optical bypass module)
- Physical Topologies (topology diagrams)
- Connecting the Line Ports to the Network (cabling charts and diagrams)

Configuring the Connection Mode

The connection mode command allows you to configure the topology of the system in one command. The connection mode is determined by the physical installation of the Cisco SCE platform.

Caution This command can only be used if the line card is in either **no-application** or **shutdown** mode. If an application is installed on the Cisco SCE platform, the command will fail with an error message and help instructions.

Options

The following topology-related parameters are included in the connection mode command. Note that some options are relevant for cascaded topologies only.

- **Connection mode—**Can be any one of the following, depending on the physical installation of the Cisco SCE platform:
  - Inline—Single Cisco SCE platform inline
  - Receive-only—Single Cisco SCE platform receive-only
  - Inline-cascade—Two cascaded Cisco SCE platforms inline
  - Receive-only-cascade—Two cascaded Cisco SCE platforms receive-only

  The default is **inline**. Inline-cascade and Receive-only-cascade options are available only if cascade topology is configured. For an overview of the cascade topology, see the Cisco Service Control Product Installation Guide.
- **sce-id**—In cascaded topologies, defines which link is connected to this Cisco SCE platform.

  The sce-id parameter, which identifies the Cisco SCE platform, replaces the physically-connected-link parameter, which identified the link. This change was required with the introduction of the Cisco SCE 8000 GBE platform, which supports multiple links. In the Cisco SCE 8000 10GBE, the number assigned to the sce-id parameter (0 or 1) will be defined as the of number of the physically-connected-link.

Note For backwards compatibility, the physically-connected-link parameter is currently still recognized.

  Possible values are '0' and '1'.

  Not applicable to single Cisco SCE platform topologies.
- **Priority**—This parameter defines which is the primary Cisco SCE platform. It is applicable only in a two Cisco SCE platform topology.

  Possible values are 'primary' and 'secondary'

  Not applicable to single Cisco SCE platform topologies
- **On-failure—**This parameter determines the behavior of the system when the Cisco SCE platform either has failed or is booting:
  - cut the traffic (cutoff)
  - bypass the traffic (bypass)
  - automatically direct traffic through the external bypass module (external-bypass)

Note If the *external-bypass* option is configured, two optical bypass devices must be properly connected, one on each link. If an optical bypass device is not detected, the command is executed but a warning is issued. The system then enters warning mode until either the command is changed, or the presence of an optical bypass device is detected.

Note If the *bypass* option is configured and the connection mode is 'inline-cascade', a single optical bypass device must be connected on link #0 (SPA bays 0 and 1). As explained above, the command is executed, but the system enters warning mode.

Default:

- inline mode: **external-bypass**
- inline-cascade mode: **bypass**

Not applicable to receive-only topologies.

Note Do not change the connection mode unless the physical installation has been changed.

Step 1 From the SCE(config if)# prompt, type **connection-mode (inline | receive-only | inline-cascade | receive-only-cascade) [sce-id (0|1)] [priority (primary | secondary)] on-failure (bypass|external-bypass|cutoff)** and press **Enter**.

Configuring the Connection Mode Examples

Example 1

This example defines defines a primary Cisco SCE 8000 in a cascaded inline topology. Link 0 is connected to this device, and the link mode on failure is bypass (default).

```
SCE(config if)# connection-mode inline-cascade sce-id 0 priority primary
```

Example 2

This example defines a single-Cisco SCE platform, dual link, receive-only topology. The link mode **on-failure, sce-id**, and **priority** options are not applicable.

```
SCE(config if)# connection-mode receive-only
```

Monitoring the Connection Mode and Related Parameters

The following table shows the commands used to monitor the connection mode.

Note that the **show** commands are entered in Viewer mode.

Command	Purpose
show interface linecard 0 connection-mode	Displays the connection mode configuration.
show interface linecard 0 sce-id	Displays the Cisco SCE-ID.
show interface linecard 0 cascade redundancy-status	Displays the current redundancy status of the Cisco SCE platform.
show interface linecard 0 cascade peer-sce-information	Displays information about the peer Cisco SCE platform.
show interface linecard 0 cascade connection-status	Displays information about the cascade connections.

Connection Mode Examples

Monitoring the Connection Mode: Examples

The following example shows the current configuration of the connection mode for a single platform.

```
SCE> enable 5
Password:<cisco>
SCE> show interface linecard 0 connection-mode

Slot 0 connection mode
Connection mode is inline
slot failure mode is external-bypass
Redundancy status is standalone
SCE>
```

The following example shows the current configuration of the connection mode for a cascaded system.

```
SCE> enable 5
Password:<cisco>
SCE> show interface linecard 0 connection-mode

Slot 0 connection mode
Connection mode is inline-cascade
slot 0 sce-id is 1
slot 0 is secondary
slot 0 is connected to peer
slot failure mode is bypass
Redundancy status is active
SCE>
```

Viewing the Cisco SCE-ID: Example

```
SCE> enable 5
Password:<cisco>
SCE> show interface linecard 0 sce-id

slot 0 sce-id is 1
```

Viewing the Current Redundancy Status of the Cisco SCE Platform: Example

The following example shows typical output of this command.

```
SCE> enable 5
Password:<cisco>
SCE> show interface linecard 0 cascade redundancy-status

Redundancy status is active
```

Viewing Information about the Peer Cisco SCE Platform: Example

The following example shows typical output of this command.

```
SCE> enable 5
Password:<cisco>
SCE> show interface linecard 0 cascade peer-sce-information


Peer SCE's IP address is 10.10.10.10
```

Monitoring the Connection Status: Examples

The following example shows the output of this command in the case of two cascaded Cisco SCE 8000 10GBE platforms where the cascade interfaces have not been connected correctly.

```
SCE> enable 5
Password:<cisco>
SCE> show interface linecard 0 cascade connection-status

SCE is improperly connected to peer SCE
Please verify that each cascade port is connected to the correct port of the peer SCE.
Note that in the current topology, the SCE must be connected to its peer as follows:
Port 3/2/0 must be connected to port 3/3/0 at peer
Port 3/3/0 must be connected to port 3/2/0 at peer
SCE>
```

The following example shows the output of this command in the case of two cascaded Cisco SCE platforms where the cascade interfaces have been connected correctly.

```
SCE> enable 5
Password:<cisco>
SCE> show interface linecard 0 cascade connection-status

SCE is connected to peer SCE
SCE>
```

Configuring the Link Mode

- About the Link Mode, page 8-6
- Options, page 8-6

About the Link Mode

The Cisco SCE platform has an internal hardware card used to maintain the links even when the Cisco SCE platform fails. This hardware card has three possible modes of operation:

- bypass
- forwarding
- cutoff

Normally, the link mode is selected by the Cisco SCE platform software according to the configured connection-mode. However, the **link mode** command can be used to enforce a specific desired mode. This may be useful when debugging the network, or in cases where we would like the Cisco SCE platform just to forward the traffic. (Note that this is only relevant to inline topologies even though the configuration is available also when in receive-only mode.)

Options

The following link mode options are available:

- **Forwarding—**Forwards traffic to the Cisco SCE platform for processing.
- **Bypass—**Stops all forwarding of traffic to the Cisco SCE platform. Traffic still flows through the Cisco SCE platform, but is not processed by it in any way.

  This does not affect the redundancy states.
- **Cutoff—**Completely cuts off flow of traffic through the Cisco SCE platform.

Recommendations and restrictions

Note the following recommendations and restrictions:

- For the Cisco SCE 8000 platform, the link mode setting is global, and cannot be set for each link separately. Therefore the **all-links** keyword must be used.
- Link mode is relevant only to inline topologies.
- The default link mode is forwarding.

  When other link modes are selected, active service control is not available and any service control configuration will not be applicable.
- It is recommended that in cascaded topologies, both Cisco SCE platforms be configured for the same link mode, otherwise the service will be unpredictable.

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

Command	Purpose
link mode all-links (forwarding|bypass|cutoff)	Enforces a specific desired mode.

External Optical Bypass

- How to Activate the External Bypass, page 8-7
- How to Deactivate the External Bypass, page 8-8
- How to Set the External Bypass to the Default State, page 8-8
- How to Display the State of the External Bypass, page 8-8

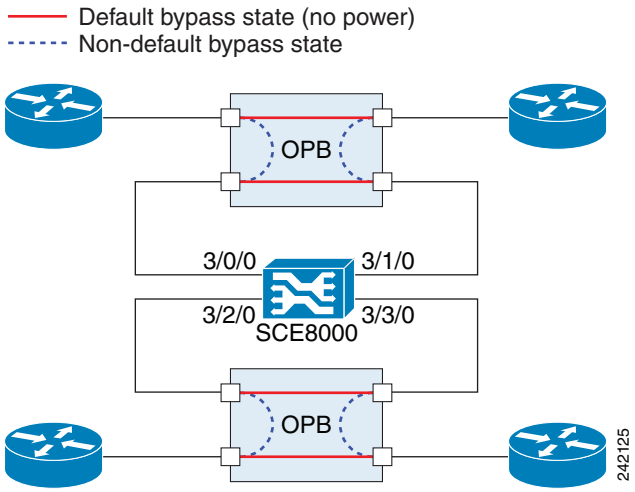
The Cisco SCE 8000 supports connection to up to two external optical bypass devices. These protect the line against power failure or total hardware failure, which prevents the hardware card from bypassing the traffic. Each external optical device protects a single traffic link passing through the Cisco SCE platform. The main objective of the external bypass is to provide automatic redundancy and failover support. However, the user can also manually enable the external bypass, assuming it is connected.

At power failure the external bypass is automatically activated. The external bypass can also be controlled by the software and by hardware in case of software failure.

In case of power failure, the bypass shortcuts the interfaces that are connected to the two sides of the Cisco SCE 8000, bypassing all the traffic, as illustrated in Figure 8-1.

The Cisco SCE 8000 can detect the presence of each external optical bypass device, and warns the user by various means (CLI **show** command, system operational-state, SNMP traps) if an expected external bypass device is not detected as present.

Figure 8-1 External Optical Bypass Connectivity



How to Activate the External Bypass

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

Command	Purpose
external-bypass	Activates the external bypass.

How to Deactivate the External Bypass

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

Command	Purpose
no external-bypass	Deactivates the external bypass.

How to Set the External Bypass to the Default State

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

Command	Purpose
default external-bypass	Sets the external bypass to the default state. The default state of the external optical bypass is deactivated.

How to Display the State of the External Bypass

From the SCE> prompt, type the following **show** command. Note that the **show** command is entered in the Viewer mode.

Command	Purpose
show interface linecard 0 external-bypass	Displays the state of the external bypass.

Output Sample: Both Optical Bypass Modules Functional

```
External bypass current state is 'not activated'.
External bypass failure state is 'activated'.
Amount of expected external bypass devices: 2 (automatically configured).
```

Output Sample: One Optical Bypass Module Not Detected

```
External bypass current state is 'not activated'.
External bypass failure state is 'activated'.
Amount of expected external bypass devices: 2 (automatically configured).
Warning: External bypass device expected but not detected on link #1
```

Hardware Bypass

- How to Enable the Hardware Bypass Mode, page 8-9
- How to Disable the Hardware Bypass Mode, page 8-10
- How to Display the Status of the Hardware Bypass Mode, page 8-10
- How to Set the Hardware Bypass Status for a Static Party, page 8-10
- How to Reset the Hardware Bypass Staus of a static party, page 8-10
- How to Display the Hardware Bypass Status of a Static Party, page 8-11
- How to Display the Startup Configuration Party Database, page 8-11
- How to Display the Currently Running Party Database Configuration, page 8-12
- How to Copy the Running Configuration Party Database to the Startup Configuration Party Database, page 8-12
- How to Copy the Startup Configuration Party Database and Create a Backup File, page 8-12

The Cisco SCE 8000 platform supports the Hardware Bypass feature for IPv4 traffic. The main objective of this feature is to bypass the traffic of the configured static parties created in the hardware bypass mode at the hardware (SIP module) level, based on their IP address or IP range. By default, the hardware bypass mode is disabled. However, users can enable the hardware bypass mode through the **hw-bypass mode** CLI command.

Note Before performing the party configurations, Engage Service Management Language (SML) 3.7.0 application should be installed in interface linecard configuration mode.

How to Enable the Hardware Bypass Mode

From the SCE(config)#> prompt, type:

Command	Purpose
hw-bypass mode	Enables the hardware bypass mode of the Cisco SCE 8000 platform. It also allows you to set the hardware bypass state for the specified static parties when these parties are configured in this mode.

How to Disable the Hardware Bypass Mode

From the SCE(config)#> prompt, type:

Command	Purpose
no hw-bypass mode	Disables the hardware bypass mode of the Cisco SCE 8000 platform. It also allows you to reset the hardware bypass state for the specified static parties when these parties are configured in this mode.

How to Display the Status of the Hardware Bypass Mode

From the SCE(config)#> prompt, type:

Command	Purpose
show hw-bypass mode	Displays the hardware bypass status of the Cisco SCE 8000 platform.

The following example shows how to display the status of the hardware bypass mode:

```
SCE>enable 15
Password:<cisco>
SCE#>show hw-bypass mode
hw-bypass mode is enabled
SCE(config)#>
```

How to Set the Hardware Bypass Status for a Static Party

From the SCE(config)#> prompt, type:

Command	Purpose
party name *party-name* **hw-bypass**	Sets the hardware bypass status for a specified static party in the hardware bypass mode of the Cisco SCE 8000 platform.

How to Reset the Hardware Bypass Staus of a static party

From the SCE(config)#> prompt, type:

Command	Purpose
no party name *party-name* **hw-bypass**	Resets the hardware bypass status of a specified static party in the hardware bypass mode of the Cisco SCE 8000platform.

How to Display the Hardware Bypass Status of a Static Party

From the SCE(config)#> prompt, type:

Command	Purpose
show party name *party-name* **hw-bypass**	Displays the hardware bypass status of a specified static party in the hardware bypass mode of the Cisco SCE platform.

Note The hardware bypass action can be performed only on the static parties created in the hardware bypass mode.

The following example shows how to display the hardware bypass of a specified static party:

```
SCE>enable 15
Password:<cisco>
SCE#>show party name [party-name] hw-bypass
SCE#> hw-bypass for party "[party-name]" is  disabled
```

How to Display the Startup Configuration Party Database

From the SCE(config)#> prompt, type:

Command	Purpose
show startup-config-party-db	Displays the contents of the startup configuration party database of the static parties that are configured in the Cisco SCE platform.

Note The configuration file contents will be displayed only if the corresponding startup configuration party database is copied from the running configuration party database.

The following example shows how to display the contents of the startup party database:

```
SCE>enable 15
Password:<cisco>
SCE#>show startup-config-party-db
#This is a party database configuration file (running-config-party-db) for static parties
only.
#Created on 13:34:02  UTC  TUE  July  12  2011
#cli-type 1
#version 1
hw-bypass mode
party name "N/A"
party name "[party-name]"
party mapping ip-address 24.11.52.128 name [party-name]
party mapping ip-address 110.10.10.10 name [party-name]
party name  [party-name] hw-bypass
SCE#>
```

How to Display the Currently Running Party Database Configuration

From the SCE(config)#> prompt, type:

Command	Purpose
show running-config-party-db	Displays the contents of the currently running party database configuration for the static parties that are configured in the Cisco SCE 8000 platform.

The following example shows how to display the contents of the running party database configuration:

```
SCE>enable 15
Password:<cisco>
SCE#>show running-config-party-db
#This is a party database configuration file (running-config-party-db) for static parties
only.
#Created on 13:34:02  UTC  TUE  July  12  2011
#cli-type 1
#version 1
hw-bypass mode
party name "N/A"
party name "[party-name]"
party mapping ip-address 24.11.52.128 name [party-name]
party mapping ip-address 110.10.10.10 name [party-name]
party name  [party-name] hw-bypass
SCE#>
```

How to Copy the Running Configuration Party Database to the Startup Configuration Party Database

From the SCE(config)#> prompt, type:

Command	Purpose
copy running-config-party-db startup-config-party-db	Enables the task of copying the currently running configuration party database to the startup configuration party database of the static parties that are configured on the Cisco SCE platform.

How to Copy the Startup Configuration Party Database and Create a Backup File

From the SCE(config)#> prompt, type:

Command	Purpose
copy startup-config-party-db *backup-file name*	Enables the task of copying the startup configuration party database and create a backup file of the configured static parties in the Cisco SCE 8000 platform.

Configuring the Static Subscribers

The following sections provide details on how to configure static subscribers for IPv4 and IPv6:

- How to the Set the IP Address for a Static Subscriber, page 8-13
- How to Set the IP Range for a Static Subscriber, page 8-13
- How to Display All Mappings to Dual Stack Static Subscriber, page 8-13
- How to Display IPv6 Mappings to Dual Stack Static Subscriber, page 8-14
- How to Display Dual Stack Static Subscriber, page 8-14

How to the Set the IP Address for a Static Subscriber

From the SCE(config)#> prompt, type:

Command	Purpose
party mapping ip-address *ip-address* **name** *party-name*	Sets the IPv4 address for the specified static party in the Cisco SCE 8000 platform.
party mapping ipv6-address *ip-address* **name** *party-name*	Sets the IPv6 address for the specified static party in the Cisco SCE 8000 platform.

How to Set the IP Range for a Static Subscriber

From the SCE(config)#> prompt, type:

Command	Purpose
party mapping ip-range *ip-address/mask-value* **name** *party-name*	Sets the IPv4 range for a specified static party in the Cisco SCE 8000 platform.
party mapping ip-range *ipv6-address:mask-value* **name** *party-name*	Sets the IPv6 range for a specified static party in the Cisco SCE 8000 platform.

Note If the mask value is not provided for the corresponding IP address, the complete mask value of 32 will be taken into consideration for the specified IP address.

How to Display All Mappings to Dual Stack Static Subscriber

From the SCE(config)#> prompt, type:

Command	Purpose
show part name *party-name* mappings all	Displays all mappings to dual stack static subscriber.

How to Display IPv6 Mappings to Dual Stack Static Subscriber

From the SCE(config)#> prompt, type:

Command	Purpose
show part name ***party-name*** **IPv6-ranges**	Displays IPv6 mappings to dual stack static subscriber.

How to Display Dual Stack Static Subscriber

From the SCE(config)#> prompt, type:

Command	Purpose
show part name ***party-name***	Displays dual stack static subscribers.
show interface LineCard 0 subscriber name ***party-name***	Displays dual stack static subscribers.

Link Failure Reflection

- How to Enable Link Failure Reflection, page 8-15
- How to Disable Link Failure Reflection, page 8-15
- Enabling and Disabling Link Failure Reflection on All Ports, page 8-15
- Configuring Link Failure Reflection in Linecard-Aware Mode, page 8-16

In some topologies, link failure on one port must be reflected to the related port to allow the higher layer redundancy protocol in the network to detect the failure and function correctly.

The **link failure-reflection** command determines the behavior of the system when there is a link problem. The link failure-reflection command enables reflection of a link failure. Use the [no] form of this command to disable failure reflection on the link.

The default value is **disabled**.

How to Enable Link Failure Reflection

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

Command	Purpose
link failure-reflection	Enables link failure-reflection.

How to Disable Link Failure Reflection

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

Command	Purpose
no link failure-reflection	Disables link failure-reflection.

Enabling and Disabling Link Failure Reflection on All Ports

- Options, page 8-16
- How to Enable Link Failure Reflection on All Ports, page 8-16
- How to Disable Link Failure Reflection on All Ports, page 8-16

The **link reflection on all ports** feature extends the link failure reflection feature. It allows the user to determine whether all ports should be taken down if a single port link fails.

In certain topologies, when a failure state occurs on one link, the link state must be reflected to all ports to signal any element using this Cisco SCE platform that the device is in a failure state, and therefore cannot be used.

In **link reflection on all ports** mode, all ports of the Cisco SCE platform are forced down and the link state of the first port is reflected on all the ports.

When recovering from the failure state, the forced down ports (the other link) are brought up only after the first failed port (link) has recovered. In addition, the reflection algorithm will not try to reflect failure for this link again for the next 15 seconds, to avoid link stability problems on auto-negotiation.

Options

The following options are available:

- The **on-all-ports** keyword enables reflection of a link failure to all ports.
- Use the **no** form of this command to disable failure reflection to all ports (the **on-all-ports** keyword is not used in the **no** form of the command).

The default value is **disabled**.

How to Enable Link Failure Reflection on All Ports

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

Command	Purpose
link failure-reflection on-all-ports	Enables failure reflection to all ports.

How to Disable Link Failure Reflection on All Ports

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

Command	Purpose
no link failure-reflection	Disables failure reflection to all ports.

Configuring Link Failure Reflection in Linecard-Aware Mode

- How to Enable Linecard-Aware Mode, page 8-17
- How to Disable Linecard-Aware Mode, page 8-17

The **linecard-aware-mode** option is an additional extension of the link failure reflection feature for use in MGSCP topologies. Use this option when the subscriber-side interface and the corresponding network-side interface of the same link are connected to the same linecard in the router.

This mode reflects a failure of one port to the other three ports of the Cisco SCE platform differently, depending on different failure conditions, as follows:

- One interface of the Cisco SCE 8000 is down: Link failure is reflected to the all other Cisco SCE platform ports.
- Two reciprocal ports of the Cisco SCE 8000 are down simultaneously, indicating a possible problem in the linecard of the router to which the Cisco SCE platform is connected: In this case the failure is not reflected to any of the other interfaces. This allows the second link in the Cisco SCE platform to continue functioning without interruption.

Use the **no** form of this command with the **linecard-aware-mode** keyword to disable the linecard aware mode without disabling link failure reflection itself.

How to Enable Linecard-Aware Mode

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

Command	Purpose
link failure-reflection on-all-ports linecard-aware-mode	Enables failure reflection to all ports with linecard aware mode.

How to Disable Linecard-Aware Mode

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

Command	Purpose
no link failure-reflection linecard-aware-mode	Disables linecard aware mode. Note that this command does not disable link failure reflection on all ports.

Asymmetric Routing Topology

- Asymmetric Routing and Other Service Control Capabilities, page 8-18
- Enabling Asymmetric Routing, page 8-18
- Monitoring Asymmetric Routing, page 8-19

In some Service Control deployments, asymmetrical routing occurs between potential service control insertion points. Asymmetrical routing can cause a situation in which the two directions of a bi-directional flow pass through different Cisco SCE platforms, resulting in each Cisco SCE platform seeing only one direction of the flow (either the inbound traffic or the outbound traffic).

This problem is typically solved by connecting the two Cisco SCE platforms through an MGSCP cluster, thereby making sure that both directions of a flow run through the same Cisco SCE platform. However, this is sometimes not feasible, due to the fact that the Cisco SCE platforms sharing the split flow are geographically remote (especially common upon peering insertion). In this type of scenario, the asymmetric routing solution enables the Cisco SCE platform to handle such traffic, allowing SCA BB to classify traffic based on a single direction and to apply basic reporting and global control features to uni-directional traffic.

Asymmetric Routing and Other Service Control Capabilities

Asymmetric routing can be combined with most other Service Control capabilities, however there are some exceptions.

Service Control capabilities that cannot be used in an asymmetric routing topology include the following:

- Subscriber redirect
- Subscriber notification
- Any kind of subscriber integration, including MPLS VPN. (Use subscriber-less mode or anonymous subscriber mode instead)
- Classical open flow mode, including the following:
  - Flow-open-mode classical explicitly enabled (ROOT level configuration)
  - VAS traffic forwarding mode enabled
  - Analysis layer transport mode enabled (ROOT level configuration)
  - ‘no TCP bypass-establishment’ mode enabled (ROOT level configuration)
  - A traffic rule is configured for certain flows to use the classical open flow mode (ROOT level configuration)

Enabling Asymmetric Routing

The asymmetric routing mode is disabled by default. It is typically enabled by the SCA-BB application when applying an appropriate service configuration.

Note that the detection of uni-directional flows is done by the Cisco SCE platform regardless of the asymmetric routing mode, but the appropriate configuration will assure that the uni-directional flows are properly classified and controlled.

For more information, see the *Cisco Service Control Application for Broadband User Guide*.

Monitoring Asymmetric Routing

Use the command below to display the following information regarding asymmetric routing:

- Current status of asymmetric routing mode (enabled or disabled)
- TCP unidirectional flows ratio: the ratio of TCP unidirectional flows to total TCP flows per traffic processor, calculated over the period of time since the Cisco SCE platform was last reloaded (or since the counters were last reset).

From the SCE> prompt, type:

Command	Purpose
show interface linecard 0 asymmetric-routing-topology	Displays the asymmetric routing information.

Monitoring Asymmetric Routing: Example

This example shows how to display the current asymmetric routing information.

```
SCE>show interface linecard 0 asymmetric routing-topology
Asymmetric Routing Topology mode is disabled
TCP Unidirectional flows ratio statistics:
=========================================
Traffic Processor 1   :   0%
Traffic Processor 2   :   0%
Traffic Processor 3   :   0%
Traffic Processor 4   :   0%
Traffic Processor 5   :   0%
Traffic Processor 6   :   0%
Traffic Processor 7   :   0%
Traffic Processor 8   :   0%
Traffic Processor 9   :   0%
Traffic Processor 10   :   0%
Traffic Processor 11   :   0%
Traffic Processor 12   :   0%

Note that the statistics are updated only if the system is configured to work in Enhanced
Open Flow (i.e. following settings are disabled:
Classical Open Flow mode, VAS, TCP no bypass est, etc.).
The statistics are updated once every two minutes
SCE>
```

Configuring a Forced Failure

Use the following commands to force a virtual failure condition, and to exit from the failure condition when performing an application upgrade.

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

Command	Purpose
force failure-condition	Forces a virtual failover condition. The system asks for confirmation `Forcing failure will cause a failover - do you want to continue? n` Type '**Y**' and press **Enter** to confirm the forced failure.
no force failure-condition	Exits from the virtual failure condition.

Configuring the Failure Recovery Mode

The **failure-recovery operation-mode** command defines the behavior of the system after boot resulting from failure.

Options

The following options are available:

- **operational—**After failure, the system will return to operational mode.
- **non-operational—**After failure, the system will remain not operational.

The default value is **operational**.

From the SCE(config)# prompt, type:

Command	Purpose
failure-recovery operation-mode operational|non-operational	Specifies the desired failure recovery mode.

Configure the Failure Recovery Mode: Examples

Example 1

This example sets the system to boot as non-operational after a failure.

```
SCE(config)#failure-recovery operation-mode non-operational
```

Example 2

This example sets the system to the default failure recovery mode.

```
SCE(config)# default failure-recovery operation-mode
```

Configuring the Cisco SCE Platform/SM Connection

- Configuring the Behavior of the Cisco SCE Platform in Case of Failure of the SM, page 8-22
- Configuring the SM-SCE Platform Connection Timeout, page 8-22

The user can configure the behavior of the Cisco SCE platform in case of failure of the Subscriber Manager (SM):

- If SM functionality is critical to the operation of the system—configure the desired behavior of the Cisco SCE platform if any loss of connection with the SM (may be due either to failure of the SM or failure of the connection itself).
- If SM functionality is not critical to the operation of the system—no action needs to be configured In this case you can specify that the system operational-status of the Cisco SCE platform should be 'warning' when the link is down.

Configuring the Behavior of the Cisco SCE Platform in Case of Failure of the SM

Options

The following options are available:

- **action**—The specified action will be performed in case of loss of connection between the Cisco SCE platform and the SM.

  Possible actions are:

  - **force-failure—**Force failure of Cisco SCE platform. The Cisco SCE platform then acts according to the behavior configured for the failure state.
  - **remove-mappings—**Remove all current subscriber mappings.
  - **shut—**The Cisco SCE platform shuts down and quits providing service.
  - **none** (default)—Take no action.
- **warning**—The system operational-status of the Cisco SCE platform should be 'warning' in case of loss of connection between the Cisco SCE platform and the SM. No action is taken.

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

Command	Purpose
subscriber sm-connection-failure action [*force-failure* | *none* | *remove-mappings* | *shut*]	Specifies the action that the Cisco SCE platform will perform if the Cisco SCE-SM connection fails.
subscriber sm-connection-failure warning	Specifies that the system operational-status of the Cisco SCE platform should be 'warning' if the Cisco SCE-SM connection fails.

Configuring the SM-SCE Platform Connection Timeout

You can also configure the timeout interval; the length of time that the SM-SCE platform connection is disrupted before a failed connection is recognized and the configured behavior is applied.

Options

The following option is available:

- **interval—**The timeout interval in seconds

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

Command	Purpose
subscriber sm-connection-failure timeout *interval*	Configures the connection timeout.



CHAPTER 9

Raw Data Formatting: The RDR Formatter and NetFlow Exporting

Revised: September 23, 2015

Introduction

Cisco Service Control is able to deliver gathered reporting data to an external application for collecting, aggregation, storage and processing over two protocols:

- RDRv1: the Service Control proprietary export protocol
- NetFlow V9: an industry standard export protocol

These two protocols can be used simultaneously in the same deployment. However, any specific destination (external collector) to which data is sent can be configured with only one protocol.

The services over IPv6 are classified to the same service IDs as the corresponding services over IPv4. RDRs that contain IP address information provide IPv4 or IPv6 addresses based on the traffic. Transaction Usage RDR, HTTP Transaction Usage RDR, Transaction RDR, Blocking RDR, and Link Usage RDR support IPv6 information.

Note Generic transaction usage RDR or anonymized transaction usage RDR is not generated for IPv6 if the Anonymized transaction usage RDR is enabled on the device.

- RDR Formatter and NetFlow Exporting Support, page 9-2
- Configuring Data Destinations and Categories, page 9-7
- Configuring the RDR Formatter, page 9-14
- Configuring NetFlow Exporting Support, page 9-15
- Configuring the Flexible Redirection Across SCOS Login, page 9-16
- Displaying Data Destination Configuration and Statistics, page 9-22
- Disabling the Linecard from Sending RDRs, page 9-25

RDR Formatter and NetFlow Exporting Support

- The RDR Formatter, page 9-2
- NetFlow, page 9-2
- Data Destinations, page 9-4

The RDR Formatter

The RDR formatter is used to gather the streams of Raw Data Records (RDRs) events passed from the application, format the data into external reporting protocol (RdrV1 or NetFlowV9), and send these reports to the appropriate destination(s). As the exporting of NetFlow traffic is done by the RDR Formatter, any of the configurations of the RDR Formatter affects the exporting of NetFlowV9 reports. For more information regarding RDR types and a description of their formats, see the *Cisco Service Control Application for Broadband Reference Guide*.

NetFlow

NetFlow reporting protocol is an industry standard for delivering gathered reporting data for external application for collecting, aggregation, storage and processing. The NetFlow protocol option integrates the Service Control solution with a wide range of existing data collectors and reporters.

Release 3.1.0 supports layer 7 application export reporters.

- NetFlow Terminology, page 9-2
- NetFlow Exporting Support, page 9-3

NetFlow Terminology

- Exporter

  A device (in this case, the RDR formatter component in the Cisco SCE platform) with NetFlow services enabled, responsible for exporting information using NetFlowV9 protocol.

- NetFlow Collector

  A device that receives records from one or more exporters. It processes the received export packet(s) by parsing and storing the record information. Records can be optionally aggregated before being stored on the hard disk.

- Export Packet

  A packet originating at the exporter, carrying the records of the exporter to the NetFlow collector.

- Packet Header

  The first part of an export packet. the packet header provides basic information about the packet such as the NetFlow version, number of records contained within the packet, sequence numbering, and the observation domain source ID.

- FlowSet

  A generic term for a collection of flow records that have a similar structure. In an export packet, one or more flowsets follow the packet header. There are two different types of flowsets:

  - Template FlowSet
  - Data FlowSet.

- Template FlowSet

  One or more template records that have been grouped together in an export packet

- Template Record

  Defines the structure and interpretation of fields in a flow data record.

- Data FlowSet

  One or more records, of the same type, that are grouped together in an export packet. Each record is either a flow data record or an options data record previously defined by a template record or an options template record.

- Flow Data Record

  A data record that contains values of the flow parameters corresponding to a template record.

NetFlow Exporting Support

The RDR formatter supports the exporting of NetFlowV9 reports and is able to send export packets to the configured destinations. The packets contain template records and data records. The template records define the format of the following data records. Each export packet may contain both types of records, or only one type of records.

NetFlow Templates

Each RDR type supported for NetFlowV9 exporting has a pre-defined mapping that allows the RDR formatter to convert it to a NetFlow V9 report and sent it over a NetFlow destination. The Cisco SCE platform maintains template records for several RDR types, with the structure of each NetFlow data record that corresponds to that RDR type. All NetFlow templates are pre-defined; users cannot create or edit the NetFlow templates.

Please note that if an RDR tag that is not supported for NetFlow exporting is configured to be sent over a NetFlow destination, this report will not be formatted and sent, and a special counter will be incremented along and a warning will be logged. (See the output of the show rdr-formatter statistics command – 'unsupported-tags'.)

Although template records and data records can be mixed in the same export packet, the template must precede any related data records. For this reason, templates are included in the first export packet, the first time a record is sent, and again at configured intervals (template refresh) to ensure that data records will be recognized and read correctly.

Data Destinations

- Categories, page 9-5
- Priority, page 9-5
- Setting DSCP for NetFlow, page 9-6
- Forwarding Modes, page 9-6
- Protocol, page 9-6
- Transport Type, page 9-6

The Cisco SCE platform can be configured with a maximum of eight destinations, three destinations per category. Each destination is defined by the following parameters:

- IP address
- port number
- protocol (RDRv1 or NetFlow)
- transport type (TCP (RDRv1) or UDP (NetFlow))

The destination is assigned a priority for each category to which it is assigned.

Figure 9-1 illustrates the simplest data destination topology, with only one category and one destination.

Figure 9-1 Data Destination Topology: One Category and One Destination

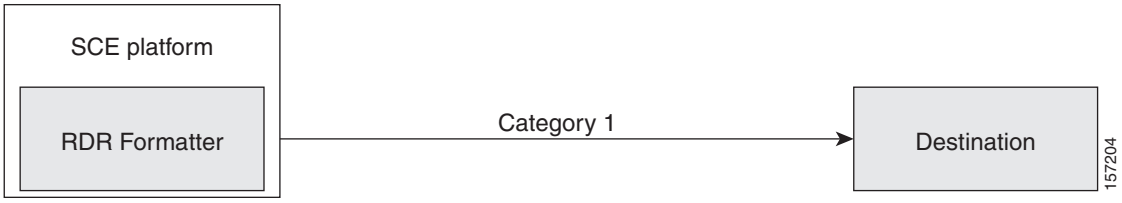
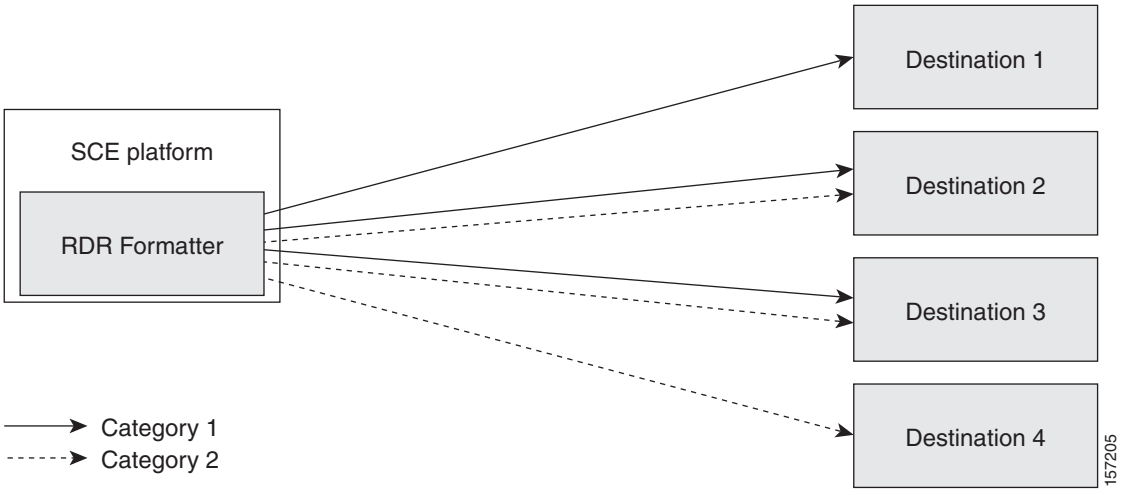


Figure 9-2 illustrates a complex topology using two categories and four destinations. Each category can send data to three of the four destinations.

Figure 9-2 Data Destination Topology: Two Categories and Four Destinations



Categories

In certain installations, data must be sent to different collector servers according to their type. For instance, in the pre-paid environment, some types of data must be sent to the pre-paid collector to get a new quota, while others should be sent to the mediation system. In this case, the data types are divided into up to four groups, and each group, or category, is assigned to a particular destination or destinations. The categories are defined by the application running on the Cisco SCE platform.

The system supports up to four categories:

- **Category 1—**Usage RDRs to the Collection Manager or mediation system
- **Category 2—**Quota RDRs to Pre-Paid Server or Subscriber Controller OSS
- **Category 3—**External events RDR or RT Signaling to various systems, such as a packet cable multi-media policy server
- **Category 4—**URL Query RDR to URL Filtering database (for example, surfControl)

Each destination must be configured regarding each category in use. Each destination may be assigned to more than one category and may be assigned the same or different priorities for each category.

It is also possible to remove a category from a destination, leaving only the desired category. If all categories are removed, the destination itself is deleted.

By default, the categories are referred to as Category 1 through Category 4. However, you can define meaningful names for the categories. This generally reduces confusion and prevents errors.

You can also configure the buffer size for each category. The total buffer size is 80 MB.

Priority

The priority value is used to indicate whether the destination should be a destination for a given category.

Priority is related to the redundant forwarding mode, in that it indicates which is the primary active connection. Priority values have no effect in multicast forwarding mode.

Each destination is assigned a priority value for each category. The first destination that is configured is automatically assigned a priority of 100 (highest priority) for all categories, unless explicitly defined otherwise.

Following are some important points to keep in mind regarding priority values:

- Two destinations may not have the same priority for one category. The priority values for destinations within a category must be unique to have any meaning.
- If only one priority value is assigned to the destination, that priority is automatically assigned to all categories for that destination.
- If only one category is assigned a priority value for a destination, no RDRs from the other categories will be sent to the specified destination.
- Assign a high priority if data from the specified category should be sent to this destination. Assign a low priority if data from the specified category should less likely to be sent to this destination.
- Redundant forwarding mode—Assign a high priority to the primary destination for the system/category. Assign a lower priority to the secondary destination for the system/category.

Setting DSCP for NetFlow

When using the NetFlowV9 protocol, priority can be defined by configuring a DSCP value to be assigned to the NetFlow packets. This DSCP value defines the DiffServ level of the NetFlow traffic to all destinations.

Forwarding Modes

When more than one destination is defined for a category, the system must decide which of these destinations is to receive the data. This is determined by the forwarding mode. There are three forwarding modes:

- Redundancy—All records are sent only to the primary (active) connection. If the primary connection fails, the records will be sent to the connected destination with the next highest priority.
- Multicast—All records are sent to all destinations. This feature may negatively affect performance in an installation with a high rate of data.
- Simple load balancing—Each successive record is sent to a different destination, one destination after the other, in a round robin manner. It is the responsibility of the collectors to aggregate the records.

  If one connection fails, the contents of the history buffer are sent to all connected destinations.



Note Some types of deployments using the NetFlow protocol require multicast forwarding mode. In a deployment where there are multiple destinations for at least one category, and at least one of those is a NetFlow destination, the multicast forwarding mode must be configured.

Protocol

The following two protocols are supported:

- The RDR Formatter: the Service Control proprietary protocol
- NetFlow V9: an industry standard protocol

These two protocols can be used simultaneously in the same deployment. However, each formatter destination can support only one protocol.

Transport Type

The following two transport types are available:

- TCP
- UDP

Currently, the transport type is linked to the configured protocol as follows:

- RDRv1 protocol requires TCP transport type
- NetFlow V9 protocol requires UDP transport type

Configuring Data Destinations and Categories

- Configuring a Data Destination, page 9-7
- Configuring the Data Categories, page 9-8
- Configuring the Forwarding Mode, page 9-13

Configuring a Data Destination

There are three general categories of CLI commands related to the configuration of data destinations:

- General commands that apply to both the RDRv1 protocol and the NetFlow protocol
- Commands that are relevant only to the RDR formatter (may affect NetFlow exporting as well)
- Commands relevant only to the NetFlowV9 protocol and the NetFlow exporting support

Options

In order for the data records, either RDRs or NetFlow export packets, from the Cisco SCE platform to arrive at the correct location, the following parameters must be configured:

- **ip-address—**The IP address of the destination
- **portnumber—**The port number
- **protocol**—The protocol used for data sent to the destination (either RDRv1 or NetFlow; if no protocol is assigned the protocol is RdrV1)
- **transport—**The transport type, TCP or UDP (optional, as this parameter is determined by the protocol)

A priority value may be assigned. Priority is important in the redundancy forwarding mode, but not crucial in multicast mode. Remember that in multicast mode, the existence of any priority value causes the destination to receive reports. The relationship between priorities and categories is addressed in "Configuring a Destination and Assigning Categories" section on page 9-9.

From the SCE(config)# prompt, type:

Command	Purpose
rdr-formatter destination *ip-address* **port** *portnumber* **[priority** *priority***] protocol** *protocol* **[transport** *transport***]**	Defines the destination. When no category is specified, the specified priority is assigned to all categories.

Configuring the Data Destinations: Examples

Example 1

This example shows how to configure a simple system with only one destination, using the NetFlow protocol. With only one destination, it is not necessary to configure a priority value.

```
SCE(config)# rdr-formatter destination 10.1.1.205 port 33000 protocol NetFlowV9 transport
udp
```

Example 2

The following example shows how to configure two destinations in a system without using the categories.

The first destination will automatically be assigned a priority of 100, and therefore the priority does not need to be explicitly defined. For the second destination, the priority must be explicitly defined.

The same priority will automatically be assigned to both categories for each destination, but since the categories will be ignored, this is irrelevant.

```
SCE(config)# rdr-formatter destination 10.1.1.205 port 33000 protocol RdrV1 transport tcp
SCE(config)# rdr-formatter destination 10.1.1.206 port 33000 priority 80 protocol RdrV1
transport tcp
```

Configuring the Data Categories

There are three steps in defining the data categories:

- Define the category names (optional).
- Configure the buffer size (optional).
- Configure the destinations with the proper priorities for each category, as well as configuring all the other destination parameters, may be approached in several different ways, and may take some planning. Refer to the examples below for illustrations of some of the issues involved in configuring categories.

Options

The following options are available:

- **category-number—**The number of the category (1-4)
- **category-name—**The name to be assigned to the category

From the SCE(config)# prompt, type:

Command	Purpose
rdr-formatter category number *category-number* **name** *category-name*	Defines the name for the specified category number. This category name can then be used in any **rdr-formatter** command instead of the category number.

Configuring the Buffer Size

Options

The following options are available:

- **category-number**—the number of the category (1-4)
- **size**—the size of the buffer in bytes

  By default, the buffer sizes are as follows:

  - Category 1—40 MB
  - Category 2—24 MB
  - Category 3—8 MB
  - Category 4—8 MB

  Total buffer size is 80 MB.

From the SCE(config)# prompt, type:

Command	Purpose
rdr-formatter category number *category-number* **buffer-size** *size*	Configures the buffer size for the specified category number.

Configuring a Destination and Assigning Categories

Options

The following options are available:

- **ip-address—**The IP address of the destination
- **portnumber—**The port number
- **category-number—**The number of the category (1-4)
- **category-name—**The name to be assigned to the category
- **priority—**The priority value assigned to this category for this destination (1-100)
- **protocol—**The protocol used for data sent to the destination (either RDRv1 or NetFlow; if no protocol is assigned the protocol is RdrV1)
- **transport—**The transport type, TCP or UDP (optional, as this parameter is determined by the protocol)

General Guidelines:

- A maximum of four categories can be configured in one command.
- The category may defined by either number or name.
- A different priority may be assigned to each category.
- Note that within each category the priorities must be unique for each destination.

From the SCE(config)# prompt, type:

Command	Purpose
rdr-formatter destination *ip-address* **port** *portnumber* **category [name** *category-name* **| number** *category-number*] **[priority** *priority*] **[category [name** *category-name* **| number** *category-number*] **[priority** *priority*]] **[category [name** *category-name* **| number** *category-number*] **[priority** *priority*]] **[category [name** *category-name* **| number** *category-number*] **[priority** *priority*]] **protocol** *protocol* **[transport** *transport*]	Defines the destination and assigns categories with optional priorities.

Configuring the Data Destinations with Categories: Examples

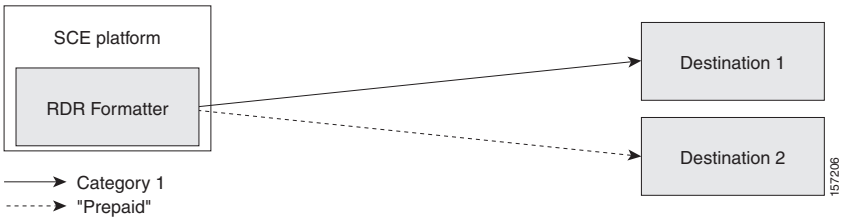
Example 1

The following example defines a name for one category, and then configures two destinations, assigning each to a different category (Figure 9-3).

The data from category 1 goes to the first destination, so a high priority was assigned to that category in the first destination, and no priority in the second.

Since all data from category 2 (prepaid) goes to the second destination, the priority assigned to category 2 is assigned only to the second destination and not to the first.

Figure 9-3 Configuring Destinations: Two Categories and Two Destinations



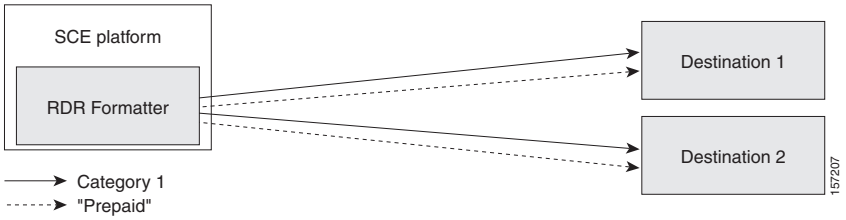
Note that if there is a loss of connection to either destination, transmission of data of the relevant category is interrupted until the connection is re-established. There is no redundant connection defined for either category.

```
SCE(config)# rdr-formatter category number 2 name prepaid
SCE(config)# rdr-formatter destination 10.1.1.205 port 33000 category number 1 priority 90
protocol RdrV1 transport tcp
SCE(config)# rdr-formatter destination 10.1.1.206 port 33000 category name prepaid
protocol RdrV1 transport tcp
```

Example 2

This example is similar to the above, but a low priority is assigned to the second category for each destination, rather than no priority (Figure 9-4). This allows each destination to function as a backup for the other in case of a problem with one of the connections (redundancy forwarding mode).

Figure 9-4 Configuring Destinations: Two Categories with Redundancy Mode




```
SCE(config)# rdr-formatter category number 2 name prepaid
SCE(config)# rdr-formatter destination 10.1.1.205 port 33000 category name prepaid
priority 90 category number 1 priority 25 protocol RdrV1 transport tcp
SCE(config)# rdr-formatter destination 10.1.1.206 port 33000 category number 1 priority 80
category name prepaid priority 20 protocol RdrV1 transport tcp
```

Example 3

This example demonstrates two methods for assigning one category to the first destination only, while the other category uses the second destination as the primary destination, and the first destination as a secondary destination.

```
SCE(config)# rdr-formatter category number 2 name prepaid
SCE(config)# rdr-formatter destination 10.1.1.205 port 33000 category name prepaid
priority 90 category number 1 priority 10 protocol RdrV1 transport tcp
SCE(config)# rdr-formatter destination 10.1.1.206 port 33000 category number 1 priority 95
protocol RdrV1 transport tcp
```

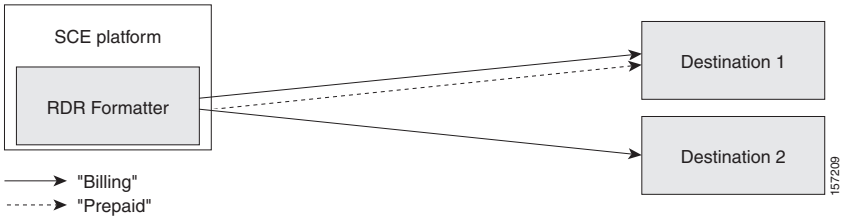
In the following example, all priority values seem quite high. However, it is the relative values of priorities for a category that determine which destination is the primary destination.

```
SCE(config)# rdr-formatter category number 2 name prepaid
SCE(config)# rdr-formatter destination 10.1.1.205 port 33000 priority 90 protocol RdrV1
transport tcp
SCE(config)# rdr-formatter destination 10.1.1.206 port 33000 priority 95 protocol RdrV1
transport tcp
SCE(config)# no rdr-formatter destination 10.1.1.206 port 33000 category name prepaid
protocol RdrV1 transport tcp
```

Example 4

This example illustrates a more complex configuration with one category (prepaid) assigned to one destination and the other (billing) being sent to both destinations, in multi-cast mode (Figure 9-5).

Figure 9-5 Configuring Destinations: Two Categories and Two Modes



The forwarding mode is defined for the entire RDR formatter, not just one category. Since the category "prepaid" goes to only one destination, the forwarding mode is irrelevant. It is relevant, however to the "billing" category, since it goes to two different destinations.

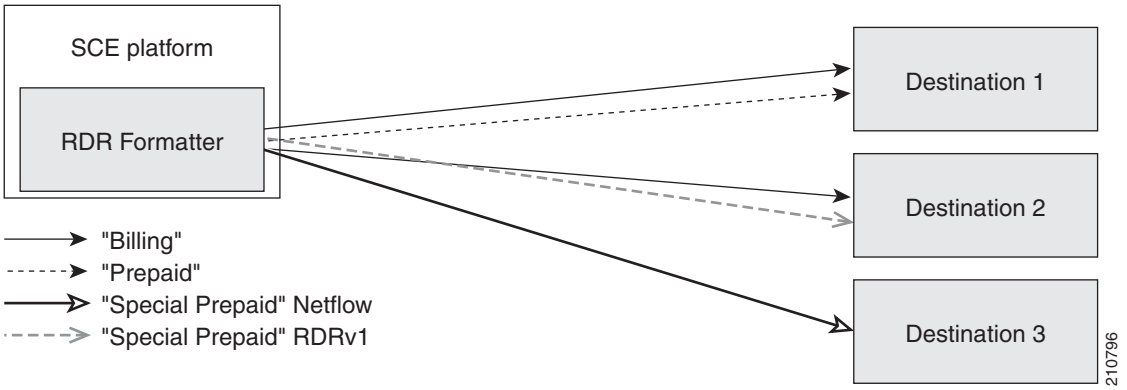
```
SCE(config)# rdr-formatter forwarding-mode multicast
SCE(config)# rdr-formatter category number 1 name billing
SCE(config)# rdr-formatter category number 2 name prepaid
SCE(config)# rdr-formatter destination 10.1.1.205 port 33000 priority 40 protocol
NetFlowV9 transport udp
SCE(config)# no rdr-formatter destination 10.1.1.205 port 33000 category name prepaid
protocol NetFlowV9 transport udp
SCE(config)# rdr-formatter destination 10.10.10.96 port 33000 category name billing
priority 90 protocol NetFlowV9 transport udp
SCE(config)# rdr-formatter destination 10.1.96.0 port 33000 category name prepaid priority
80prepaid priority 80 protocol NetFlowV9 transport udp
```

Example 5

Finally, the following example illustrates a configuration with three categories and three destinations, as follows (Figure 9-6):

- Category 1: 'Billing", RDRv1 protocol, goes to Destination 1
- Category 2: "Prepaid", RDRv1 protocol, goes to Destinations 1 and 2
- Category 3: "Special Prepaid", NetFlow V9 protocol, goes to Destination 3, RDRv1 protocol goes to Destination 2

Figure 9-6 Configuring Destinations: Three Categories and Two Protocols




```
SCE(config)# rdr-formatter forwarding-mode multicast
SCE(config)# rdr-formatter category number 1 name billing
SCE(config)# rdr-formatter category number 2 name prepaid
SCE(config)# rdr-formatter category number 3 name special-prepaid
SCE(config)# rdr-formatter destination 10.1.1.205 port 33000 category name billing
priority 90 category name prepaid priority 80 protocol RdrV1 transport tcp
SCE(config)# rdr-formatter destination 10.10.10.96 port 33000 category name prepaid
priority 90 category name special-prepaid priority 80 protocol RdrV1 transport tcp
SCE(config)# rdr-formatter destination 10.1.1.206 port 33000 category name special-prepaid
priority 90 protocol NetFlowV9 transport udp
```

Configuring the Forwarding Mode

In a deployment where there are multiple destinations for at least one category, and at least one of those is a NetFlow destination, the multicast forwarding mode must be configured.

Options

The following forwarding modes are available:

- **redundancy—**All records are sent only to the primary (active) connection. If the primary connection fails, the records will be sent to the connected destination with the next highest priority.
- **multicast—**All records are sent to all destinations. This feature may negatively affect performance in an installation with a high rate of data.
- **load-balancing—**Each successive record is sent to a different destination, one destination after the other, in a round robin manner. It is the responsibility of the collectors to aggregate the records.

From the SCE(config)# prompt, type:

Command	Purpose
rdr-formatter forwarding-mode *mode*	Configures the specified forwarding mode.

Configuring the Forwarding Mode: Example

The following example shows how to set the forwarding-mode to multicast.

```
SCE(config)# rdr-formatter forwarding-mode multicast
```

Configuring the RDR Formatter

- Options, page 9-14
- How to Configure the Size of the RDR Formatter History Buffer, page 9-14

Options

The following options are relevant specifically to the RDR formatter:

- Enabling and disabling the RDR formatter
- Setting the size of the RDR formatter history buffer.
- Dynamic mapping of RDRs to categories (see "Configuring the Flexible Redirection Across SCOS Login" section on page 9-16)

Use the following commands to enable or disable the RDR Formatter:

From the SCE(config)# prompt, type:

Command	Purpose
service rdr-formatter	Enables the RDR formatter.
no service rdr-formatter	Disables the RDR formatter.

How to Configure the Size of the RDR Formatter History Buffer

The following option is available:

- **size**—size of the history buffer in bytes. Maximum buffer size is 64 KB.

From the SCE(config)# prompt, type:

Command	Purpose
rdr-formatter history-size *size*	Sets the size of the RDR formatter history buffer.

Configuring NetFlow Exporting Support

- Options, page 9-15
- How to Configure a DSCP Value for NetFlow, page 9-15
- How to Configure the Template Refresh Interval, page 9-15

Options

The following options are relevant specifically to NetFlow exporting support (within the RDR-Formatter):

- Assigning a DSCP value to the NetFlow export packets to a specified destination for priority configuration.

  The DSCP value must be between 0 and 63, and be entered in HEX format.
- Configuring the frequency of exporting the template records (template refresh interval)

How to Configure a DSCP Value for NetFlow

Options

The following option is available:

- **dscp-value—**DSCP value to be assigned to the NetFlow packets over all destinations (0-63 in HEX format)

From the SCE(config)# prompt, type:

Command	Purpose
rdr-formatter protocol NetFlowV9 dscp *dscp-value*	Configures the DSCP value for NetFlow exporting support.

How to Configure the Template Refresh Interval

Options

The following options are available:

- **ip-address—**The destination IP address.
- **port-number—**The destination port number
- **timeout-value—**The frequency of exporting the template records in seconds (1 – 86400.)

From the SCE(config)# prompt, type:

Command	Purpose
rdr-formatter destination *ip-address* **port** *port-number* **protocol NetFlowV9 template data timeout** *timeout-value*	Sets the template refresh interval.

Configuring the Flexible Redirection Across SCOS Login

The REDIRECTION_HISTORY_RDR is generated whenever a redirection happens for the subscriber on breaching the quota.

The following pre-configuration are required to enable the “Flexible redirection across login” feature:

- Apply PQB.

  For more information, see “Subscriber Redirection” section in Cisco Subscriber Manger user guide.

  For more information, see “Redirection History RDR” section in Cisco Service Control Application for Broadband reference guide.

- Change the redirection history RDR from Category -1 to Category - 3 (as Subscriber Manager receives Category-3 RDR only).

```
SCE8000#>conf
SCE8000(config)#>rdr-formatter rdr-mapping tag-ID 0xf0f0f061 category-number 3
SCE8000(config)#>no rdr-formatter rdr-mapping tag-ID 0xf0f0f061 category-number 1
```

- Configure the RDR-formatter.

```
SCE8000#>configure
SCE8000(config)#>RDR-formatter forwarding-mode multicast
SCE8000(config)#>RDR-formatter destination [SM-IP Address] port 33145 category number
3 priority 100
```

Note The default port number for redirection is 33145. If you want to change the port number other than the default port number, you need to change the port number both on the Subscriber Manager side and RDR-formatter side.

Note One subscriber can have up to 128 profiles redirected.

In redirection section, the following two parameters will be added in p3sm.cfg file of Subscriber Manager:

- enable—This parameter is used to enable the Profile Redirection Feature. The values are yes/no.
- redirection_rdr_port—This parameter validates the port number. The default value is 33145.

Sample Configuration

```
#The following parameters define Profile Redirection feature options
[Redirection]
#To enable Profile Redirection Feature this should be used. By default no.
enable=no
#The port on which redirect rdr server to be run. By default 33145
redirection_rdr_port=33145
```

For more information, see *Cisco Service Control Subscriber Manager User Guide, Release 5.1.x.* and *Cisco Service Control Application for Broadband User Guide, Release 5.1.x*

Configuring Redirection Rate Limit

- In 5.1.0, Redirection Rate Limit feature configuration is done per redirection profile. This feature controls the redirections based on per second per profile.
- In this design, timer task will check every second to allow redirection or not as per configuration. If redirections exceed the configured Redirection Rate Limit value in a specific time, further redirections are not allowed. In other words, redirections will continue provided that the rate limit value not exceeding the configured redirection rate limit value.
- This feature will work with all subscriber redirection options – Only Once , Always, Periodically and Once in a Day/Week/Month.
- The Rate Limit value is configured for all 12 TP's. Based on this, each TP will get redirections as per the configured value (configured rate limit value divided by 12).

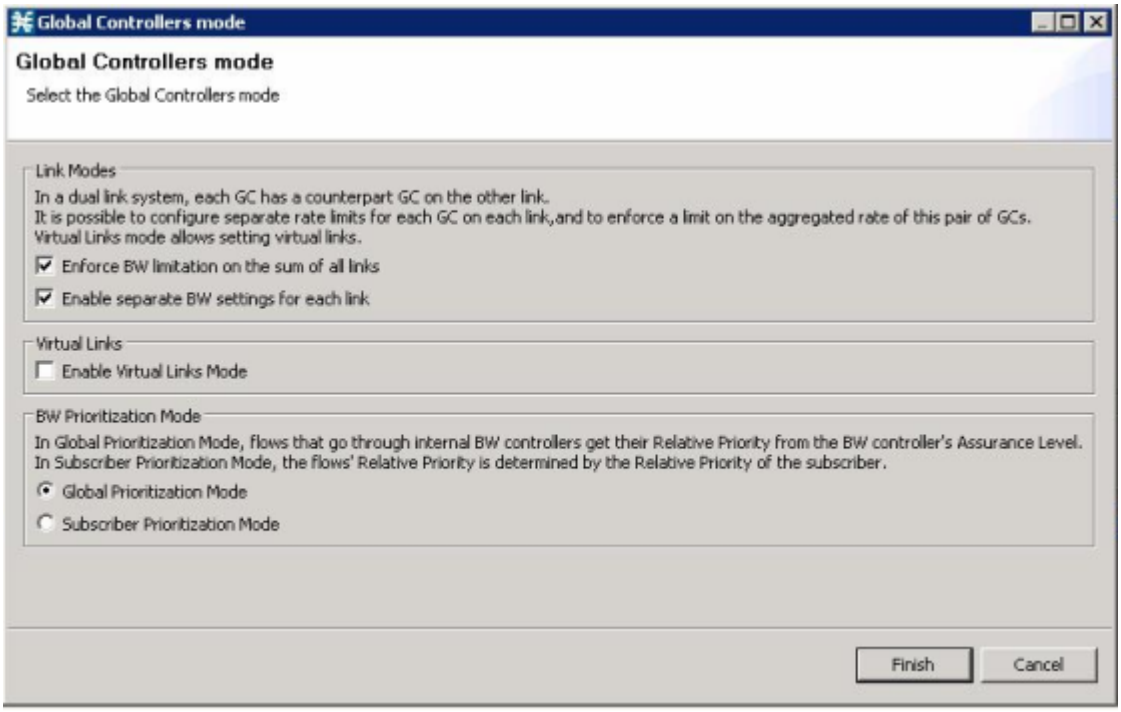
Configuring Total Link Limit

This feature enables SCE to control the total traffic rate in each link per system level. By using this feature, each SCE link can be rate limited with configured equal rate or different rate in AGC / GC mode to cater to the Customer requirements.

To configure to AGC/GC mode in SCE:

Step 1 Select the checkboxes **Enforce BW limitation on the sum of all links** and **Enable separate BW settings for each link** in SCABB console.

Figure 9-7 Configuring AGC/GC mode in SCE

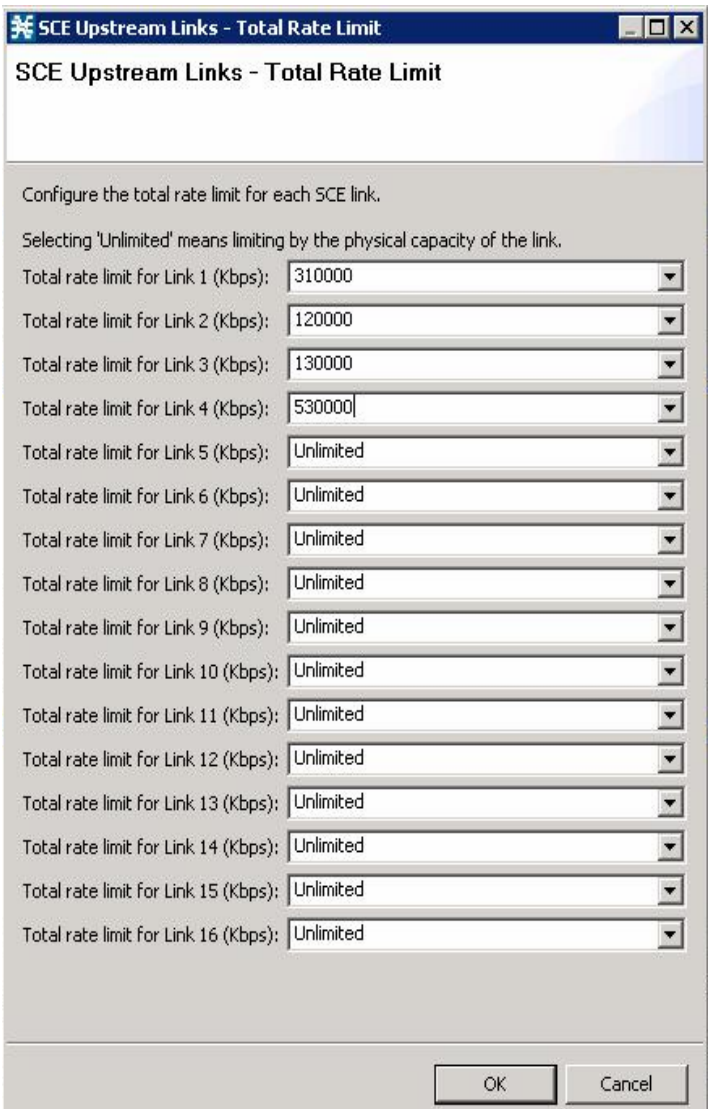


Step 2 Configure Edit Link Rate for both upstream and downstream.

To configure same rate for each link or different rate for each link:

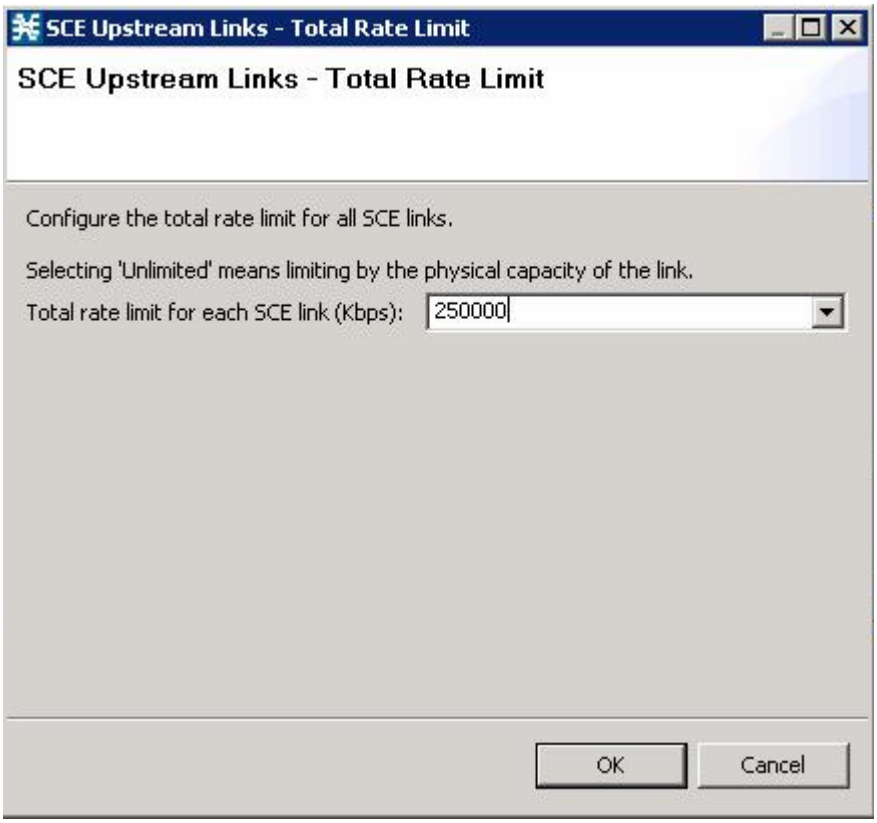
a. Disable the checkbox "Enable separate BW settings for each link" to configure same rate for each link.

Figure 9-8 Configuring same rate for each link



b. Enable the checkbox "Enable separate BW settings for each link" to configure different rate for each link.

Figure 9-9 Configuring different rate for each link



Step 3 Apply PQB to SCE.

Note Make sure to check the insight reports to confirm the total traffic rate is achieved per link.

Configuring Dynamic Mapping of RDRs to Categories

Dynamic configuration of RDRs to multiple categories is supported.

Each RDR tag has a list of categories. The default category is the one that was assigned when application was loaded.

The configuration of categories to RDR tags is done by adding and removing mappings. A user can add a mapping of an RDR tag to a category and can remove a mapping, including the default mapping. If all categories are removed from a tag, the tag is ignored as long as it remains with no mapped categories.

The user must provide the RDR tag ID and the category number to add or remove. The configuration is saved as part of the application configuration.

Configuring Mappings

Use these command to add or remove a mapping.

- Options, page 9-21

- How to Restore the Default Mapping for a Specified RDR Tag, page 9-21

Options

The following options are available:

- **tag-umber—**The complete 32 bit value given as an hexadecimal number. The RDR tag must be already configured in the Formatter by the application.
- **category-number—**Number of the category (1-4) to which to map the RDR tag.

From the SCE(config)# prompt, type:

Command	Purpose
rdr-formatter rdr-mapping (tag-id *tag-number* **category-number** *category-number***)**	Adds a mapping to a category. If the table already contains a mapping with the same tag and category number, an error is issued and nothing is done.
no rdr-formatter rdr-mapping (tag-id *tag-number* **category-number** *category-number***)**	Removes a mapping from a category.

How to Restore the Default Mapping for a Specified RDR Tag

From the SCE(config)# prompt, type:

Command	Purpose
default rdr-formatter rdr-mapping tag-id *tag-number*	Restores the default mapping for a specified RDR tag.

Displaying Data Destination Configuration and Statistics

- How to the Display the Current RDR Formatter Configuration, page 9-22
- How to the Display the Current RDR Formatter Statistics, page 9-23

The following commands can be used to display the RDR formatter configuration and statistics:

- **show rdr-formatter**
- **show rdr-formatter connection-status**
- **show rdr-formatter counters**
- **show rdr-formatter destination**
- **show rdr-formatter enabled**
- **show rdr-formatter forwarding-mode**
- **show rdr-formatter rdr-mapping**
- **show rdr-formatter statistics**
- **show rdr-formatter protocol NetFlowV9 dscp**

For a complete description of the other s**how rdr-formatter** commands, see the *Cisco SCE8000 CLI Command Reference* .

How to the Display the Current RDR Formatter Configuration

The system can display the complete data destination configuration, or just specific parameters.

From the SCE> prompt, type:

Command	Purpose
show rdr-formatter	Displays the current RDR formatter configuration.

Displaying the RDR Formatter Configuration: Example

The following example shows how to display the current RDR formatter configuration.

```
SCE#show rdr-formatter
Status: enabled
Connection is: up
Forwarding mode: redundancy
Connection table:
-------------------------------------------------------------------------|
Collector    | Port|Status|           Priority per Category:               |
IP Addres /  |     |      |-----------------------------------------------|
Host-Name    |     |      |Category1  |Category2  |Category3  |Category4  |
-------------------------------------------------------------------------|
10.1.1.205   |33000|Up    |100 primary|100 primary|100 primary|100 primary|
10.1.1.206   |33000|Down  |60         |60         |60         |60         |
10.12.12.12  |33000|Up    |40         |40         |40         |40         |
-------------------------------------------------------------------------|
RDR:    queued:    0 , sent:4460807,  thrown:   0,  format-mismatch:0
UM:     queued:    0 , sent:      0,  thrown:   0
Logger: queued:    0 , sent:     39,  thrown:   0
Errors: thrown:    0
Last time these counters were cleared: 20:23:05  IST  WED March 14 2007
```

How to the Display the Current RDR Formatter Statistics

From the SCE> prompt, type:

Command	Purpose
show rdr-formatter statistics	Displays the current RDR formatter statistics.

Displaying the Current RDR Formatter Statistics: Example

The following example shows how to display the current statistics in a deployment using both RDRv1 and NetFlow protocols.

```
SCE#show rdr-formatter statistics
RDR-formatter statistics:
=========================
Category 1:
sent:               1794517
in-queue:           0
thrown:             0
format-mismatch:    0
unsupported-tags:   1701243
rate:               2 RDRs per second
max-rate:           64 RDRs per second
Category 2:
sent:               12040436
in-queue:           0
thrown:             0
format-mismatch:    0
unsupported-tags:   0
rate:               12 RDRs per second
max-rate:           453 RDRs per second
Category 3:
sent:               0
in-queue:           0
```

```
thrown:             0
format-mismatch:    0
unsupported-tags:   0
rate:               0 RDRs per second
max-rate:           0 RDRs per second
Category 4:
sent:               0
in-queue:           0
thrown:             0
format-mismatch:    0
unsupported-tags:   0
rate:               0 RDRs per second
max-rate:           0 RDRs per second
Destination:    10.56.201.50 Port: 33000 Status: up
Sent:   13835366
Rate:   211  Max:   679
Last connection establishment:  17 hours, 5 minutes, 14 seconds
Destination:     10.56.204.7 Port: 33000 Status: up
Sent:   12134054
Rate:   183  Max:   595
Sent Templates:         13732
Sent Data Records:   12134054
Refresh Timeout (Sec):      5
Last connection establishment:  17 hours, 5 minutes, 15 seconds
```

Disabling the Linecard from Sending RDRs

The **silent** command disables the linecard from issuing data records. Both RDRs and NetFlow export packets are suppressed.

Use the **no** form of this command if you want the linecard to send records.

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

Command	Purpose
silent	Disables the linecard from issuing data records.
no silent	Enables the linecard to produce data records.

Disabling RDR Aggregation

In large deployments, if each traffic processor sends its own records separately to the CM, the number of RDRs reaching the CM becomes enormous. Therefore, the Cisco SCE platform aggregates certain RDRs, thus reducing the load on the CM without affecting the usability of the information provided. In essence, the control processor receives records from all traffic processors, but it only sends one record for each reporting period, containing the aggregated data of all CPUs together.

The RDR aggregation feature is relevant only to global records. More specifically, only periodic records are aggregated, because other records relate to events like a single transaction or flow, and cannot be aggregated across processors – if they are aggregated, they loose the required granularity.

Currently the following RDRs are aggregated:

- Virtual Link Usage RDRs (VLURs)
- Link Usage RDRs (LURs)
- Package Usage RDRs (PURs)

You can disable RDR aggregation for either a specific RDR type, or for all aggregated RDRs.

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

Command	Purpose
no periodic-records aggregate-by-cpu [all | LUR | PUR | VLUR]	Disables RDR aggregation for either a specific RDR type, or for all aggregated RDRs.
periodic-records aggregate-by-cpu [all | LUR | PUR | VLUR]	Re-enables RDR aggregation.

To display the current RDR aggregation configuration, use the following command:

Command	Purpose
show interface LineCard 0 periodic-records aggregation	Displays the current RDR aggregation configuration.



Managing Subscribers

Revised: September 23, 2015

Introduction

The Cisco SCE platform is subscriber aware, that is, it can relate traffic and usage to specific customers. This ability to map between IP flows and a specific subscriber allows the system to do the following:

- Maintain the state of each subscriber transmitting traffic through the platform
- Provide usage information for specific subscribers
- Enforce the appropriate policy on subscriber traffic (each subscriber can have a different policy)

- Information About Subscribers, page 10-2
- Importing and Exporting Subscriber Information, page 10-10
- Removing Subscribers and Templates, page 10-13
- Importing and Exporting Anonymous Groups, page 10-19
- Monitoring Subscribers, page 10-20
- Configuring Subscriber Aging, page 10-34
- Managing VPNs and VPN Subscriber Mappings, page 10-37
- Configuring the Cisco SCE Platform/SM Connection, page 10-39

Note Subscriber-aware features are not available for IPv6 traffic. IPv6 traffic is handled in the subscriber-less mode. All traffic is classified to the default subscriber, with a default package.

Information About Subscribers

- What is a Subscriber?, page 10-2
- Subscriber Modes in Service Control Solutions, page 10-3
- Subscriber Mapping Limits, page 10-5
- Aging Subscribers, page 10-5
- Anonymous Groups and Subscriber Templates, page 10-7
- Subscriber Files, page 10-8

What is a Subscriber?

In the Service Control solution, a subscriber is defined as a managed entity on the subscriber side of the Cisco SCE Platform to which accounting and policy are applied individually.

Table 10-1 lists several examples of subscribers in Service Control solutions.

Table 10-1 Subscriber Examples

The Subscriber	**Subscriber Characteristics**	
	Managed Entity	**Subscriber (Entity) Identified By**
DSL residential subscriber	DSL residential user IP address	The list of IP addresses is allocated by a Radius server
Cable residential subscriber	Cable residential user	IP address The list of IP addresses of the CPEs is allocated dynamically by a DHCP server
Owner of a 3G-phone that is subscribed to data services	3G-phone owner	The MS-ISDN, which is dynamically allocated by a Radius server.
A corporate/enterprise customer of the service provider	The corporate/enterprise and the traffic it produces	The set of NAT-ed IP addresses, which are allocated statically
A CMTS	The CMTS and the broadband traffic of the Cable Modem users that connect to the Internet through the CMTS	• A range of IP addresses • A group of VLAN tags

Table 10-1 Subscriber Examples (continued)

The Subscriber	Subscriber Characteristics	
VPN-based subscriber (all or part of a VPN)	VLAN-based subscriber	• A set of IP addresses or ranges in a certain VPN • A VLAN tag or range of VLAN tags
SCMP subscriber	SCMP subscriber	A combination of the following three items: • IP address or range • Manager-Id of the SCMP peer device • Subscriber ID including the GUID Each subscriber is assigned a Manager-Id based on the management entity that created the subscriber. The possible managers are the SM, CLI and an SCMP peer device.

Subscriber Modes in Service Control Solutions

Service Control solutions support several modes of handling subscribers:

- Subscriber-less mode
- Anonymous subscriber mode
- Static subscriber aware mode
- Dynamic subscriber aware mode

Note that not all the Service Control solutions support all modes.

The most basic mode is **Subscriber-less** mode. In this mode, there is no notion of subscriber in the system, and the entire link where the Cisco SCE platform is deployed is treated as a single subscriber. Global Application level analysis (such as total p2p, browsing) can be conducted, as well as global control (such as limiting total p2p to a specified percentage). From a configuration stand point, this is a turnkey system and there is no need to integrate or configure the system from a subscriber perspective.

In **Anonymous subscriber** mode, analysis is performed on an incoming network ID (IP address, VLAN, or VPN ID), as the Cisco SCE platform creates an 'anonymous/on-the-fly' record for each subscriber. This permits analyzing traffic at an individual network ID level (for example, to identify/monitor what a particular 'subscriber' IP is currently doing) as well as control at this level (for example, to limit each subscriber's bandwidth to a specified amount, or block, or redirect). Anonymous-subscriber allows quick visibility into application and protocol usage without OSS integration, and permits the application of a uniform control scheme using predefined templates.

There are two possible **Subscriber Aware** modes. In these modes, subscriber IDs and currently used network IDs are provisioned into the Cisco SCE platform. The Cisco SCE platform can then bind usage to a particular subscriber, and enforce per-subscriber policies on the traffic. Named reports are supported (such as top subscribers with the OSS IDs), quota-tracking (such as tracking a subscriber-quota over time even when network IDs change) as well as dynamic binding of packages to subscribers. The two Subscriber Aware modes are:

- **Static subscriber aware—**The network IDs are static. The system supports the definition of static-subscribers directly to the Cisco SCE platform. This is achieved by using the Cisco SCE platform CLI, and defining the list of subscribers, their network IDs and policy information using interactive configuration or import/export operations.
- **Dynamic subscriber aware—**The network IDs change dynamically for each subscriber login into the Service Provider's network. In this case, subscriber awareness is achieved by integrating with external provisioning systems (either directly or through the SM) to dynamically learn network-ID to subscriber mappings, and distribute them to the Cisco SCE platforms.

Subscriber Database: Capacity and Limits

The capacity of the subscriber database depends on three variables:

The capacity of the subscriber database depends on three variables:

- Subscriber context size—Determined by the specific SML application loaded to the Cisco SCE platform. This size is multiplied by the number of subscribers.
- Available memory per traffic processor—The main memory consumers in a traffic processor are flows and subscribers. The total number of subscribers that can be supported is the number of subscribers per traffic processor multiplied by the number of traffic processors.
- Available memory in the control processor—The control processor holds one entry per subscriber. However, the control database is usually not the limiting factor regarding the number of subscribers, since the control processor entry (context size) is much smaller than the traffic processor entry.

Table 10-2 contains the maximum subscribers capacities for the Cisco SCE platform. Note the following:

- These capacities are the maximum limits imposed by the SCOS. Usually actual numbers would be lower due to large subscriber context size.
- There is a difference between the maximum number of network ID entries and the numbers of specific types of network IDs due to hardware limitations.

Table 10-2 Maximum Number of Subscribers and Network IDs

Subscribers	Network IDs	IP addresses	IP ranges	VLAN tags	Static Subscribers	Virtual Gi with VPN
1,000,000	1,000,000	1,000,000	250,000	4096	250,000	250,000

The maximum number of subscribers are same for all system modes. The maximum number of IPv4 subscribers, IPv6 subscribers, and network IDs in Dual stack mode varies based on the Dual stack mode configuration. If the const-db value of the Dual stack mode is configured as 20, the device supports a maximum of 800,000 IPv4 subscribers and 200,000 IPv6 subscribers. When the Dual stack mode is enabled and all subscribers are dual stack subscribers (subscribers with one IPv4 and one IPv6 address), then the device supports only a maximum of 500,000 dual stack subscribers. For details on configuring the system modes, see the "Configuring the System Mode" section on page 3-17.

Similarly, IP addresses and IP ranges also varies based on the const db value of the Dual stack mode.

The number of static subscribers and VLAN subscribers are not based on the const db value of the Dual stack mode. These are dependent on the type of subscriber logged in first.

Working with Large Numbers of Subscribers

A very large number of subscribers, approaching the upper limit of the SCOS capacity, is more typical in mobile installations. As the actual number of subscriber increases, the impact is expected to be approximately four flows per subscriber.

Actual Maximum Number of Subscribers

As shown in Table 10-2, there is an absolute maximum numbers of subscribers based on SCOS capacity. However, the actual maximum number of subscribers supported is based on the limit specified while loading the SCA BB application.

There are two ways to specify the actual maximum number of subscribers:

- The capacity options mechanism: This mechanism is based on various capacity options supplied by the application. The actual capacity option used is either specified while loading the application or a previously-configured default capacity option is used.
- Specific capacity CLI command (see "Configuring the Actual Maximum Number of Subscribers" section on page 10-32): This specific command overrides the capacity option configured when loading the application. It provides the following options:

  100K, 250K, 500K, 1M

Subscriber Mapping Limits

Refer to Table 10-3 for the maximum number of IP mappings permitted per single subscriber. An IP mapping may be either a single IP address or a range of addresses.

Table 10-3 Maximum Number of IP Mappings per Single Subscriber

Mode	Pure IP	Private IP
Standalone	1024	200
Cascade	200	50

Rate of Creating Anonymous Subscribers

The maximum rate for creation of anonymous subscribers is 360 per second.

Aging Subscribers

Subscribers can be aged automatically by the Cisco SCE platform. 'Aging' is the automatic removal of a subscriber, performed when no traffic sessions assigned to it have been detected for a certain amount of time. The most common usage for aging is for anonymous subscribers, since this is the easiest way to ensure that anonymous subscribers that have logged-out of the network are removed from the Cisco SCE platform and are no longer occupying resources. Aging time can be configured individually for introduced subscribers and for anonymous subscribers.

VPN-Based Subscribers

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. A VPN-based subscriber is VLAN-based.

Most VPN-based subscriber functionality is managed via the SM, with the role of the Cisco SCE platform CLI being more limited.

The Cisco SCE platform CLI can be used to do the following:

- Display VPN-related mappings
- View all automatic VLAN VPNs
- Clear all automatic VLAN VPNs (only VPNs that have no active subscriber mappings).

Automatic VLAN VPNs

The Cisco SCE platform will automatically create a new VPN under the following conditions

- The VPN name does not currently exist

  and

- The VPN name is a number in the range [0 to 4095]

The number is used as the VLAN mapping of the newly created VPN. VLAN mappings cannot be added to automatic VPNs.

Synchronizing Subscriber Information in a Cascade System

In a hot standby, cascade setup with full redundancy, the external provisioning server updates only the active Cisco SCE platform. However, the standby Cisco SCE platform must always be updated with the latest subscriber-related information (login, logout). This is required to minimize information loss in case of failover. In general, the only entity that is allowed to change subscriber information in the standby Cisco SCE platform is the active Cisco SCE platform. The standby Cisco SCE platform does not accept any subscriber operations (it returns a STANDBY_VIOLATION error instead), and it also does not generate any asynchronous subscriber notifications (such as pull-response or logout-notification).

There are only two exceptions to this rule:

- Standby Cisco SCE platform can change subscriber information of the default subscriber.
- Standby Cisco SCE platform can perform subscriber aging

Therefore, when working as a pair, the active Cisco SCE platform constantly updates the standby Cisco SCE platform with external data information. In addition, the standby Cisco SCE platform constantly requests external data information from the active Cisco SCE platform. The synchronization is bi-directional to ensure that the subscriber databases in both Cisco SCE platforms are identical.

Note that external data is only relevant for introduced subscribers (both static and dynamic). It has no meaning for anonymous subscribers or the default subscriber. No more than two minutes of external data information will be lost by the standby Cisco SCE platform if a failover occurs.

The following subscriber information is considered as external data:

- subscriber name
- IP mappings

- tunables
- manager name
- aging time
- lease time
- is-static flag

Only the active Cisco SCE platform communicates with the SM. The SM is aware of the active/standby state of each Cisco SCE platform, and is also aware of a failover.

Specifically, this means the following:

- In push mode, the SM pushes events to the active Cisco SCE platform, which updates the standby Cisco SCE platform.
- In pull mode, only the active Cisco SCE platform pulls subscribers from the SM.
- The standby Cisco SCE platform can create anonymous subscribers based on the updates received from the active Cisco SCE platform, but does not generate pull-response for them.
- If Cisco SCE-SM connection failure, the SM handles the Cisco SCE recovery of the active Cisco SCE platform only. The active Cisco SCE platform propagates the information to the standby Cisco SCE platform.

Anonymous Groups and Subscriber Templates

An anonymous group is a specified IP range, possibly assigned a subscriber template. When an anonymous group is configured, the Cisco SCE platform generates anonymous subscribers for that group when it detects traffic with an IP address that is in the specified IP range. If a subscriber template has been assigned to the group, the anonymous subscribers generated have properties as defined by that template. If no subscriber template has been assigned, the default template is used.

Anonymous groups can have overlapping IP ranges. When the Cisco SCE detects traffic for an IP address which is contained in more than one anonymous group, the group with the longest prefix is used to create the anonymous subscriber for that IP address.

Subscriber templates are identified by a number from 0-199. Subscriber templates 1-199 are defined in csv formatted subscriber template files. However, template #0 cannot change; it always contains the default values.

If an anonymous group is not explicitly assigned a template, the group uses template #0.

Important Information

- Maximum number of anonymous groups—5000 for IPv4 and 1000 for IPv6
- Maximum rate of creating anonymous subscribers—360 per second
- Maximum number of subscriber templates—200 (numbered 0-199)

Subscriber Files

- Subscriber Default csv File Format, page 10-8
- Subscriber Anonymous Groups csv File Format, page 10-9

Individual subscribers, anonymous groups, and subscriber templates may all be defined in csv files. A csv file is a text file in a comma-separated-values format. Microsoft Excel™ can be used to view and create such files. The subscriber data is imported into the system using the appropriate CLI command. The Cisco SCE platform can also export the currently configured subscribers, subscriber templates and anonymous groups to csv-formatted files.

Subscriber csv files and subscriber template csv files are application-specific. Refer to the relevant application documentation for the definition of the file format.

Each line in a csv file should contain either a comment (beginning with the character '#'), or a list of comma-separated fields.

Subscriber csv files are application-specific, but a default format is defined by the Cisco SCE, which is used when the application does not choose to over-ride it. The application might over-ride the format when additional data is desired for each subscriber or subscriber template. Refer to the relevant Service Control Application documentation to see if the application defines a different format.

Subscriber template csv files are application-specific. Refer to the relevant Service Control Application documentation of the file format.

Anonymous groups csv files are not application specific. Their format is described below.

Subscriber Default csv File Format

Each line has the following structure:

name, mappings, packageId

- **Name**—is the subscriber name
- **Mappings**—contains one or more mappings, specifying the Tunnel IDs or IP addresses mapped to this subscriber. Multiple mappings are separated by semi-colon. Tunnel IDs and IP address/range cannot be specified for the same subscriber. The following mapping formats are supported:
  - IP address—in dotted decimal notation. Example: 10.3.4.5
  - IP address range—dotted decimal, followed by the amount of significant bits. Note that the non-significant bits (as determined by the mask) must be set to zero. Example: 10.3.0.0/16. Example for a bad range: 10.1.1.1/24 (Should be 10.1.1.0/24).
- **PackageId**—the ID of the package to which the subscriber is assigned

Here is an example of a subscriber csv file in the default format:

```
# A comment line
sub7, 10.1.7.0/24, 1
sub8, 10.1.11.32, 1
sub9, 10.2.22.10, 2
sub10, 10.3.33.10, 2
sub11, 10.4.44.10, 1
sub12, 10.1.11.90; 10.3.0.0/16, 2
```

IPv6 Subscriber csv File Format

Each line has the following structure:

name, mappings, mappings_ipv6, packageId

- **name**—The is the subscriber name.
- **mappings**—This contains one or more mappings, specifying the Tunnel IDs or IP addresses mapped to the corresponding subscriber. Multiple mappings are separated by a semi-colon. Tunnel IDs and IP address or IP range cannot be specified for the same subscriber. The following mapping formats are supported:
  - IP address (In dotted decimal notation. Example: 10.3.4.5.)
  - IP address range (Dotted decimal, followed by the amount of significant bits. Note that the nonsignificant bits (as determined by the mask) must be set to zero. Example: 10.3.0.0/16. Example of an incorrect range: 10.1.1.1/24 (Should be 10.1.1.0/24).)
- **mappings_IPv6**—This contains one or more IPv6 mappings. Multiple mappings are separated by a semi-colon. The following mapping format is supported:
  - IPv6 address—(In hexadecimal colon notation. Example: 2001:a:d:f::/64. The valid range is from 32 to 64.)
- **packageId**—(The ID of the package to which the subscriber is assigned.)

Here is an example of a subscriber csv file in the default format:

```
# A comment line
sub7, 10.1.7.0/24, 2001:a:d:f::/64,1
sub8, 10.1.11.32, 2001:a:d:d::/64,1
sub9, 10.2.22.10, ,2
sub10, , 2001:a:d:f::/64, 2
sub11, 10.4.44.10, 1
sub12, 10.1.11.90; 10.3.0.0/16, 2
```

Subscriber Anonymous Groups csv File Format

Each line has the following structure:

name, IP-range or IPv6-range, template-index, manager-name (optional)

- **name**—This is the anonymous group name.
- **IP-range**—IPv4 or IPv6, followed by the amount of significant bits. In case of IPv6, the range is between 32 to 64. Example: 10.3.0.0/16, 2001:a:d:f::/64.
- **template-index**—This is the index of the subscriber template to be used by subscribers belonging to this anonymous group.
- **manager-name** (optional)—This is either SM or the name of the SCMP peer. Use "SM" to pull subscribers from the SM (if it exists). If not specified, "SM" is assumed.

Here is an example of an anonymous groups csv file:

```
# Yet another comment line
anon1, 10.1.1.0/24, 1, 1
anon2, 2001:a:d:f::/64, 2, 2
anon3, 10.1.3.0/32, 3, 3
anon4, 2001:a:d:d::/64, 3, 3
anon5, 10.1.5.0/31, 2
anon6, 10.1.6.0/30, 1
anon7, 0.0.0.0/0, 1
```

Importing and Exporting Subscriber Information

- Editing the subaware.pro File, page 10-10
- Options, page 10-11
- How to Import Subscriber Information, page 10-11
- How to Export Subscriber Information, page 10-11
- How to Export a Subscriber Template, page 10-12

Editing the subaware.pro File

You can import and export IPv6 subscribers. Before importing and exporting IPv6 subscribers for the first time, you need to modify the subaware.pro file located at /apps/data/scos/system/p3hidden/um/ssu/. To modify the file, follow these steps:

Step 1 From the Linux shell, open the file subaware.pro:

```
#>vi /apps/data/scos/system/p3hidden/um/ssu/subaware.pro
```

Step 2 Remove the pound symbol before the following line:

```
#smm.ssu.maindecoder.fields=name,mappings,mappings_ipv6,tuneable.packageId,tuneable.upVlinkId,tuneable.downVlinkId,tuneable.monitor
```

Step 3 Include a pound symbol before the following line:

```
smm.ssu.maindecoder.fields=name,mappings,tuneable.packageId,tuneable.upVlinkId,tuneable.downVlinkId,tuneable.monitor
```

Step 4 Save the subaware.pro file.

Step 5 From the SCE(config if)# prompt, type the following command to load the subaware.pro file:

```
SCE8000(config if)#>subscriber load property-file
```

Use the following commands to import subscriber data from csv files and to export subscriber data to these files:

- **subscriber import csv-file**
- **subscriber export csv-file [all]**
- **subscriber anonymous-group import csv-file**
- **subscriber anonymous-group export csv-file**
- **subscriber template import csv-file**
- **subscriber template export csv-file**

These subscriber management commands are LineCard interface commands. Make sure that you are in LineCard Interface command mode.

Options

The following options are available:

filename—The name of the csv file.

additivemapping—Appends new subscriber mappings with the existing ones.

How to Import Subscriber Information

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

Command	Purpose
subscriber import csv-file *filename*	Imports the subscriber information from the specified file. If the information in the imported file is not valid, the command will fail during the verification process before it is actually applied.
subscriber import csv-file *filename additivemapping*	Imports the subscriber information from the specified file and appends the new subscriber mappings with the existing ones.

Note The csv file used to import subscriber information should have the same fields as those in the subaware.pro file.

How to Export Subscriber Information

Note To export subscribers managed by the SM, use the SM GUI or CLU (see the *Cisco Service Control Subscriber Manager User Guide.*)

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

Command	Purpose
subscriber export csv-file *filename* **[all]**	Exports all the static and dynamic subscribers information to the specified file.
subscriber export csv-file *filename*	Exports only the static subscribers information to the specified file.

How to Import a Subscriber Template

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

Command	Purpose
subscriber template import csv-file *filename*	Imports the subscriber template from the specified file.

How to Export a Subscriber Template

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

Command	Purpose
subscriber template export csv-file *filename*	Exports the subscriber template to the specified file.

Removing Subscribers and Templates

- How to Remove a Specific Subscriber, page 10-13
- How to Remove All Introduced Subscribers, page 10-14
- How to Remove a Specific Anonymous Subscriber Group, page 10-14
- How to Remove All Anonymous Subscriber Groups, page 10-14
- How to Remove All the Anonymous Subscribers, page 10-14
- How to Remove All Subscriber Templates, page 10-15
- How to Remove Subscribers by Device, page 10-15

Use the following commands to remove all subscribers, anonymous groups, or subscriber templates from the system.

- **no subscriber all**
- **no subscriber anonymous-group all**
- **clear interface linecard subscriber anonymous**
- **default subscriber template all**

Use the following commands to remove a specific subscriber or anonymous group from the system.

- **no subscriber name**
- **no subscriber anonymous-group name**

These subscriber management commands are LineCard interface commands (with the exception of the **clear interface linecard subscriber anonymous** command, which is a Privileged Exec command). Make sure that you are in LineCard Interface command mode, and that the SCE(config if)# prompt appears in the command line.

How to Remove a Specific Subscriber

Options

The following option is available:

subscriber-name—The name of the subscriber to be removed

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

Command	Purpose
no subscriber name *subscriber-name*	Removes the specified subscriber.

How to Remove All Introduced Subscribers

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

Command	Purpose
no subscriber all	Removes all introduced subscribers. • *subscriber-name* specifies the name of the subscriber to be removed.

How to Remove a Specific Anonymous Subscriber Group

Options

The following option is available:

group-name—The name of the anonymous subscriber group to be removed

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

Command	Purpose
no subscriber anonymous-group name *group-name*	Removes the specified anonymous subscriber group.

How to Remove All Anonymous Subscriber Groups

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

Command	Purpose
no subscriber anonymous-group all	Removes all anonymous subscriber groups.

How to Remove All the Anonymous Subscribers

From the SCE# prompt, type:

Command	Purpose
clear interface linecard 0 subscriber anonymous all	Removes all anonymous subscribers. **Note** The **clear subscriber anonymous** command is a **Privileged Exec** command.

Caution Because the **clear interface linecard subscriber anonymous all** command clears all the anonymous subscribers in the Cisco SCE, do not use the command in a production environment. Using this command in a production environment impacts anonymous subscribers' accountability. Use the command only when the linecard interface is shut down.

How to Remove All Subscriber Templates

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

Command	Purpose
default subscriber template all	Removes all subscriber templates. All anonymous subscribers will be assigned to the default subscriber template.

Removing VPN-based Subscribers

All VPN-based subscribers must be cleared to change the tunneling mode. If there are VPN-based subscribers that the SM cannot remove for some reason (for example, if there is no communication between the SM and the Cisco SCE platform), use this command.

Note Use this command ONLY when the Cisco SCE platform is disconnected from the SM.

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

Command	Purpose
no subscriber all with-vpn-mappings	Clears all VPN-based subscribers.

How to Remove Subscribers by Device

- How to Remove Subscribers from the SM, page 10-15
- How to Remove Subscribers from a Specified SCMP Peer Device, page 10-17

You can remove all subscribers managed by a specified device. The device can be either of the following:

- The SM
- A specified SCMP peer device

How to Remove Subscribers from the SM

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

Command	Purpose
no subscriber sm all	Clears all subscribers from the SM.

How to Remove Subscribers from a Specified SCMP Peer Device

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

Command	Purpose
no subscriber scmp name *peer-device-name* **all**	Clears all subscribers from the specified SCMP peer device. • *peer-device-name* specifies the name of the SCMP peer device from which to clear the subscribers.

Creating Anonymous Groups

You can create anonymous groups in two ways:

- Define the group, along with the related IP address range, using the **subscriber anonymous-group** command. The Cisco SCE platform then generates anonymous subscribers for that group when it detects traffic with an IP address that is in the specified IP range.
- Create the group by importing anonymous groups from a csv file.

  Groups can also be exported to a csv file.
- Maximum creation rate of anonymous subscribers is 360 per second on the Cisco SCE 8000 and 180 per second on the Cisco SCE2020.

For information on deleting anonymous groups, see "How to Remove a Specific Anonymous Subscriber Group" section on page 10-14 and "How to Remove All Anonymous Subscriber Groups" section on page 10-14

Defining Anonymous Groups

Use this command to define an anonymous group, assigning the following to the group created:

- group name
- range of IP addresses
- subscriber template to be assigned to all subscribers within that IP range (optional)

How to Define an Anonymous Group

Options

The following options are available:

- **group-name—**Name to be assigned to the anonymous group.
- **range—**Range of IP addresses that defines this group.
- **template**—Number of the subscriber template to assign to all subscribers in this group.

  Default = 0

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

Command	Purpose
subscriber anonymous-group name *group-name* **ip-range** ***range*** **[template** ***template*****]**	Defines an anonymous group for IPv4. If no template is specified, the default template is applied to all subscribers in this group.
subscriber anonymous-group name *group-name* **ipv6-range** *ipv6-address/prefix* **[template** *template*]	Defines an anonymous group for IPv6. If no template is specified, the default template is applied to all subscribers in this group.

Importing and Exporting Anonymous Groups

- How to Import Anonymous Groups, page 10-19
- How to Export Anonymous Groups, page 10-19

How to Import Anonymous Groups

Options

The following option is available:

- **filename—**Name of the csv file.

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

Command	Purpose
subscriber anonymous-group import csv-file *filename*	Creates anonymous groups by importing anonymous groups from the specified csv file. Imported anonymous groups information is added to the existing anonymous groups information. It does not overwrite the existing data. The Cisco SCE platform can support a maximum of 5000 anonymous groups.

How to Export Anonymous Groups

Options

The following option is available:

- **filename—**Name of the csv file.

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

Command	Purpose
subscriber anonymous-group export csv-file *filename*	Exports all existing anonymous groups to the specified csv file.

Monitoring Subscribers

- How to Monitor the Subscriber Database, page 10-20
- Displaying Subscribers, page 10-22
- How to Display Subscriber Information, page 10-27
- Displaying Anonymous Subscriber Information, page 10-29

The CLI provides several commands that allow you to monitor subscribers. These commands can be used to display information regarding the following:

- Subscriber Database
- All subscribers meeting various criteria
- Individual subscriber information, such as properties and mappings
- Anonymous subscribers

Subscribers may be introduced to the Cisco SCE platform via the Cisco SCE platform CLI or via the Subscriber Manager. The monitoring commands may be used to monitor all subscribers and subscriber information, regardless of how the subscribers were introduced to the system.

Note that these commands are all in Viewer mode. Make sure that you are in the proper mode and that the SCE> prompt appears in the command line. Note also that you must specify ‘linecard 0’ in these commands.

How to Monitor the Subscriber Database

- How to Display the Subscriber Database Counters, page 10-21
- Clearing the Subscriber Database Counters, page 10-22

Use the following commands to display statistics about the subscriber database, and to clear the “**total**” and “**maximum**” counters.

- **show interface linecard 0 subscriber db counters**

  The following counters are displayed:

  - Current number of subscribers
  - Current number of introduced subscribers
  - Current number of anonymous subscribers
  - Current number of active subscribers (with active traffic sessions)
  - Current number of subscribers with mappings
  - Current number of IP mappings
  - Current number of vlan mappings
  - Max number of subscribers that can be introduced
  - Max number of subscribers with mappings
  - Max number of subscribers with mappings date / time
  - Total aggregated number introduced

  - Total number of aged subscribers
  - Total number of pull events
  - Number of traffic sessions currently assigned to the default subscriber
- clear interface linecard 0 subscriber db counters

How to Display the Subscriber Database Counters

From the SCE# prompt, type:

Command	Purpose
show interface linecard 0 subscriber db counters	Displays the subscriber database counters.

Monitoring the Subscriber Database: Example

The following example shows the output from this command.

```
SCE#show interface linecard 0 subscriber db counters
Current values:
===============
Subscribers: 249999 used out of 999999 max.
Introduced/Pulled subscribers: 999999.
Anonymous subscribers: 0.
Subscribers with mappings: 0 used out of 249999 max.
Single non-VPN IP mappings: 0.
Non-VPN IP Range mappings: 0.
IP Range over VPN mappings: 0.
Single IP over VPN mappings: 0.
VLAN based VPNs with subscribers: 0 used out of 4095.
Subscribers with open sessions: 243562.
Subscribers with TIR mappings: 0.
Sessions mapped to the default subscriber: 2.
Peak values:
============
Peak number of subscribers with mappings: 249999
Peak number occurred at: 15:54:06 UTC TUE May 13 2008
Peak number cleared at: 07:47:49 UTC SUN May 11 2008
Event counters:
===============
Subscriber introduced: 249999.
Subscriber pulled: 0.
Subscriber aged: 0.
Pull-request notifications sent: 0.
Pull-request by ID notifications sent: 0.
Subscriber pulled by ID: 0.
State notifications sent: 0.
Logout notifications sent: 0.
Subscriber mapping TIR contradictions: 0.
```

Clearing the Subscriber Database Counters

From the SCE# prompt, type:

Command	Purpose
clear interface linecard 0 subscriber db counters	Clears the "**total**" and "**maximum**" counters.

Displaying Subscribers

- Displaying Subscribers: All Current Subscriber Names, page 10-22
- Displaying Subscribers: By Subscriber Property or Prefix, page 10-23
- How to Display Subscribers: By Mapping (IP Address, VPN, or VLAN ID), page 10-25

You can display the names of all subscribers.

You can also display specific subscriber name(s) that meet various criteria:

- A subscriber property is equal to, larger than, or smaller than a specified value.
- Subscriber name matches a specific prefix or suffix.
- Mapped to a specified IP address range.
- Mapped to a specified VLAN ID.

Use the following commands to display subscribers:

- show interface linecard 0 subscriber all-names
- show interface linecard 0 subscriber [amount] [prefix 'prefix'] [property 'propertyname' equals|greater-than|less-than 'property-val']
- show interface linecard 0 subscriber [amount] prefix 'prefix'
- show interface linecard 0 subscriber [amount] suffix 'suffix'
- show interface linecard 0 subscriber mapping IP 'iprange' [VPN 'vpn-name']
- show interface linecard 0 subscriber [amount] mapping intersecting IP 'iprange [VPN 'vpn-name']
- show interface linecard 0 subscriber mapping VLAN-id 'VLAN-id'

Displaying Subscribers: All Current Subscriber Names

You can display the names of all subscribers currently in the Cisco SCE subscriber database.

From the SCE> prompt, type:

Command	Purpose
show interface linecard 0 subscriber all-names	Displays the names of all subscribers currently in the Cisco SCE subscriber database.

Displaying Subscribers: By Subscriber Property or Prefix

You can search for all subscribers that match a specified value of one of the subscriber properties, or are greater than or less than the specified value. You can also search for all subscribers that match a specified prefix. You can also find out how many subscribers match any one of these criteria, rather than displaying all the actual subscriber names.

- How to display subscribers that match a specified value of a subscriber property, page 10-23
- How to display subscribers that are greater than or less than a specified value of a subscriber property, page 10-23
- How to display subscribers that match a specified prefix, page 10-24
- How to display subscribers that match a specified suffix, page 10-24
- How to display the number of subscribers that match a specified value of a subscriber property, page 10-24
- How to display the number of subscribers that are greater than or less than a specified value of a subscriber property, page 10-24
- How to display the number of subscribers that match a specified prefix, page 10-25

How to display subscribers that match a specified value of a subscriber property

Options

The following options are available:

- **propertyname—**Name of the subscriber property to match
- **property-val—**Value of that subscriber property to match

From the SCE> prompt, type:

Command	Purpose
show interface linecard 0 subscriber property *propertyname* **equals** *property-val*	Displays subscribers that match a specified value of a subscriber property.

How to display subscribers that are greater than or less than a specified value of a subscriber property

Options

The following options are available:

- **propertyname—**Name of the subscriber property to match
- **property-val—**Value of that subscriber property to match

From the SCE> prompt, type:

Command	Purpose
show interface linecard 0 subscriber property *propertyname* **greater-than|less-than** *property-val*	Displays subscribers that are greater than or less than a specified value of a subscriber property.

How to display subscribers that match a specified prefix

Options

The following options are available:

- **prefix—**Subscriber prefix to match

From the SCE> prompt, type:

Command	Purpose
show interface linecard 0 subscriber prefix *prefix*	Displays subscribers that match a specified prefix.

How to display subscribers that match a specified suffix

Options

The following options are available:

- **suffix—**Subscriber suffix to match

From the SCE> prompt, type:

Command	Purpose
show interface linecard 0 subscriber suffix *suffix*	Displays subscribers that match a specified suffix.

How to display the number of subscribers that match a specified value of a subscriber property

Options

The following options are available:

- **propertyname—**Name of the subscriber property to match
- **property-val—**Value of that subscriber property to match

From the SCE> prompt, type:

Command	Purpose
show interface linecard 0 subscriber amount property *propertyname* **equals** *property-val*	Displays the number of subscribers that match a specified value of a subscriber property.

How to display the number of subscribers that are greater than or less than a specified value of a subscriber property

Options

The following options are available:

- **propertyname—**Name of the subscriber property to match
- **property-val—**Value of that subscriber property to match

From the SCE> prompt, type:

Command	Purpose
show interface linecard 0 subscriber amount property *propertyname* **greater-than|less-than** *property-val*	Displays the number of subscribers that are greater than or less than a specified value of a subscriber property.

How to display the number of subscribers that match a specified prefix

Options

The following options are available:

- **prefix—**Subscriber prefix to match

From the SCE> prompt, type:

Command	Purpose
show interface linecard 0 subscriber amount prefix *prefix*	Displays the number of subscribers that match a specified prefix.

How to Display Subscribers: By Mapping (IP Address, VPN, or VLAN ID)

- How to display subscribers that are mapped to a specified IP address, or range of IP addresses, page 10-26
- How to display subscribers that are mapped to IP addresses that are included in a given IP address or IP range, page 10-26
- How to display subscribers that are mapped to a specified VLAN ID, page 10-26
- How to display subscribers with no mapping, page 10-26
- How to display the number of subscribers that are mapped to a specified VLAN ID, page 10-27
- How to display the number of subscribers with no mapping, page 10-27

You can display the subscribers who are mapped to any of the following:

- A specified IP address, or range of IP addresses
- IP addresses intersecting a given IP address or IP range
- A specified VLAN ID
- A specified VPN
- no mapping

You can also display just the number of subscribers with a specified mapping, rather than listing the actual subscribers.

How to display subscribers that are mapped to a specified IP address, or range of IP addresses

Options

The following options are available:

- **ip-range—**IP address (x.x.x.x) or range of IP addresses (x.x.x.x/y) to match
- **vpn-name** (optional)—The name of the VPN in which to search for the IP address

From the SCE> prompt, type:

Command	Purpose
show interface linecard 0 subscriber mapping IP *ip-range* [**VPN** *vpn-name*]	Displays subscribers that are mapped to a specified IP address, or range of IP addresses.

How to display subscribers that are mapped to IP addresses that are included in a given IP address or IP range

Options

The following options are available:

- **ip-range—**IP address (x.x.x.x) or range of IP addresses (x.x.x.x/y) to match
- **vpn-name** (optional)—The name of the VPN in which to search for the IP address

From the SCE> prompt, type:

Command	Purpose
show interface linecard 0 subscriber mapping included-in IP *ip-range* [**VPN** *vpn-name*]	Displays subscribers that are mapped to IP addresses that are included in a given IP address or IP range.

How to display subscribers that are mapped to a specified VLAN ID

Options

The following options are available:

- **VLAN-id—**VLAN ID to match

From the SCE> prompt, type:

Command	Purpose
show interface linecard 0 subscriber mapping VLAN-id *VLAN-id*	Displays subscribers that are mapped to a specified VLAN ID.

How to display subscribers with no mapping

From the SCE> prompt, type:

Command	Purpose
show interface linecard 0 subscriber mapping none	Displays subscribers with no mapping.

How to display the number of subscribers that are mapped to a specified VLAN ID

Options

The following options are available:

- **VLAN-id—**VLAN ID to match

From the SCE> prompt, type:

Command	Purpose
show interface linecard 0 subscriber amount mapping VLAN-id *VLAN-id*	Displays the number of subscribers that are mapped to a specified VLAN ID.

How to display the number of subscribers with no mapping

From the SCE> prompt, type:

Command	Purpose
show interface linecard 0 subscriber amount mapping none	Displays the number of subscribers with no mapping.

How to Display Subscriber Information

- How to Display a Listing of Subscriber Properties, page 10-28
- How to Display Complete Information for a Specified Subscriber, page 10-28
- How to Display Values of Subscriber Properties for a Specified Subscriber, page 10-28
- How to Display Mappings for a Specified Subscriber, page 10-28
- How to Display OS Counters for a Specified Subscriber, page 10-29

You can display the following information about a specified subscriber:

- values of the various subscriber properties
- mappings (IP address or VLAN-ID)
- OS counters:
  - current number of flows
  - bandwidth

Use the following commands to display subscriber information:

- **show interface linecard 0 subscriber properties**
- **show interface linecard 0 subscriber name ‘name’**
- **show interface linecard 0 subscriber name ‘name’ mappings**
- **show interface linecard 0 subscriber name ‘name’ counters**
- **show interface linecard 0 subscriber name ‘name’ properties**

How to Display a Listing of Subscriber Properties

From the SCE> prompt, type:

Command	Purpose
show interface linecard 0 subscriber properties	Displays a listing of subscriber properties.

How to Display Complete Information for a Specified Subscriber

Options

The following options are available:

- **name—**Subscriber name

From the SCE> prompt, type:

Command	Purpose
show interface linecard 0 subscriber name *name*	Displays complete information for a specified subscriber, including all values of subscriber properties and mappings.

How to Display Values of Subscriber Properties for a Specified Subscriber

Options

The following options are available:

- **name—**Subscriber name

From the SCE> prompt, type:

Command	Purpose
show interface linecard 0 subscriber name *name* **properties**	Displays values of subscriber properties for a specified subscriber.

How to Display Mappings for a Specified Subscriber

Options

The following options are available:

- **name—**Subscriber name

From the SCE> prompt, type:

Command	Purpose
show interface linecard 0 subscriber name *name* **mappings**	Displays mappings for a specified subscriber.

How to Display OS Counters for a Specified Subscriber

Options

The following options are available:

- **name—**Subscriber name

From the SCE> prompt, type:

Command	Purpose
show interface linecard 0 subscriber name *name* **counters**	Display OS counters for a specified subscriber.

Displaying Anonymous Subscriber Information

- How to Display Currently Configured Anonymous Groups, page 10-29
- How to Display Currently Configured Templates for Anonymous Groups, page 10-30
- How to Display Current Configuration for a Specified Anonymous Group, page 10-30
- How to Display Subscribers in a Specified Anonymous Group, page 10-30
- How to Display All Subscribers Currently in Anonymous Groups, page 10-30
- How to Display the Number of Subscribers in a Specified Anonymous Group, page 10-31
- How to Display the Total Number of Subscribers in All Anonymous Groups, page 10-31

You can display the following information regarding the anonymous subscriber groups:

- aging (see "How to Display Aging for Anonymous Group Subscribers" section on page 10-35)
- currently configured anonymous groups
- currently configured subscriber templates
- configuration of a specified anonymous group
- number of subscribers in a specified anonymous group, or in all anonymous groups

Use the following commands to display anonymous subscriber information:

- **show interface linecard 0 subscriber templates [index]**
- **show interface linecard 0 subscriber anonymous-group [all] [name 'groupname']**
- **show interface linecard 0 subscriber amount anonymous [name 'groupname']**
- **show interface linecard 0 subscriber anonymous [name 'groupname']**

How to Display Currently Configured Anonymous Groups

From the SCE> prompt, type:

Command	Purpose
show interface linecard 0 subscriber anonymous-group all	Displays currently configured anonymous groups.

How to Display Currently Configured Templates for Anonymous Groups

From the SCE> prompt, type:

Command	Purpose
show interface linecard 0 subscriber templates	Display currently configured templates for anonymous groups.

How to Display Current Configuration for a Specified Anonymous Group

Options

The following options are available:

- **group-name—**Name of the anonymous subscriber group.

From the SCE> prompt, type:

Command	Purpose
show interface linecard 0 subscriber anonymous-group name *group-name*	Displays current configuration for a specified anonymous group.

How to Display Subscribers in a Specified Anonymous Group

Options

The following options are available:

- **group-name—**Name of the anonymous subscriber group.

From the SCE> prompt, type:

Command	Purpose
show interface linecard 0 subscriber anonymous name *group-name*	Displays subscribers in a specified anonymous group.

How to Display All Subscribers Currently in Anonymous Groups

From the SCE> prompt, type:

Command	Purpose
show interface linecard 0 subscriber anonymous	Displays all subscribers currently in anonymous groups.

How to Display the Number of Subscribers in a Specified Anonymous Group

Options

The following options are available:

- **group-name—**Name of the anonymous subscriber group.

From the SCE> prompt, type:

Command	Purpose
show interface linecard 0 subscriber amount anonymous name *group-name*	Displays the number of subscribers in a specified anonymous group.

How to Display the Total Number of Subscribers in All Anonymous Groups

From the SCE> prompt, type:

Command	Purpose
show interface linecard 0 subscriber amount anonymous	Displays the total number of subscribers in all anonymous groups.

Configuring the Actual Maximum Number of Subscribers

The actual maximum number of subscribers supported is based on the capacity option specified when loading the SCA BB application. If no capacity option is specified, the user-configured default capacity is used. However, you can override this capacity option using the following commands.

Note the following:

- You must configure the override before you load the application (PQI file). The configured maximum number of subscribers takes effect when the next **load** command is executed.
- If you have disabled the capacity option and then the next time you load a new application you want to use the capacity option, you must re-enable the capacity option before loading the application file (see "How to Restore the Configured Capacity Option" section on page 10-33).
- Use the **show subscriber max-subscribers** command to see what the current maximum number of subscribers is and whether the capacity option is enabled or disabled (see "How to Monitor the Maximum Number of Subscribers" section on page 10-33.)

How to Override the Configured Capacity Option

The default maximum number of subscribers is 250, 000.

Step 1 If a policy configuration (PQB file) has been applied on Cisco SCE platform, use the SCA BB console to retrieve it and save it before proceeding.

Step 2 Select the maximum number of subscribers:

From the SCE(config if)# prompt, type **subscriber max-subscribers (40K | 80K | 120K | 200K)** and press **Enter**.

Step 3 Disable the application capacity option:

From the SCE(config if)# prompt, type **subscriber capacity-options disable** and press **Enter**.

Step 4 Install the new application (PQI) file. (The configured subscriber maximum takes effect only after a new application file has been loaded.)

Step 5 If you saved the policy configuration (PQB file), apply it to the Cisco SCE platform using the SCA BB console.

How to Override the Configured Capacity Option in a Cascade Setup

The default maximum number of subscribers is 250, 000. To change the maximum actual number of subscribers in cascade topology, configure both the standby and active Cisco SCE devices. To change this value in cascade topology, follow these steps:

Step 1 If a policy configuration (PQB file) has been applied on the standby Cisco SCE platform, use the Cisco SCA BB console to retrieve it and save it before proceeding.

Step 2 Select the maximum number of subscribers:

```
SCE(config if)# subscriber max-subscribers {40K | 80K | 120K | 200K}
```

Step 3 Disable the application capacity option:

```
SCE(config if)# subscriber capacity-options disable
```

Step 4 Install the new application (PQI) file. (The configured subscriber maximum takes effect only after a new application file has been loaded.)

Step 5 If you saved the policy configuration (PQB file), apply it to the Cisco SCE platform using the Cisco SCA BB console.

If a policy configuration (PQB file) has been applied on the active Cisco SCE platform, use the Cisco SCA BB console to retrieve it and save it before proceeding. Repeat from Step 1 to Step 5 on the active Cisco SCE platform.

Note When you install the PQI file, the Package ID parameter of all the existing subscribers become 0. The subscribers are automatically deleted when the aging period of each subscriber elapses.

How to Restore the Configured Capacity Option

Step 1 If a policy configuration (PQB file) has been applied on Cisco SCE platform, use the SCA BB console to retrieve it and save it before proceeding.

Step 2 From the SCE(config if)# prompt, type **subscriber capacity-options enable** and press **Enter**.

Step 3 Install the new application (PQI) file. (The configured subscriber maximum takes effect only after a new application file has been loaded.)

Step 4 If you saved the policy configuration (PQB file), apply it to the Cisco SCE platform using the SCA BB console.

How to Monitor the Maximum Number of Subscribers

From the SCE> prompt, type:

Command	Purpose
show interface linecard 0 subscriber max-subscribers	Shows the configured maximum number of subscribers and whether the capacity options are in effect or have been overridden

Configuring Subscriber Aging

- How to Enable Aging for Anonymous Group Subscribers, page 10-34
- How to Enable Aging for Introduced Subscribers, page 10-34
- How to Disable Aging for Anonymous Group Subscribers, page 10-34
- How to Disable Aging for Introduced Subscribers, page 10-35
- How to Set the Aging Timeout Period for Anonymous Group Subscribers, page 10-35
- How to Set the Aging Timeout Period for Introduced Subscribers, page 10-35
- How to Display Aging for Anonymous Group Subscribers, page 10-35
- How to Display Aging for Introduced Subscribers, page 10-36

As explained previously ("Aging Subscribers" section on page 10-5, aging is the automatic removal of a subscriber when no traffic sessions assigned to it have been detected for a certain amount of time. Aging may be enabled or disabled, and the aging timeout period (in minutes) can be specified.

Aging can be configured separately for introduced subscribers and for anonymous subscribers.

Use the following commands to configure and monitor aging.

- **[no] subscriber aging**
- **subscriber aging timeout**
- **show interface linecard 0 subscriber aging**

How to Enable Aging for Anonymous Group Subscribers

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

Command	Purpose
subscriber aging anonymous	Enables aging for anonymous group subscribers.

How to Enable Aging for Introduced Subscribers

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

Command	Purpose
subscriber aging introduced	Enables aging for introduced subscribers.

How to Disable Aging for Anonymous Group Subscribers

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

Command	Purpose
no subscriber aging anonymous	Disables aging for anonymous group subscribers.

How to Disable Aging for Introduced Subscribers

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

Command	Purpose
no subscriber aging introduced	Disables aging for introduced subscribers.

How to Set the Aging Timeout Period for Anonymous Group Subscribers

Options

The following option is available:

- **aging-time—**The time interval, in minutes, after which an inactive subscriber will be aged (2-14400).

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

Command	Purpose
no subscriber aging anonymous timeout *aging-time*	Sets the aging timeout period for anonymous group subscribers.

How to Set the Aging Timeout Period for Introduced Subscribers

Options

The following option is available:

- **aging-time—**The time interval, in minutes, after which an inactive subscriber will be aged (2-14400).

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

Command	Purpose
no subscriber aging introduced timeout *aging-time*	Sets the aging timeout period for introduced subscribers.

How to Display Aging for Anonymous Group Subscribers

From the SCE> prompt, type:

Command	Purpose
show interface linecard 0 subscriber aging anonymous	Displays aging of anonymous group subscribers.

How to Display Aging for Introduced Subscribers

From the SCE> prompt, type:

Command	Purpose
show interface linecard 0 subscriber aging introduced	Displays aging for introduced subscribers.

Managing VPNs and VPN Subscriber Mappings

- How to Display VPN-Related Mappings, page 10-37
- How to Clear Automatic VPNs, page 10-37

How to Display VPN-Related Mappings

- How to Display Mappings for a Specified VPN, page 10-37
- How to Display a Listing of All VPNs, page 10-37

Use the following Viewer commands to display VPNs and VPN 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 currently logged-in VPNs that were created automatically

How to Display Mappings for a Specified VPN

Options

The following option is available:

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

From the SCE> prompt, type:

Command	Purpose
show interface linecard 0 VPN name *vpn-name*	Displays mapping for a specified VPN.

How to Display a Listing of All VPNs

From the SCE> 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

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

How to Clear Automatic VPNs

From the SCE# prompt, type:

Command	Purpose
clear interface linecard 0 VPN automatic	Removes all VLAN VPNs that were created automatically by the Cisco SCE platform. (Only removes VPNs that have no active subscriber mappings).

Configuring the Cisco SCE Platform/SM Connection

- Configuring the Behavior of the Cisco SCE Platform in Case of Failure of the SM, page 10-39
- Configuring the SM-SCE Platform Connection Timeout, page 10-40

The user can configure the behavior of the Cisco SCE platform in case of failure of the Subscriber Manager (SM):

- If SM functionality is critical to the operation of the system—Configure the desired behavior of the Cisco SCE platform if any loss of connection with the SM (may be due either to failure of the SM or failure of the connection itself).
- If SM functionality is not critical to the operation of the system—No action needs to be configured In this case you can specify that the system operational-status of the Cisco SCE platform should be 'warning' when the link is down.

Configuring the Behavior of the Cisco SCE Platform in Case of Failure of the SM

Options

The following options are available:

- **action**—The specified action will be performed in case of loss of connection between the Cisco SCE platform and the SM.

  Possible actions are:

  - **force-failure—**Force failure of Cisco SCE platform. The Cisco SCE platform then acts according to the behavior configured for the failure state.
  - **remove-mappings—**Remove all current subscriber mappings.
  - **shut—**The Cisco SCE platform shuts down and quits providing service.
  - **none** (default)—Take no action.
- **warning**—The system operational-status of the Cisco SCE platform should be 'warning' in case of loss of connection between the Cisco SCE platform and the SM. No action is taken.

To specify the action that the Cisco SCE platform will perform if the Cisco SCE-SM connection fails, use this command.

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

Command	Purpose
subscriber sm-connection-failure action [force-failure|none|remove-mappings|shut]	Specifies the action that the Cisco SCE platform will perform if the Cisco SCE-SM connection fails.
subscriber sm-connection-failure warning	Specifies that the system operational-status of the Cisco SCE platform should be 'warning' if the Cisco SCE-SM connection fails.

Configuring the SM-SCE Platform Connection Timeout

You can also configure the timeout interval; the length of time that the SM-SCE platform connection is disrupted before a failed connection is recognized and the configured behavior is applied.

Options

The following option is available:

- **interval—**The timeout interval in seconds

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

Command	Purpose
subscriber sm-connection-failure timeout *interval*	Configures the connection timeout.



Redundancy and Failover

Revised: September 23, 2015

Introduction

- Redundancy and Failover, page 11-2
- Link Failure Reflection, page 11-5
- Hot Standby and Failover, page 11-6
- Recovery, page 11-11
- CLI Commands for Cascaded Systems, page 11-13
- Configuring Forced Failure, page 11-18
- System Upgrades, page 11-19

Redundancy and Failover

- Terminology and Definitions, page 11-2
- Redundant Topologies, page 11-2
- External Bypass, page 11-3
- In-line Dual Link Redundant Topology, page 11-3
- Failure Detection, page 11-3

This chapter presents the Failover and redundancy capabilities of the Cisco SCE 8000 platform. It first defines relevant terminology, as well as pertinent theoretical aspects of the redundancy and failover solution. It then explains specific recovery procedures for both single and dual link topologies. It also explains specific update procedures to be used in a cascaded Cisco SCE platform deployments. When failover is required in a deployment, a topology with two cascaded Cisco SCE 8000 platforms is used. This cascaded solution provides both network link failover, and failover of the functionality of the Cisco SCE platform, including updated subscriber state.

Terminology and Definitions

Following is a list of definitions of terms used in the chapter as they apply to the Cisco failover solution, which is based on cascaded Cisco SCE platforms.

- Failover—A situation in which the Cisco SCE platform experiences a problem that makes it impossible for it to provide its normal functionality, and a second Cisco SCE platform device immediately takes over for the failed Cisco SCE platform.
- Hot standby—When two Cisco SCE platforms are deployed in a failover topology, one Cisco SCE platform is active, while the second Cisco SCE platform is in standby, receiving from the active Cisco SCE platform all subscriber state updates and keep alive messages.
- Primary/Secondary—The terms Primary and Secondary refer to the default status of a particular Cisco SCE platform. The Primary Cisco SCE Platform is active by default, while the Secondary device is the default standby. Note that these defaults apply only when both devices are started together. However, if the primary Cisco SCE platform fails and then recovers, it will not revert to active status, but remains in standby status, while the secondary device remains active.
- Subscriber state failover—A failover solution in which subscriber state is saved.

Redundant Topologies

The Cisco SCE 8000 includes SPA Interface Processor card with an internal electrical bypass module, which provides the capability of preserving the network link in case of failure. However, preserving the Cisco SCE platform functionality in case of a failure, requires a redundant Cisco SCE platform. Cisco provides a unique solution for this scenario, through deploying two cascaded Cisco SCE platforms.

The cascading is implemented by connecting the two Cisco SCE platforms using two of the data links. In each Cisco SCE platform, two of the four data interfaces are connected to each of the network links, while the other two data interfaces are used for cascading between the Cisco SCE platforms. (See the *Cisco SCE8000 10GBE Installation and Configuration Guide* for specific cabling procedures for redundant topologies.) The cascade ports are used for transferring network traffic, keep-alive messages and subscriber state updates.

External Bypass

The Cisco SCE 8000 platform can control an external bypass device, which bypasses the traffic during a power failure and also under specific control command from the Cisco SCE 8000. The Cisco SCE 8000 automatically activates the external bypass device during reload for the short period (less than 10 seconds) in which the SPA Interface Processor card does not forward traffic between traffic ports. In addition, the Cisco SCE 8000 can be configured to activate the external bypass device in the following cases:

- After executing the **external-bypass** command , until the **no external-bypass** command is executed
- When the Cisco SCE 8000 is in failure state.

Note that in a cascaded configuration, an external bypass device should be connected only for the traffic ports. The cascade ports should be directly connected between the two Cisco SCE 8000 platforms (see Figure 11-1).

Hardware Bypass

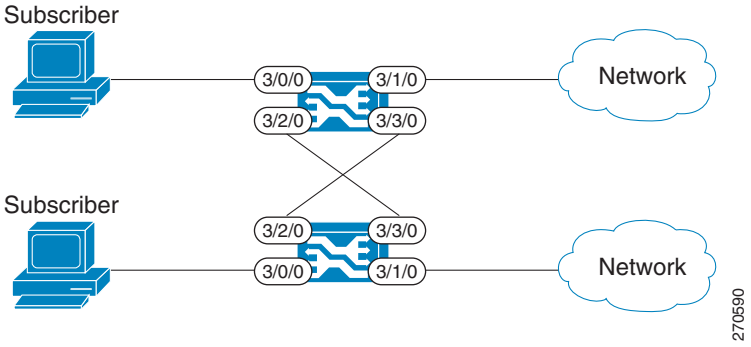
The Cisco SCE 8000 platform can support the hardware bypass, which bypasses the traffic of the configured static parties created in the hw-bypass mode at the hardware (SIP module) level based on their IP address or the IP address range.

Note The static parties created only in the hw-bypass mode are bypassed.

In-line Dual Link Redundant Topology

This topology serves inline deployments where the Cisco SCE platform functionality should be preserved in case of a failure, in addition to preserving the network link (Figure 11-1).

Figure 11-1 In-line Dual Link Redundant Topology



Failure Detection

The Cisco SCE 8000 platform has several types of mechanisms for detecting failures:

- Internal failure detection—The Cisco SCE platform monitors for hardware and software conditions such as overheating and fatal software errors.
- Inter-device failure detection—The Cisco SCE platform sends periodic keep-alive messages via the cascade ports
- The Cisco SCE platform-Subscriber Manager (SM) communication failure detection—A failure to communicate with the SM may be regarded as a cause for failover. However, this communication failure is not necessarily a problem in the Cisco SCE platform. If the connection to the SM of the active Cisco SCE platform has failed, while the connection to the SM of the standby Cisco SCE platform is alive, a failover process will be initiated to allow the Cisco SCE platform proper exchange of information between the Cisco SCE platforms and the SM.
- Link failure—The system monitors all three types of links for failures:
  - Traffic port link failure—Traffic cannot flow through the Cisco SCE platform.
  - Cascade port link failure—Traffic cannot flow between the Cisco SCE platforms through the cascade ports.
  - Management port link failure—This is not a failure that interrupts traffic on the link in and of itself. However, when SM is used, management port link failure will cause an SM connection failure and this, in turn, will be declared as a failure of the Cisco SCE platform.

  This type of failure, in most cases, does not require reboot of the Cisco SCE platform. When the connection with the SM is re-established the Cisco SCE platform is again ready for hot standby. If both Cisco SCE platforms lose their connections with the SM, it is assumed that it is the SM which has failed, thus, no action will be taken in the Cisco SCE platform.

Link Failure Reflection

The Cisco SCE platforms are transparent at Layers 2 and 3. The Cisco SCE 8000 platform operates in promiscuous mode, and the network elements on both sides of the Cisco SCE platform, are using the MAC address of the other network element when forwarding traffic.

To assist the network elements on both sides of the Cisco SCE platform to identify the link failures as quickly as possible, the Cisco SCE platform supports a functionality of reflecting to the other side of the Cisco SCE platforms events of link failure. When the link on one side of the Cisco SCE platform fails, the corresponding link on the other side is forced down, to reflect the failure. Link failure reflection is done on the traffic ports. When operating in deployments of single Cisco SCE platform with two data links, link failure is reflected between the two ports of each link.

When working with two cascaded Cisco SCE platforms, link failure is reflected in two cases:

- Reflection between the traffic ports of each Cisco SCE platform.
- If there is a failure in the cascade port link, the two Cisco SCE platforms can no longer support proper processing of the two links, since the traffic flowing on the standby Cisco SCE platform's link must be forwarded to the active Cisco SCE platform for processing. In this case the link failure is reflected from the cascade ports to the traffic ports of the standby Cisco SCE platform, in order to force the network to switch all the traffic only through the link of the active Cisco SCE platform.

Hot Standby and Failover

The failover solution requires two Cisco SCE platforms connected in a cascade manner.

- Hot Standby, page 11-6
- Failover, page 11-6
- Hardware Crash Mode, page 11-8
- Failure in the Cascade Connection, page 11-9
- Installing a Cascaded System, page 11-9

Hot Standby

In failover solution, one of the Cisco SCE platforms is used as the active Cisco SCE platform and the other is used as the standby. Although traffic enters both the active and the standby Cisco SCE platforms, all traffic processing takes place in the Cisco SCE platform which is currently the active one. The active Cisco SCE platform processes the traffic coming on both links, its own link and the link connected to the standby Cisco SCE platform, as follows

- All traffic entering the active Cisco SCE platform through its traffic ports is processed in that Cisco SCE platform and then forwarded to the line.
- All traffic entering the standby Cisco SCE platform through its traffic ports is forwarded through the cascade ports to the active Cisco SCE platform where it is processed, and then returned to the standby Cisco SCE platform through the cascade ports to be forwarded to the original line from which it came.

Since only one Cisco SCE platform processes all traffic at any given time, split flows, which are caused by asymmetrical routing, that exist in the two data links are handled correctly.

To support subscriber-state failover, both Cisco SCE platforms hold subscriber states for all parties, and subscriber state updates are exchanged between the active Cisco SCE platform and the standby. This way, if the active Cisco SCE platform fails, the standby Cisco SCE platform is able to start serving the line immediately with a minimum loss of subscriber-state.

The two Cisco SCE platforms also use the cascade channel for exchanging periodic keep-alive messages.

Failover

In failover solution, the two Cisco SCE platforms exchange keep alive messages via the cascade ports. This keep alive mechanism enables fast detection of failures between the Cisco SCE platforms and fast failover to the standby Cisco SCE platform when required.

If the active Cisco SCE platform fails, the standby Cisco SCE platform then assumes the role of the active Cisco SCE platform.

The failed Cisco SCE platform uses its electrical bypass mechanism, which is a hardware entity that is separate from the main board and processors, to forward traffic to the other Cisco SCE platform, and to forward processed traffic back to the link. The previously standby Cisco SCE platform now processes all the traffic of this other link that is forwarded to it by the previously active Cisco SCE platform in addition to the traffic of its own link.

When the failed Cisco SCE platform recovers, it will remain in standby, while the previously standby Cisco SCE platform remains active. Switching the Cisco SCE platforms back to their original roles may be performed manually, if required, after the failed Cisco SCE platform has either recovered or been replaced.

If the failure is in the standby Cisco SCE platform, it will continue to forward traffic to the active Cisco SCE platform and back to its link, while the active Cisco SCE platform continues to provide its normal processing functionality to the traffic of the two links.

Note For information regarding the synchronization of subscriber information between cascaded Cisco SCE platforms and the effect of failover on the subscriber databases, see "Anonymous Groups and Subscriber Templates" section on page 10-7.

There are three user-configurable options that are relevant in a situation when a Cisco SCE platform fails:

- **Bypass—**Maintain the link in bypass mode (continue sending traffic to the other Cisco SCE platform, and then continue forwarding the processed traffic back to the link). The incoming traffic in the failed Cisco SCE platform is forwarded to the working Cisco SCE platform, where it is processed and then sent back to the original Cisco SCE platform and back to the link. This is the default configuration.
  - Effect on the network link—Negligible.
  - Effect on the Cisco SCE platform functionality—The effect on the Cisco SCE platform functionality is dependent on the failed Cisco SCE platform.
  - If the failure is in the standby Cisco SCE platform—The active Cisco SCE platform continues providing its normal functionality, processing the traffic of the two links.
  - If the failure is in the active Cisco SCE platform—The standby Cisco SCE platform takes over processing the traffic, and becomes the active Cisco SCE platform.
- **Cutoff—**Change the link of the failed Cisco SCE platform to cutoff (layer 1) forcing the network to switch all traffic through the line of the working Cisco SCE platform. This will, of course, decrease the network capacity by 50%, but may be useful in some cases.

  This option is available for use in special cases, and requires special configuration.
  - Effect on the network link—The network loses 50% of its capacity (until the failed Cisco SCE platform has recovered).
  - Effect on the Cisco SCE platform functionality—The effect on the Cisco SCE platform functionality is dependent on the failed Cisco SCE platform.
  - If the failure is in the standby Cisco SCE platform—The active Cisco SCE platform continues providing its normal functionality, processing the traffic of the two links.
  - If the failure is in the active Cisco SCE platform—The standby Cisco SCE platform takes over processing the traffic, and becomes the active Cisco SCE platform.
- **External-bypass—**Activate the external bypass device connected to the failed Cisco SCE platform, passing all traffic through the line without being serviced by the failed Cisco SCE platform. Although this may cause the traffic passing through the other link (that of the non-failed Cisco SCE platform) to get service in split-flow conditions if asymmetric routing is present, it may be useful in some cases.

  This option is available for use in special cases, and requires special configuration.
  - Effect on the network link—Negligible

  - Effect on the Cisco SCE platform functionality—Since the active Cisco SCE platform only sees the traffic of a single link, split-flow effects might occur. The link connected to the failed Cisco SCE platform gets no service.

Hardware Crash Mode

There are three hardware components that operate together to produce the desired behavior upon failure of the Cisco SCE 8000 platform:

- external optical bypass: In a cascade setup, if the traffic links are connected to external optical bypass modules, the optical bypass may be either activated or deactivated by the hardware during a failure.

  The external optical bypasses protect against a second Cisco SCE 8000 platform failure. In the case of a second failure, if a bypass module is connected to the last Cisco SCE 8000 to fail, it will be enabled. This preserves one of the network links, assuming the *on-failure* configuration is is **bypass** . If the *on-failure* configuration is **external-bypass** , the external optical bypass is activated even during a single failure by the failed Cisco SCE platform.

- internal electrical bypass: The Cisco SCE 8000 contains internal electrical bypasses connecting SPA modules 0 and 2 and modules 1 and 3. These bypasses transmit the traffic to and from the cascade connections between the platforms.

- SPA modules: Under some conditions (such as if the *on-failure* configuration is **cutoff**), the SPA modules are disabled when failure occurs.

The collective behavior of these three components is known as the hardware crash mode and is dependent on the configuration of the *on-failure* parameter of the **connection-mode** command, as well as whether the platform is the active or standby platform.

In the standby platform, hardware crash mode behavior is as follows for:

For **on-failure bypass**:

- The external optical bypass, if installed, is deactivated (traffic is sent to the platform).
- The electrical bypass is enabled (cascade ports transmit traffic to the active platform for processing)
- The SPA modules are enabled (all ports and links are functioning)

This means that whether the standby platform is operational or has failed, it transmits traffic to the active platform for processing via the electrical bypasses.

For **on-failure cutoff**:

- The external optical bypass, if installed, is deactivated (traffic is sent to the platform).
- The electrical bypass is disabled.
- The SPA modules are powered off.

This means that if the standby platform fails, the link connected to it Cisco SCE is severed.

For **on-failure external-bypass**:

- The external optical bypass is activated (traffic is bypassed).
- The electrical bypass is enabled between ports 0 and 1. (This is a special configuration, done just in case the optical bypass device does not function, which has a very low probability.)
- The SPA modules are enabled (all ports and links are functioning)

This means that when the standby Cisco SCE platform is failed, the external optical bypass is used to ensure link continuity at the expense of not servicing the traffic on that link.

In the active platform, the hardware crash mode behavior is exactly the same as on a standby platform. The active platform assumes that when it fails, the standby platform will take over and process the traffic.

If the standby platform has failed, failure of the active platform means that the entire system has failed (this state is also called a 'second failure'). The hardware crash mode behavior in the active platform in this case depends on the configuration of the *on-failure* parameter, which determines whether traffic is bypassed via the external bypass, if installed, maintaining traffic flow through one link, although with no processing, or whether traffic is completely cut off.

If the standby platform has failed and the *on-failure* configuration is **bypass** or **external-bypass**, hardware crash mode behavior in the active platform is as follows:

- The external optical bypass is activated (traffic is bypassed via the external bypass).
- The electrical bypass is enabled between ports 0 and 1. (This is a special configuration, done just in case the optical bypass device does not function, which has a very low probability)
- The SPA modules are enabled (all ports and links are functioning)

Since neither platform is operational at this point, there is no processing taking place and traffic is simply bypassed via the external optical bypass. If external optical bypass modules are not installed, traffic is cut off.

If the standby platform has failed and the *on-failure* configuration is **cutoff**, hardware crash mode behavior in the active platform is as follows:

- The external optical bypass, if installed, is deactivated (traffic is sent to the platform).
- The electrical bypass is disabled (cascade ports do not transmit traffic to the standby platform since it is also not operational)
- The SPA modules are disabled (ports and links are not functioning)

Since neither platform is operational at this point, there is no processing taking place. In addition, although traffic reaches the platform, since the internal bypasses as well as all the ports are disabled, this results in a 'dead end' , cutting off all traffic on both links.

Failure in the Cascade Connection

The effect of a failure in the cascade connection between the two Cisco SCE platforms depends on whether one or both connections fail:

- Only one cascade connection is down—In this case, both Cisco SCE platforms can still communicate, so each still knows the status of the peer.

  As long as one cascade connection remains up, the standby will cut off its traffic links so that all traffic is routed via the active Cisco SCE platform. Therefore, split flow is avoided, but at the expense of half line capacity.
- Both cascade links are down—In this case, neither Cisco SCE platform knows anything about the status of the peer. Each platform then works in standalone mode, which means that each Cisco SCE platform processes on its own traffic only. This results in split flows.

Installing a Cascaded System

This section outlines the installation procedures for a redundant solution with two cascaded Cisco SCE platforms.

For information on topologies and connections, see the *Cisco SCE8000 10GBE Installation and Configuration Guide* .

For details of the CLI commands, see the *Cisco SCE8000 CLI Command Reference*.

Note When working with two Cisco SCE platforms with split-flow and redundancy, it is extremely important to follow this installation procedure.

Step 1 Install both Cisco SCE platforms, power them up, and perform the initial system configuration.

Step 2 If external optical bypass modules are installed, make sure they are connected correctly and are operational. Use the **show external-bypass** command.

Step 3 Connect both Cisco SCE platforms to the management station.

Step 4 Connect the cascade ports.

The cascade ports may be connected either directly in Layer 1 (dark fibers) or through a switch. When connecting the cascade ports through a switch, it is important to assign each cascade link to a different VLAN, otherwise the traffic will be forwarded incorrectly (between different links) by the switch.

Step 5 Set topology configurations for each Cisco SCE platform via the connection-mode options. (See "Topology-Related Parameters for Redundant Topologies" section on page 11-13)

Step 6 Make sure that the Cisco SCE platforms have synchronized and active Cisco SCE platform was selected. Use the **show interface linecard 0 connection-mode** command.

Step 7 If you want to start in bypass mode, change the link mode to bypass in both Cisco SCE platforms. The bypass mode will be applied only to the active Cisco SCE platform. (See "About the Link Mode" section on page 8-6.)

Step 8 Verify the link mode configuration. (See "Monitoring the System" section on page 11-14.) Use the **show interface linecard 0 link mode** command.

Step 9 Connect the traffic port of Cisco SCE platform #0. This will cause a momentary down time until the network elements from both sides of the Cisco SCE platform auto-negotiate with it and start working (when working inline).

Step 10 Connect the traffic port of Cisco SCE platform #1, this will cause a momentary down time until the network elements from both sides of the Cisco SCE platform auto-negotiate with it and start working (when working inline).

Step 11 When full control is needed, change the link mode on both Cisco SCE platforms on both links to 'forwarding'. It is recommended to first configure the active Cisco SCE platform and then the standby. (See "About the Link Mode" section on page 8-6.)

Step 12 You can now start working with the Subscriber Manager.

Note Cisco SCE devices does not support downgrading from Cisco SCE Release 3.8.0 and later to Cisco SCE Releases earlier than Release 3.8.0 when the devices are configured in a cascade setup. To downgrade the devices to Cisco SCE Releases earlier than Release 3.8.0, remove the devices from the cascade setup and configure the devices as standalone devices, and then start the downgrade procedure. Note that this downgrade procedure may cause a service disruption.

Recovery

- Replacing the Cisco SCE Platform (Manual Recovery), page 11-11
- Reboot Only (Fully Automatic Recovery), page 11-12

This section specifies the procedure for recovery after a failure. The purpose of the recovery procedure is to restore the system to fully functional status. After the recovery procedure, the behavior of the system is the same as after installation.

A failed Cisco SCE platform may either recover automatically or be replaced (manual recovery). Whether recovery is automatic or manual depends on the original cause of the failure:

- Power failure—Manual or automatic recovery can be implemented.
- Any failure resulting in a reboot—Manual or automatic recovery can be implemented (this is configurable).
- 3-consecutive reboots within half an hour—Manual recovery only
- Cascade ports link-failure—Automatic recovery when link revives.
- Traffic link failure—Automatic recovery when link revives.
- Failure in the communications with the SM—Automatic by SM decisions after connection is re-established.
- Hardware malfunction—Manual recovery, after replacing the malfunctioning Cisco SCE platform.

Replacing the Cisco SCE Platform (Manual Recovery)

This is done in two stages, first manual installation steps performed by the technician, and then automatic configuration steps performed by the system.

- Manual Steps, page 11-11
- Automatic Steps (in parallel with the manual steps, requires no user intervention):, page 11-12

Manual Steps

Step 1 Disconnect the failed Cisco SCE platform from the network

Step 2 Connect a new Cisco SCE platform to the management link and the cascade links (leave network ports disconnected.)

Step 3 Configure the Cisco SCE platform.

Step 4 Basic network configurations done manually (first time).

Step 5 Load application software (*Service Control Application for Broadband*) to the Cisco SCE platform. Provide application configuration.

Step 6 Connect the traffic ports to the network links.

Automatic Steps (in parallel with the manual steps, requires no user intervention):

Step 1 Establishment of inter-Cisco SCE platform communication.

Step 2 Synchronization with the SM.

Step 3 Copying updated subscriber states from the active Cisco SCE platform to the standby.

Reboot Only (Fully Automatic Recovery)

Step 1 Reboot of the Cisco SCE platform.

Step 2 Basic network configurations.

Step 3 Establishment of inter-Cisco SCE platform communication.

Step 4 Selection of the active Cisco SCE platform.

Step 5 Synchronization of the recovered Cisco SCE platform with the SM.

Step 6 Copying updated subscriber states from the active Cisco SCE platform to the standby.

CLI Commands for Cascaded Systems

- Topology-Related Parameters for Redundant Topologies, page 11-13
- Configuring the Connection Mode, page 11-13
- Monitoring the System, page 11-14

This section presents CLI commands relevant to the configuration and monitoring of a redundant system.

Topology-Related Parameters for Redundant Topologies

All four of the topology-related parameters are required when configuring a redundant topology.

- **Connection mode—**Redundancy is achieved by cascading two Cisco SCE platforms. Therefore the connection mode for both Cisco SCE platforms may be either:
  - Inline-cascade
  - Receive-only-cascade
- **sce-id—**For each of the cascaded Cisco SCE platforms, this parameter defines the number of the link (0 or 1) connected to this Cisco SCE platform.

  The sce-id parameter, which identifies the Cisco SCE platform, replaces the physically-connected-link parameter, which identified the link. This change was required with the introduction of the Cisco SCE 8000 GBE platform, which supports multiple links. In the Cisco SCE 8000 10GBE, the number assigned to the sce-id parameter (0 or 1) will be defined as the of number of the physically-connected-link.

Note For backwards compatibility, the physically-connected-link parameter is currently still recognized.

- **Priority—**For each of the cascaded Cisco SCE platforms, this parameter defines whether it is the primary or secondary device.
- **On-failure—**For each of the cascaded Cisco SCE platforms, this parameter determines whether the system cuts the traffic or bypasses it when the Cisco SCE platform either has failed or is booting.

Configuring the Connection Mode

Use the following command to configure the connection mode, including the following parameters.

- inline/receive only
- sce-id (physically connected link)
- behavior upon failure of the Cisco SCE platform
- primary/secondary

To configure the connection mode, use the following command.

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

Command	Purpose
connection-mode (inline-cascade|receive-only-cascade) sce-id {0|1}priority {primary|secondary} on-failure {bypass | external-bypass|cutoff}	Configures the connection mode.

Examples

EXAMPLE 1

Use the following command to configure the primary Cisco SCE platform in a two-Cisco SCE platform inline topology. Link 1 is connected to this Cisco SCE platform and the behavior of the Cisco SCE platform if a failure occurs is *bypass* , which is the default.

```
SCE(config-if)#connection-mode inline-cascade sce-id 1 priority primary
```

EXAMPLE 2

Use the following command to configure the Cisco SCE platform that might be cascaded with the Cisco SCE platform in Example 1. This Cisco SCE platform would have to be the secondary Cisco SCE platform, and Link 0 would be connected to this Cisco SCE platform, since Link 1 was connected to the primary. The connection mode would be the same as the first, and the behavior of the Cisco SCE platform if a failure occurs is external-bypass.

```
SCE(config-if)# connection-mode inline-cascade sce-id 0 priority secondary on-failure
external-bypass
```

Monitoring the System

Use the following commands to view the current connection mode and link mode parameters.

How to View the Current Connection Mode

From the SCE# prompt, type:

Command	Purpose
show interface linecard 0 connection-mode	Displays the current connection mode.

Monitoring the Connection Mode: Examples

The following example shows the current configuration of the connection mode for a single platform.

```
SCE>enable 5
Password:<cisco>
SCE>show interface linecard 0 connection-mode
Slot 0 connection mode
Connection mode is inline
slot failure mode is external-bypass
Redundancy status is standalone
SCE>
```

The following example shows the current configuration of the connection mode for a cascaded system.

```
SCE>enable 5
Password:<cisco>
SCE>show interface linecard 0 connection-mode
Slot 0 connection mode
Connection mode is inline-cascade
slot 0 sce-id is 1
slot 0 is secondary
slot 0 is connected to peer
slot failure mode is bypass
Redundancy status is active
SCE>
```

How to View the Cisco SCE-ID

From the SCE# prompt, type:

Command	Purpose
show interface linecard 0 sce-id	Displays the Cisco SCE-ID.

Viewing the Cisco SCE-ID: Example

```
SCE>enable 5
Password:<cisco>
SCE>show interface linecard 0 sce-id
slot 0 sce-id is 1
```

How to View the Current Redundancy Status of the Cisco SCE Platform

From the SCE# prompt, type:

Command	Purpose
show interface linecard 0 cascade **redundancy-status**	Displays the current redundancy status of the Cisco SCE platform.

Viewing the Current Redundancy Status of the Cisco SCE Platform: Example

The following example shows typical output of this command.

SCE>enable 5
Password:<cisco>
SCE>**show interface linecard 0 cascade redundancy-status**
Redundancy status is active

How to View Information about the Peer Cisco SCE Platform

From the SCE# prompt, type:

Command	Purpose
show interface linecard 0 cascade **peer-sce-information**	Displays information about the peer Cisco SCE platform.

Viewing Information about the Peer Cisco SCE Platform: Example

The following example shows typical output of this command.

```
SCE>enable 5
Password:<cisco>
SCE>show interface linecard 0 cascade peer-sce-information
Peer SCE's IP address is 10.10.10.10
```

How to View Information about the Cascade Connections

From the SCE> prompt, type:

Command	Purpose
show interface linecard 0 cascade **connection-status**	Displays information about the cascade connections.

Monitoring the Connection Status: Examples

The following example shows the output of this command in the case of two cascaded Cisco SCE 8000 GBE platforms where the cascade interfaces have not been connected correctly.

```
SCE>enable 5
Password:<cisco>
SCE>show interface linecard 0 cascade connection-status
SCE is improperly connected to peer SCE
Please verify that each cascade port is connected to the correct port of the peer SCE.
Note that in the current topology, the SCE must be connected to its peer as follows:
Port 3/2/0 must be connected to port 3/2/0 at peer
Port 3/3/0 must be connected to port 3/3/0 at peer
SCE>
```

The following example shows the output of this command in the case of two cascaded Cisco SCE platforms where the cascade interfaces have been connected correctly.

```
SCE>enable 5
Password:<cisco>
SCE>show interface linecard 0 cascade connection-status
SCE is connected to peer SCE
SCE>
```

How to View the Current Link to Port Mappings

From the SCE> prompt, type:

Command	Purpose
show interface linecard 0 link-to-port-mapping	Displays the current link to port mappings.

Viewing the Current Link to Port Mappings: Example

The following example shows the link-to-port mapping.

```
SCE>enable 5
Password:<cisco>
SCE>show interface linecard 0 link-to-port-mappings
Link Id | Upstream Port <Out>| Downstream Port <Out>
------------------------------------------------------------
0       | 0/2                | 0/1
SCE>
```

How to View the Current Link Mode

From the SCE# prompt, type:

Command	Purpose
show interface linecard 0 link mode	Displays the current link mode.

Configuring Forced Failure

Use the following commands to force a virtual failure condition, and to exit from the failure condition when performing an application upgrade.

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

Commands	Purpose
force failure-condition	Forces the Cisco SCE platform into a virtual failure state.
no force failure-condition	Exits from the virtual failure state.

System Upgrades

- Firmware Upgrade (package installation), page 11-19
- Application Upgrade, page 11-19
- Simultaneous Upgrade of Firmware and Application, page 11-20

In a redundant solution, it is important that firmware and/or application upgrades be performed in such a way that line and service are preserved.

Refer to the following sections for instructions on how to perform these procedures on two cascaded Cisco SCE platforms:

- Upgrade the firmware only
- Upgrade the application only
- Upgrade both the firmware and the application at the same time

Note When upgrading only one component (either firmware only or application only), always verify that the upgraded component is compatible with the component that was not upgraded.

Firmware Upgrade (package installation)

Step 1 Install package on both Cisco SCE platforms (open the package and copy configuration).

Step 2 Reload the standby Cisco SCE platform.

Step 3 Wait until the standby finishes synchronizing and is ready to work.

Step 4 Make sure that the connection mode configurations are correct.

Step 5 Reload the active Cisco SCE platform.

Step 6 After the former active Cisco SCE platform reboots and is ready to work manually, it may be left as standby or we can manually switch the Cisco SCE platforms to their original state.

Application Upgrade

Step 1 Unload the application in the standby Cisco SCE platform.

Step 2 Load new application to the standby Cisco SCE platform.

Step 3 Make sure that the connection mode configurations are correct.

Step 4 Wait until the standby Cisco SCE platform finishes synchronizing and is ready to work.

Step 5 Force failure condition in the active Cisco SCE platform.

Step 6 Upgrade the application in the former active Cisco SCE platform.

Step 7 Remove the force failure condition in that platform.

Step 8 After the former active Cisco SCE platform recovers and is ready to work, it may remain the standby or can be manually switched back to active.

Simultaneous Upgrade of Firmware and Application

Step 1 In the standby Cisco SCE platform:

- **a.** Uninstall the application.
- **b.** Upgrade the firmware (this includes a reboot).
- **c.** Install the new application.

Step 2 Force-failure in the active Cisco SCE platform.

This makes the updated Cisco SCE platform the active one, and it begins to give the NEW service.

Step 3 Repeat step 1 for the (now) standby Cisco SCE platform.

Since this includes a reboot, it is not necessary to undo the force failure command.



Identifying and Preventing Distributed Denial-of-Service Attacks

Revised: September 23, 2015

Introduction

This chapter describes the ability of the Cisco SCE platform to identify and prevent DDoS attacks, and the various procedures for configuring and monitoring the Attack Filter Module.

- Attack Filtering and Attack Detection, page 12-2
- Configuring Attack Detectors, page 12-8
- Subscriber Notifications, page 12-20
- Preventing and Forcing Attack Detection, page 12-21
- Monitoring Attack Filtering, page 12-24

Attack Filtering and Attack Detection

- Attack Filtering, page 12-2
- Specific Attack Filtering, page 12-2
- Attack Detection, page 12-4
- Attack Detection Thresholds, page 12-4
- Attack Handling, page 12-5
- Hardware Filtering, page 12-6

Attack Filtering

The Cisco SCE platform includes extensive capabilities for identifying DDoS attacks, and protecting against them.

Attack filtering is performed using specific-IP attack detectors. A specific-IP attack detector tracks the rate of flows (total open and total suspected) in the Cisco SCE platform for each combination of IP address (or pair of IP addresses), protocol (TCP/UDP/ICMP/Other), destination port (for TCP/UDP), interface and direction. When the rates satisfy user-configured criteria, it is considered an attack, and a configured action can take place (report/block, notify subscriber, send SNMP trap).

This mechanism is enabled by default, and can be disabled and enabled for each attack type independently.

There are 32 different attack types:

- **1—**TCP flows from a specific IP address on the subscriber side, regardless of destination port
- **2—**TCP flows to a specific IP address on the subscriber side, regardless of destination port
- **3-4—**Same as 1 and 2, but for the opposite direction (subscriber network)
- **5—**TCP flows from a specific IP address on the subscriber side to a specific IP address on the network side
- **6—**Same as 5, but for the opposite direction (from the network side to the subscriber side)
- **7-12—**Same as 1-6 but with a specific destination port common to all flows of the attack (1-6 are port-less attack types, 7-12 are port-based attack types)
- **13-24—**Same as 1-12 but for UDP instead of TCP
- **25-28—**Same as 1-4 but for ICMP instead of TCP
- **29-32—**Same as 1-4 but for Other protocols instead of TCP

Specific Attack Filtering

When the specific IP attack filter is enabled for a certain attack type, two rates are measured per defined entity:

- Rate of new flows
- Rate of suspected flows (In general, suspected flows are flows for which the SCOS did not see proper establishment (TCP) or saw only a single packet (all other protocols)).

Separate rate meters are maintained both for each IP address separately (single side) and for IP address pairs (the source and destination of a given flow), so when a specific IP is attacking a specific IP, this pair of IP addresses defines a single incident (dual-sided).

Based on these two metrics, a specific-IP attack is declared if either of the following conditions is present:

- The new flows rate exceeds a certain threshold
- The suspected flows rate exceeds a configured threshold and the ratio of suspected flows rate to total new flow rate exceeds a configured threshold.

When the rates stop satisfying this criterion, the end of that attack is declared.

Note that specific attack filtering is configured in two steps:

- Enabling specific IP filtering for the particular attack type.
- Configuring an attack detector for the relevant attack type. Each attack detector specifies the thresholds that define an attack and the action to be taken when an attack is detected.

  In addition to specific attack detectors, a default detector exists that can be configured with user-defined thresholds and action, or system defaults may be retained.

In addition, the user can manually override the configured attack detectors to either force or prevent attack filtering in a particular situation.

Specific IP filtering for selected attack types is enabled with the following parameters. These parameters control which of the 32 attack types are being filtered for:

- **Protocol—**TCP, UDP, ICMP, or Other
- **Attack direction—**The direction of the attack may be identified by only one IP address or by two IP addresses:
  - **single side—**The attack is identified by either the source IP address or the destination address only.

  The filter definition may specify the specific side, or may include any single side attack, regardless of side (both).
  - **dual side** (TCP and UDP protocols only)—The attack is identified by both the source and destination IP addresses. In other words, when a specific IP attacks a specific IP, this is detected as one incident rather than as two separate incidents.
- **Destination port** (TCP and UDP protocols only)—Defines whether specific IP detection is enabled or disabled for port-based or port-less detections. Enable port-based detection for TCP/UDP attacks that have a fixed destination port or ports.

  The list of destination ports for port-based detection is configured separately. (See "Specific Attack Detectors" section on page 12-14.)

Attack Detection

Specific IP detections are identified with the following parameters:

- Specific IP address (or two IP addresses for dual-sided detections)
- Protocol—TCP, UDP, ICMP or Other
- Port—For TCP/UDP attacks that have a fixed destination port
- Side—Interface (Subscriber/Network) from which attack packets are sent
- Attack-direction—If a single IP address is specified, the IP address is an attack-source or an attack-destination address.

The system can identify a maximum of 1000 independent, simultaneous attacks.

Once an attack is identified, the system can be instructed to perform any of the following actions:

- **Report—**By default, the attack beginning and end are always reported.
- **Block—**The system will block all attack traffic for the duration of the attack. (The traffic is from or to the attack IP address, depending on whether the IP address is an attack-source or attack-destination)
- **Notify—**Subscriber notification. When the IP address identified is mapped to a particular subscriber context, the system can be configured to notify the subscriber of the fact that he is under an attack (or a machine in his network is generating such an attack), using HTTP Redirect.
- **Alarm—**The system will generate an SNMP trap each time an attack starts and stops.

Attack detection and handling are user-configurable. The remainder of this chapter explains how to configure and monitor attack detection.

Attack Detection Thresholds

There are three thresholds that are used to define an attack. These thresholds are based on meters that are maintained by the Cisco SCE platform for each IP address or pair of addresses, protocol, interface and attack-direction.

- **open flow rate—**A flow for which some traffic was seen. Any packet seen for a new flow is enough to declare this flow an open flow.

  The rate is measured in new flows per second.
- **suspected flow rate—**A suspected flow is one that was opened, but did not become an established flow.

  The rate is measured in new flows per second.
- **suspected flow ratio—**The ratio of the suspected flow rate to the open flow rate.

As explained above, a specific-IP attack is declared if either of the following conditions is present:

- The open flows rate exceeds the threshold
- The suspected flows rate exceeds the threshold and the suspected flows ratio exceeds the threshold.

The values for each attack type will have a separate configured default value.

In general, for a given protocol, the suspected flows rate threshold should be lower for a port-based detection than for a port-less detection. This is because flows with a given IP address and a common destination port are metered twice:

- By themselves—To detect a port-based attack
- Together with flows with the same IP address and different destination ports—to detect a port-less attack

If a port-based attack occurs, and the rate of flows is above both thresholds (port-based thresholds and the port-less thresholds), it is desirable for the port-based attack to be detected before the port-less attack. Similarly, this threshold should be lower for dual-IP detections then for single-IP detections.

The user may define values for these thresholds that override the preset defaults. It is also possible to configure specific thresholds for certain IP addresses and ports (using access lists and port lists). This enables the user to set different detection criteria for different types of network entities, such as a server farm, DNS server, or large enterprise customer.

Attack Handling

Attack handling can be configured as follows

- **Configuring the action:**
  - Report—Attack packets are processed as usual, and the occurrence of the attack is reported.
  - Block—Attack packets are dropped by the Cisco SCE platform, and therefore do not reach their destination.

  Regardless of which action is configured, two reports are generated for every attack: one when the start of an attack is detected, and one when the end of an attack is detected.
- **Configuring subscriber-notification (notify):**
  - Enabled—If the subscriber IP address is detected to be attacked or attacking, the subscriber is notified about the attack.
  - Disabled—The subscriber is not notified about the attack.
- **Configuring sending an SNMP trap (alarm):**
- Enabled—An SNMP trap is sent when attack begins and ends.

  The SNMP trap contains the following information fields:
  - A specific IP address or
  - **Protocol** (TCP, UDP, ICMP or Other)
  - **Interface** (User/Network) behind which the detected IP address is found. This is referred to below as the attack ‘side’
  - **Attack direction** (whether the IP address is the attack source or the attack destination).
  - **Type of threshold breached** (open- flows / ddos- suspected- flows) [‘attack- start’ traps only]
  - **Threshold value breached** [‘attack- start’ traps only]
  - **Action taken** (report, block) indicating what was the action taken by the Cisco SCE platform in response to the detection
  - **Amount of attack flows blocked/ reported** providing the total number of flows detected during the attack [‘attack- stop’ traps only]
- Disabled—No SNMP trap is sent

Subscriber Notification

When an attack is identified, if the IP address is detected on the subscriber side and is mapped to a subscriber, the system notifies the application about the attack. This enables the application to notify the subscriber about the attack on-line by redirecting HTTP requests of this subscriber to a server that will notify it of the attack.

In addition, when blocking TCP traffic, the system can be configured not to block a specified port to make this redirection possible. This port is then considered to be un-blockable.

Note that subscriber-notification can only function if supported by the Service Control Application currently loaded to the Cisco SCE platform, and the application is configured to activate this capability. To verify whether the application you are using supports attack subscriber notification, and for details about enabling attack subscriber notification in the application, please refer to the documentation of the relevant Service Control Application.

Hardware Filtering

The Cisco SCE platform has two ways of handling an attack: by software or by hardware. Normally, attacks are handled by software. This enables the Cisco SCE platform to accurately measure the attack flows and to instantly detect that an attack has ended.

However, very strong attacks cannot be handled successfully by the software. If the software cannot adequately handle an attack, the resulting high CPU load will harm the service provided by the Cisco SCE platform (normal traffic classification and control). An attack that threatens to overwhelm the software will, therefore, be automatically filtered by the hardware.

When the hardware is used to filter the attack, the software has no knowledge of the attack packets, and therefore the following side effects occur:

- The number of attack flows is greatly under-estimated by the software. This means that the total amount of flows in the attack reported by the CLI (show interface linecard attack-filter current-attacks) is considerably lower than the actual amount.
- Similarly, the reported attack flow rate (also reported by the CLI) is also considerably lower than the actual rate. Usually a rate of 0 is measured by the software.
- There is considerable delay in detecting the end of the attack. The delay in detecting the end of attack is limited by two upper bounds.
  - The first upper bound depends on the configured action, as follows:
  - Report—A delay of no more than 8 minutes
  - Block—A delay of no more than 64 minutes
  - A second upper bound for the delay is one minute more than actual duration of the attack (for example, maximum delay for detecting the end of an attack lasting three minutes is four minutes).
- The following examples illustrate the interaction of these two upper bounds:
  - For an attack lasting two minutes, the maximum delay in detecting the end will be three minutes, regardless of configured action
  - For an attack lasting two hours whose configured action is 'report', the maximum delay in detecting the end will be eight minutes
  - For an attack lasting two hours whose configured action is 'block', the maximum delay in detecting the end will be 64 minutes

Hardware attack filtering is an automatic process and is not user-configurable. However, due to the effects of hardware attack filtering on attack reporting, it is important to be aware of when hardware processing has been activated, and so monitoring of hardware filtering is essential. There are two ways to do this (see "Monitoring Attack Filtering" section on page 12-24):

- Check the " *HW-filter* " field in the **show interface linecard attack-filter current-attacks** command.
- Check the " *HW-filter* " field in the attack log file.

Configuring Attack Detectors

- Enabling Specific-IP Detection, page 12-10
- Configuring the Default Attack Detector, page 12-11
- Specific Attack Detectors, page 12-14
- Sample Attack Detector Configuration, page 12-18

The Cisco attack detection mechanism is controlled by defining and configuring special entities called Attack Detectors.

There is one attack detector called ‘default’, which is always enabled, and 99 attack detectors (numbered 1-99), which are disabled by default. Each detector (both the default and detectors 1-99) can be configured with a separate action and threshold values for all 32 possible attack types.

When detectors 1-99 are disabled, the default attack detector configuration determines the thresholds used for detecting an attack, and the action taken by the Cisco SCE platform when an attack is detected. For each attack type, a different set of thresholds and action can be set. In addition, subscriber-notification and SNMP traps (alarm) can be enabled or disabled in the same granularity.

The default attack detector should be configured with values that reflect the desired Cisco SCE platform behavior for the majority of the traffic flows flowing through it. However, it is not feasible to use the same set of values for all the traffic that traverses through the Cisco SCE platform, since there might be some network entities for which the characteristics of their normal traffic should be considered as an attack when coming from most other network elements. Here are two common examples:

- A DNS server is expected to be the target of many short DNS queries. These queries are typically UDP flows, each flow consisting of two packets: The request and the response. Normally, the Cisco SCE platform considers all UDP flows that are opened to the DNS server as DDoS-suspected flows, since these flows include less than 3 packets. A DNS server might serve hundreds of DNS requests per second at peak times, and so the system should be configured with a suitable threshold for DDoS-suspected flows for protocol = UDP and direction = attack-destination. A threshold value of 1000 flows/second would probably be suitable for the DNS server. However, this threshold would be unsuitable for almost all other network elements, since, for them, being the destination of such large rate of UDP flows would be considered an attack. Therefore setting a threshold of 1000 for all traffic is not a good solution.
- The subscriber side of the Cisco SCE platform might contain many residential subscribers, each having several computers connected through an Internet connection, and each computer having a different IP address. In addition, there might be a few business subscribers, each using a NAT that hides hundreds of computers behind a single IP address. Clearly, the traffic seen for an IP address of a business subscriber contains significantly more flows than the traffic of an IP address belonging to a residential subscriber. The same threshold cannot be adequate in both cases.

To let the Cisco SCE platform treat such special cases differently, the user can configure non-default attack detectors in the range of 1-99. Like the default attack detector, non-default attack detectors can be configured with different sets of values of action and thresholds for every attack type. However, to be effective, a non-default attack detector must be enabled and must be assigned an ACL (access control list). The action and thresholds configured for such attack detector are effective only for IP addresses permitted by the ACL. Non-default attack-detectors can be assigned a label for describing their purpose, such as ‘DNS servers’ or ‘Server farm’.

Non-default attack detectors are effective only for attack types that have been specifically configured. This eliminates the need to duplicate the default attack detector configuration into the configuration non-default attack detectors, and is best illustrated with an example: Suppose an HTTP server on the subscriber side of the Cisco SCE platform is getting many requests, which requires the use of a non-default attack detector for configuring high threshold values for incoming TCP flow rates. Assume

attack detector number 4 is used for this purpose; hence it is enabled, and assigned an ACL which permits the IP address of the HTTP server. Also suppose that it is desirable to protect subscribers from UDP attacks, hence the default attack detector is configured to block UDP attacks coming from the network (The default configuration is only to report attacks, not block them). If the HTTP server is attacked by a UDP attack from the network, the configuration of the default attack detector will hold for this HTTP server as well, since attack detector number 4 was not configured for UDP attacks.

For each of the non-default attack detectors, for each of the 32 attack types, there are four configurable settings:

- Threshold
- Action
- Subscriber-notification
- Alarm

Each of these four settings can be either configured (with a value or set of values) or not configured. The default state is for all them is not configured.

For each attack type, the set of enabled attack detectors, together with the default attack detector, forms a database used to determine the threshold and action to take when an attack is detected. When the platform detects a possible attack, it uses the following algorithm to determine the thresholds for attack detection.

- Enabled attack detectors are scanned from low to high numbers.
- If the IP address is permitted by the ACL specified by the attack detector, and a threshold is configured for this attack type, then the threshold values specified by this attack detector are used. If not, the scan continues to the next attack detector.
- If no attack detector matches the IP address/protocol combination, then the values of the default attack detector are used.

The same logic is applied when determining the values to use for the remaining settings: action, subscriber-notification and alarm. The value that is used is the one specified by the lowest-numbered enabled attack detector that has a configured value for the attack type. If none exists, the configuration of the default attack detector is used.

Use the following commands to configure and enable attack detection:

- **[no] attack-filter protocol** *protocol* **attack-direction** *direction*
- **attack-detector (default| *number*) protocol** *protocol* **attack-direction** *direction* **side** *side* **action** *action* **[open-flows** *number* **suspected-flows-rate** *number* **suspected-flows-ratio** *number***]**
- **attack-detector (default| *number*) protocol** *protocol* **attack-direction** *direction* **side** *side* **(notify-subscriber|don't-notify-subscriber)**
- **attack-detector (default| *number*) protocol** *protocol* **attack-direction** *direction* **side** *side* **(alarm|no-alarm)**
- **default attack-detector (default| *number*) protocol** *protocol* **attack-direction** *direction* **side** *side*
- **default attack-detector default**
- **default attack-detector** *number*
- **default attack-detector (all-numbered|all)**
- **attack-detector** *number* **access-list comment**
- **attack-detector** *number* **(TCP-dest-ports|UDP-dest-ports) (all|(*port1 [port2 …*]))**
- **[no] attack-filter subscriber-notification ports** *port1*

Enabling Specific-IP Detection

- Options, page 12-10
- How to Enable Specific-IP Detection, page 12-10
- How to Enable Specific-IP Detection for the TCP Protocol Only for all Attack Directions, page 12-11
- How to Enable Specific-IP Detection for the TCP Protocol for Port-Based Detections Only for Dual-Sided Attacks, page 12-11
- How to Disable Specific-IP Detection for Protocols Other than TCP, UDP, and ICMP for all Attack Directions, page 12-11
- How to Disable Specific-IP Detection for ICMP for Single-Sided Attacks Defined by the Source IP, page 12-11

By default, specific-IP detection is enabled for all attack types. You can configure specific IP detection to be enabled or disabled for a specific, defined situation only, depending on the following options:

- For a selected protocol only.
- For TCP and UDP protocols, for only port-based or only port-less detections.
- For a selected attack direction, either for all protocols or for a selected protocol.

Options

The following options are available:

- **protocol—**The specific protocol for which specific IP detection is to be enabled or disabled.
  - Default—All protocols (no protocol specified)
- **attack direction—**Defines whether specific IP detection is enabled or disabled for single sided or dual sided attacks.
  - Default—All directions
- **destination port** (TCP and UDP protocols only)—Defines whether specific IP detection is enabled or disabled for port-based or port-less detections.
  - Default—Both port-based or port-less
- Use the **no** form of the command to disable the configured specific-IP detection.

How to Enable Specific-IP Detection

From the SCE(config if)# prompt, enter:

Command	Purposes
attack-filter [protocol (((TCP|UDP) [dest-port (specific|not-specific|both)])|ICMP|other)] [attack-direction (single-side-source|single-side-destination|single-side-both|dual-sided|all)]	Enables specific-IP detection.

How to Enable Specific-IP Detection for the TCP Protocol Only for all Attack Directions

From the SCE(config if)# prompt, enter:

Command	Purpose
attack-filter protocol TCP	Enables specific-IP detection for the TCP protocol only for all attack directions.

How to Enable Specific-IP Detection for the TCP Protocol for Port-Based Detections Only for Dual-Sided Attacks

From the SCE(config if)# prompt, enter:

Command	Purpose
attack-filter protocol TCP dest-port specific attack-direction dual-sided	Enables specific-IP detection for the TCP protocol for port-based detections only for dual-sided attacks.

How to Disable Specific-IP Detection for Protocols Other than TCP, UDP, and ICMP for all Attack Directions

From the SCE(config if)# prompt, enter:

Command	Purpose
no attack-filter protocol other	Disables specific-IP detection for protocols other than TCP, UDP, and ICMP for all attack directions.

How to Disable Specific-IP Detection for ICMP for Single-Sided Attacks Defined by the Source IP

From the SCE(config if)# prompt, enter:

Command	Purpose
no attack-filter protocol ICMP attack-direction single-side-source	Disable specific-IP detection for ICMP for single-sided attacks defined by the source IP.

Configuring the Default Attack Detector

- Options, page 12-12
- How to Define the Default Action and Optionally, the Default Thresholds, page 12-13
- How to Reinstate the System Defaults for a Selected Set of Attack Types, page 12-13
- How to Reinstate the System Defaults for All Attack Types, page 12-14

Use these commands to configure the values for the default attack detector for the following parameters:

- Attack handling action
- Thresholds
- Subscriber notification
- Sending an SNMP trap

If a specific attack detector is defined for a particular attack type, it will override the configured default attack detector.

Options

The following options are available:

- **attack-detector—**The attack detector being configured; in this case, the default attack detector.
- **protocol—**Defines the protocol to which the default attack detector applies.
- **attack-direction—**Defines whether the default attack detector applies to single sided or dual sided attacks.
- **destination port** {TCP and UDP protocols only)—Defines whether the default attack detector applies to port-based or port-less detections.
- **side—**Defines whether the default attack detector applies to attacks originating at the subscriber or network side.
- **action—**Default action:
  - **report** (default)—Report beginning and end of the attack by writing to the attack-log.
  - **block—**Block all further flows that are part of this attack, the Cisco SCE platform drops the packets.
- **Thresholds**:
  - **open-flows-rate—**Default threshold for rate of open flows. suspected-flows-rate—Default threshold for rate of suspected DDoS flows.
  - **suspected-flows-ratio—**Default threshold for ratio of suspected flow rate to open flow rate.
- Use the appropriate keyword to enable or disable subscriber notification by default:
  - **notify-subscriber—**Enable subscriber notification.
  - **don't-notify-subscriber—**Disable subscriber notification.
- Use the appropriate keyword to enable or disable sending an SNMP trap by default:
  - **alarm—**Enable sending an SNMP trap.
  - **no-alarm—**Disable sending an SNMP trap.

How to Define the Default Action and Optionally, the Default Thresholds

Defaults

The default values for the default attack detector are:

- Action—Report
- Thresholds—Varies according to the attack type
- Subscriber notification—Disabled
- Sending an SNMP trap—Disabled

Step 1 From the SCE(config if)# prompt, type a**ttack-detector default protocol (((TCP|UDP) [dest-port (specific|not- specific|both)])|ICMP|other|all) attack-direction (single-side-source|single-side-destination|single-side-both|dual-sided|all) side (subscriber|network|both) [action (report|block)] [open-flows-rate** *number* **suspected-flows-rate** *rate* **suspected-flows-ratio** *ratio*] and press **Enter**.

Configures the default attack detector for the defined attack type.

Step 2 From the SCE(config if)# prompt, type **attack-detector default protocol (((TCP|UDP) [dest-port (specific|not- specific|both)])|ICMP|other|all) attack-direction (single-side-source|single-side-destination|single-side-both|dual-sided|all) side (subscriber|network|both) (notify-subscriber|don't-notify-subscriber)** and press **Enter**.

Enables or disables subscriber notification by default for the defined attack type.

The attack type must be defined the same as in Step 1.

Step 3 From the SCE(config if)# prompt, type **attack-detector default protocol (((TCP|UDP) [dest-port (specific|not- specific|both)])|ICMP|other|all) attack-direction (single-side-source|single-side-destination|single-side-both|dual-sided|all) side (subscriber|network|both) (alarm|no-alarm)** and press Enter.

Enables or disables sending an SNMP trap by default for the defined attack type.

The attack type must be defined the same as in Step 1.

How to Reinstate the System Defaults for a Selected Set of Attack Types

Use the following command to delete user-defined default values for action, thresholds, subscriber notification, and sending an SNMP trap for a selected set of attack types, and reinstate the system defaults.

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

Command	Purpose
default attack-detector default protocol (((TCP|UDP) [dest-port (specific|not-specific|both)])|ICMP|other|all) attack-direction (single-side-source|single-side-destination|single-side-both|dual-sided|all) side (subscriber|network|both)	Reinstates the system defaults for the defined attack types.

How to Reinstate the System Defaults for All Attack Types

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

Command	Purpose
default attack-detector default	Reinstates the system defaults for the defined attack types.

Specific Attack Detectors

Use these commands to define thresholds, actions, subscriber notification setting, and sending an SNMP trap for a specific attack detector for selected set of attack types.

- Options, page 12-14
- How to Enable a Specific Attack Detector and Assign it an ACL, page 12-15
- How to Define the Action and Optionally the Thresholds for a Specific Attack Detector, page 12-16
- How to Define the Subscriber Notification Setting for a Specific Attack Detector, page 12-16
- How to Define the SNMP Trap Setting for a Specific Attack Detector, page 12-16
- How to Define the List of Destination Ports for TCP or UDP Protocols for a Specific Attack Detector, page 12-17
- How to Delete User-Defined Values, page 12-17
- How to Disable a Specific Attack Detector, page 12-17
- How to Disable All Non-default Attack Detectors, page 12-18
- How to Disable All Attack Detectors, page 12-18

Options

A specific attack detector may be configured for each possible combination of protocol, attack direction, and side. The Cisco SCE platform supports a maximum of 100 attack detectors. Each attack detector is identified by a number (1-100). Each detector can be either disabled (default) or enabled. An enabled attack detector must be configured with the following parameters:

- **access-list—**The number of the Access-Control List (ACL) associated with the specified attack detector. The ACL identifies the IP addresses selected by this detector. (See "Configuring Access Control Lists (ACLs)" section on page 5-32.)
  - For dual-ip detections, the destination IP address is used for matching with the ACL.
  - Use the "none" keyword to indicate that all IP addresses are permitted by this attack-detector.

  This option is useful when using the command to define a port list, and the desired configuration should be set for all IP addresses.
- **comment—**For documentation purposes.

In addition, an enabled attack detector may contain the following settings:

- **TCP-port-list/UDP-port-list—**Destination port list for the specified protocol. TCP and UDP protocols may be configured for specified ports only. This is the list of specified destination ports per protocol.

  Up to 15 different TCP port numbers and 15 different UDP port numbers can be specified.

  Configuring a TCP/UDP port list for a given attack detector affects only attack types that have the same protocol (TCP/UDP) and are port-based (i.e. detect a specific destination port). Settings for other attack types are not affected by the configured port list(s).

The following settings are configurable for each attack type in each attack detector. Each setting can either be in a 'not configured' state (which is the default), or be configured with a specific value.

- **action—**Action:
  - **report** (default)—Report beginning and end of the attack by writing to the attack-log.
  - **block—**Block all further flows that are part of this attack, the Cisco SCE platform drops the packets.
- **Thresholds**:
  - **open-flows-rate—**Default threshold for rate of open flows. suspected-flows-rate—Default threshold for rate of suspected DDoS flows.
  - **suspected-flows-ratio—**Default threshold for ratio of suspected flow rate to open flow rate.
- Use the appropriate keyword to enable or disable subscriber notification by default:
  - **notify-subscriber—**Enable subscriber notification.
  - **don't-notify-subscriber—**Disable subscriber notification.
- Use the appropriate keyword to enable or disable sending an SNMP trap by default:
  - **alarm—**Enable sending an SNMP trap.
  - **no-alarm—**Disable sending an SNMP trap.

How to Enable a Specific Attack Detector and Assign it an ACL

From the SCE(config if)# prompt, type;

Command	Purpose
attack-detector *number* **access-list (***aclnumber* |**none) [comment** *comment***]**	Enables the attack detector and assigns it the specified ACL.

How to Define the Action and Optionally the Thresholds for a Specific Attack Detector

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

Command	Purpose
attack-detector *number* **protocol (((TCP|UDP) [dest-port (specific|not-specific|both)])|ICMP|other|all) attack-direction (single-side-source|single-side-destination|single-side-both|dual-sided|all) side (subscriber|network|both) [action (report|block)] [open-flows-rate** *number* **suspected-flows-rate** *rate* **suspected-flows-ratio** *ratio*]	Defines the action of the specified attack detector.

How to Define the Subscriber Notification Setting for a Specific Attack Detector

Use the following command to set the subscriber notification setting for a given attack detector and selected set of attack types.

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

Command	Purpose
attack-detector *number* **protocol (((TCP|UDP) [dest-port (specific|not-specific|both)])|ICMP|other|all) attack-direction (single-side-source|single-side-destination|single-side-both|dual-sided|all) side (subscriber|network|both) (notify-subscriber|don't-notify-subscriber)**	Defines the subscriber notification setting for the specified attack detector.

How to Define the SNMP Trap Setting for a Specific Attack Detector

Use the following command to enable or disable sending an SNMP trap for a given attack detector and selected set of attack types.

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

Command	Purpose
attack-detector *number* **protocol (((TCP|UDP) [dest-port (specific|not-specific|both)])|ICMP|other|all) attack-direction (single-side-source|single-side-destination|single-side-both|dual-sided|all) side (subscriber|network|both) (alarm|no-alarm)**	Defines the SNMP trap setting for the specified attack detector.

How to Define the List of Destination Ports for TCP or UDP Protocols for a Specific Attack Detector

Use the following command to define the list of destination ports for specific port detections for TCP or UDP protocols.

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

Command	Purpose
attack-detector *number* **TCP-port-list|UDP-port-list (all|(***port1* **[,** *port2, port3…* **])**	Defines the port list for the specified protocol and attack detector.

How to Delete User-Defined Values

Use the following command to remove settings of action, thresholds, subscriber notification, and sending an SNMP trap for a specific attack detector and selected set of attack types.

Removing these settings for a given attack type restores them to the default 'not configured' state, which means that the attack detector does not take part in determining the response for attacks of this attack type.

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

Command	Purpose
default attack-detector *number* **protocol (((TCP|UDP) [dest-port (specific|not-specific|both)])|ICMP|other|all) attack-direction (single-side-source|single-side-destination|single-side-both|dual-sided|all) side (subscriber|network|both)**	Deletes the configured attack detector settings for the specified attack type.

How to Disable a Specific Attack Detector

Use the following command to disable a specific attack detector, configuring it to use the default action, threshold values and subscriber notification for all protocols, attack directions and sides.

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

Command	Purpose
default attack-detector *number*	Disables the specified attack detector.

How to Disable All Non-default Attack Detectors

Use the following command to disable all non-default attack detectors, configuring them to use the default values.

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

Command	Purpose
default attack-detector all-numbered	Disables all non-default attack detectors.

How to Disable All Attack Detectors

Use the following command to disable all attack detectors, configuring them to use the default values.

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

Command	Purpose
default attack-detector all	Disables all attack detectors.

Sample Attack Detector Configuration

The following configuration changes the default user threshold values used for detecting ICMP attacks, and configures an attack-detector with high thresholds for UDP attacks, preventing false detections of two DNS servers (10.1.1.10 and 10.1.1.13) as being attacked.

Step 1 From the SCE(config)# prompt, type **interface linecard 0** and press **Enter**.

Enters linecard interface configuration mode

Step 2 From the SCE(config if)# prompt, type **attack-detector default protocol ICMP attack-direction** *single-side-source* **side** *both* **action** *report* **open-flow-rate** *1000* **suspected-flows-rate** *100* **suspected-flows-ratio** *10* and press **Enter**.

Configures the default ICMP threshold and action.

Step 3 From the SCE(config if)# prompt, type **attack-detector** *1* **access-list** *3* **comment** *"DNS servers"* and press **Enter**.

Enables attack detector #1 and assigns ACL #3 to it.

Step 4 From the SCE(config if)# prompt, type **attack-detector** *1* **UDP-ports-list** *53*

Defines the list of UDP destination ports for attack detector #1 with one port, port 53

Step 5 From the SCE(config if)# prompt, type **attack-detector** *1* **protocol UDP dest-port** *specific* **attack-direction** *single-side-destination* **side** *both* **action** *report* **open-flow-rate** *1000000* **suspected-flows-rate** *1000000* and press **Enter**.

Defines the thresholds and action for attack detector #1.

Step 6 From the SCE(config if)# prompt, type **attack-detector** *1* **protocol UDP dest-port** *specific* **attack-direction** *single-side-destination* **side** *subscriber* **notify-subscriber** and press **Enter**.

Enables subscriber notification for attack detector #1.

Step 7 From the SCE(config if)# prompt, type **exit** and press **Enter**.

Exits the linecard interface configuration mode.

Step 8 Configure ACL #3, which has been assigned to the attack detector.

```
SCE(config)# access-list 3 permit 10.1.1.10
SCE(config)# access-list 3 permit 10.1.1.13
```

Subscriber Notifications

- Configuring the Subscriber Notification Port, page 12-20
- How to Remove the Subscriber Notification Port, page 12-20

Subscriber notification is a capability used- for notifying a subscriber in real-time about current attacks involving IP addresses mapped to that subscriber. Subscriber notification is configured on a per-attack-detector level, as explained above, and must also be enabled and configured by the application loaded to the Cisco SCE platform, as explained in the *Cisco Service Control Application for Broadband User Guide*.

In the current solutions, the Cisco SCE Platform notifies the subscriber about the attack by redirecting HTTP flows originating from the subscriber to the service provider's server, that should notify the subscriber that he is under attack. This raises a question regarding TCP attacks originating from the subscriber that are configured with block action. Such attacks cannot normally be notified to the subscriber using HTTP redirection, since all HTTP flows originating from the subscriber are TCP flows, and they are therefore blocked along with all other attack flows. To enable effective use of HTTP redirect, there is a CLI command that prevents blocking of TCP flows originating from the subscriber to a specified TCP port, even when the above scenario occurs.

Configuring the Subscriber Notification Port

You can define a port to be used as the subscriber notification port. The attack filter will never block TCP traffic from the subscriber side of the Cisco SCE platform to this port, leaving it always available for subscriber notification.

Options

The following option is available:

- **portnumber—**The number of the port to be used as the subscriber notification port

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

Commands	Purpose
attack-filter subscriber-notification ports *portnumber*	Defines a port to be used as the subscriber notification port.

How to Remove the Subscriber Notification Port

Commands	Purpose
no attack-filter subscriber-notification ports	Removes the subscriber notification port.

Preventing and Forcing Attack Detection

- Options, page 12-21
- Preventing Attack Filtering, page 12-21
- Forcing Attack Filtering, page 12-22

After configuring the attack detectors, the Cisco SCE platform automatically detects attacks and handles them according to the configuration. However, there are scenarios in which a manual intervention is desired, either for debug purposes, or because it is not trivial to reconfigure the Cisco SCE platform attack-detectors properly. For example:

- The Cisco SCE platform has detected an attack, but the user knows this to be a false alarm. The proper action that should be taken by the user is to configure the system with higher thresholds (for the whole IP range, or maybe for specific IP addresses or ports). However, this might take time, and, if attack handling is specified as 'Block', the user may wish to stop the block action for this specific attack quickly, leaving the configuration changes for a future time when there is time to plan the needed changes properly.

  Use the **dont-filter** command described below for this type of case.

- An ISP is informed that one of his subscribers is being attacked by a UDP attack from the network side. The ISP wants to protect the subscriber from this attack by blocking all UDP traffic to the subscriber, but unfortunately the Cisco SCE platform did not recognize the attack. (Alternatively, it could be that the attack was recognized, but the configured action was 'report' and not 'block').

  Use the **force-filter** command described below for this type of case.

The user can use the CLI attack filtering commands to do the following:

- Configure a **dont-filter** command to prevent or stop filtering of an attack related to a specified IP address
- Configure a **force-filter** command to force filtering (with a specific action) of an attack related to a specified IP address

Use the following commands to either force or prevent attack filtering:

- [no] attack-filter dont-filter
- [no] attack-filter force-filter

Options

In addition to the attack detector options described above, the following options are available:

- **ip-address—**The IP address for which to prevent attack filtering.

  If **attack -direction** is dual-sided, an IP address must be configured for both the source (*source-ip-address*) and the destination (*dest-ip-address*) sides.

Preventing Attack Filtering

Attack filtering can be prevented for a specified IP address and attack type by executing a dont-filter CLI command. If filtering is already in process, it will be stopped. When attack filtering has been stopped, it remains stopped until explicitly restored by another CLI command (either **force**-filter or **no dont-filter)**.

- "How to Remove All dont-filter Settings" section on page 12-22

Use the following commands to configure or remove a dont-filter setting for or from a specified situation:

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

Commands	Purpose
attack-filter dont-filter protocol (((TCP|UDP) [dest-port (port-number |not-specific))|ICMP|other) attack-direction (((single-side-source|single-side-destination|single-side-both) (ip *ip-address*)|(**dual-sided source-ip** *source-ip-address* **destination-ip** *dest-ip-address*)) **side (subscriber|network|both)**	Configures a dont-filter setting for a specified situation.
no attack-filter dont-filter protocol (((TCP|UDP) [dest-port (port-number |not-specific))|ICMP|other) attack-direction (((single-side-source|single-side-destination|single-side-both) (ip *ip-address*)|(**dual-sided source-ip** *source-ip-address* **destination-ip** *dest-ip-address*)) **side (subscriber|network|both)**	Removes a dont-filter setting from a specified situation.

How to Remove All dont-filter Settings

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

Command	Purpose
no attack-filter dont-filter all	Removes all dont-filter settings.

Forcing Attack Filtering

Attack filtering can be forced for a specified IP address/protocol. Forced attack filtering will continue until undone by an explicit CLI command (either **no force-filter** or **dont-filter**).

- How to Remove All force-filter Settings, page 12-23

Use the following commands to configure or remove a force-filter setting for or from a specified situation:

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

Command	Purpose
attack-filter force-filter action (block|report) protocol (((TCP|UDP) [dest-port (port-number |not-specific))|ICMP|other) attack-direction (((single-side-source|single-side-destination|single-side-both) (ip *ip-address*)**|(dual-sided source-ip** *source-ip-address* **destination-ip** *dest-ip-address*)) **side (subscriber|network|both)[notify-subscriber]**	Configures a force-filter setting for a specified situation.
no attack-filter force-filter protocol (((TCP|UDP) [dest-port (port-number |not-specific))|ICMP|other) attack-direction (((single-side-source|single-side-destination|single-side-both) (ip *ip-address*)**|(dual-sided source-ip** *source-ip-address* **destination-ip** *dest-ip-address*)) **side (subscriber|network|both)**	Removes a force-filter setting from a specified situation.

How to Remove All force-filter Settings

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

Command	Purpose
no attack-filter force-filter all	Removes all force-filter settings.

Monitoring Attack Filtering

- Monitoring Attack Filtering Using SNMP Traps, page 12-24
- Monitoring Attack Filtering Using CLI Commands, page 12-26
- Viewing the Attack Log, page 12-32

There are three options for monitoring attack filtering and detection:

- CLI show commands
- SNMP attack detection traps
- Attack log

Monitoring Attack Filtering Using SNMP Traps

The system sends a trap at the start of a specific attack detection event, and also when a specific detection event ends, as follows:

- STARTED_FILTERING trap – String with the attack information
- STOPPED_FILTERING
  - String with the attack information
  - String with the reason for stopping

The format of the attack-information string sent when an attack begins is:

- If attack was detected in the traffic:

```
Attack detected: Attack ‘IP-info’ from ‘side’ side, protocol ‘protocol’. ‘rate1’ open
flows per second detected, ‘rate2’ Ddos-suspected flows per second detected. Action
is: 'action'.
```

- If attack was declared as a result of a **force-filter** command:

```
Attack Filter: Forced ‘forced-action’ ‘IP-info’ from ‘side’ side, protocol ‘protocol’.
Attack forced using a force-filter command.
```

The format of the attack-information string sent when an attack ends is:

- If attack was detected in the traffic:

```
End-of-attack detected: Attack ‘IP-info’ from ‘side’ side , protocol ‘protocol’.
Action is: ‘action’ Duration ‘duration’ seconds, ‘total-flows’ ‘hw-filter’
```

- If the end of the attack was declared as a result of a **no force-filter** command or a new **don't-filter** command:

```
Attack Filter: Forced to end 'action2' 'IP-info' from 'side' side, protocol
'protocol'. Attack end forced using a 'no force-filter' or a 'don't-filter' command.
```

The format of the reason string sent when an attack begins is:

- If attack end was detected in the traffic:

```
Detected attack end
```

- If the end of the attack was declared as a result of a **no force-filter** command or a new **don't-filter** command:

```
Forced attack end
```

Following are the possible values that may appear in the fields indicated in the information strings (''):

- 'action'
  - Report
  - Block
- 'forced-action' is one of the following values, depending on the configured force-filter action.
  - block of flows
  - report
- 'IP-info' is in one of the following formats, depending on the direction of the attack, and whether one or two IP addresses were detected
  - from IP address A.B.C.D
  - on IP address A.B.C.D
  - from IP address A.B.C.D to IP address A.B.C.D
- 'side'
  - subscriber
  - network
- 'protocol'
  - TCP
  - UDP
  - ICMP
  - other
- 'rate1' and 'rate2' are numbers
- 'duration' is a number.
- 'total-flows' is one of the following strings, depending on the attack action:
  - If 'action' is block: 'number' flows blocked.
  - If 'action' is report: attack comprised of 'number' flows.
- 'hw-filter'
  - If the attack was not filtered by a hardware filter: empty string
  - If the attack was filtered by a hardware filter: HW filters used, actual attack duration is probably smaller than reported above, actual amount of flows handled is probably larger than reported above.

Monitoring Attack Filtering Using CLI Commands

- How to Display a Specified Attack Detector Configuration, page 12-26
- How to Display the Default Attack Detector Configuration, page 12-28
- How to Display All Attack Detector Configurations, page 12-28
- How to Display Filter State (Enabled or Disabled), page 12-29
- How to Display Configured Threshold Values and Actions, page 12-29
- How to Display the Current Counters, page 12-30
- How to Display all Currently Handled Attacks, page 12-31
- How to Display all Existing Force-Filter Settings, page 12-31
- How to Display all Existing Don't-Filter Settings, page 12-31
- How to Display the List of Ports Selected for Subscriber Notification, page 12-31
- How to Find out Whether Hardware Attack Filtering has been Activated, page 12-32

Use these commands to monitor attack detection and filtering:

- **show interface linecard 0 attack-detector**
- **show interface linecard 0 attack-filter**
- **show interface linecard 0 attack-filter query**
- **show interface linecard 0 attack-filter current-attacks**
- **show interface linecard 0 attack-filter don't-filter**
- **show interface linecard 0 attack-filter force-filter**
- **show interface linecard 0 attack-filter subscriber-notification ports**

How to Display a Specified Attack Detector Configuration

- Options, page 12-27
- Example, page 12-27

The following information is displayed:

- Protocol Side—Whether the attack detector applies to attacks originating at the subscriber or network side.
- Direction—Whether the attack detector applies to single sided or dual sided attacks. Action to take if an attack is detected.
- Thresholds:
  - open-flows-rate—Default threshold for rate of open flows (new open flows per second).
  - suspected-flows-rate—Default threshold for rate of suspected DDoS flows (new suspected flows per second).
  - suspected-flows-ratio—Default threshold for ratio of suspected flow rate to open flow rate.
- Subscriber notification—Enabled or disabled.
- Alarm: sending an SNMP trap enabled or disabled.

Options

The following option is available:

- **number—**The number of the attack detector to display

From the SCE> prompt, type:

Command	Purpose
show interface linecard 0 attack-detector *number*	Displays a specified attack detector configuration.

Example

```
SCE>show interface LineCard 0 attack-detector 1
Detector #1:
Comment: 'Sample'
Access-list: 1
Effective only for TCP port(s) 21,23,80
Effective for all UDP ports
Protocol|Side|Direction  ||Action|     Thresholds                |Sub- |Alarm
        |   |           ||       |Open flows|Ddos-Suspected flows|notif|
        |   |           ||       |rate      |rate        |ratio  |     |
--------|----|-----------||------|----------|------------|-------|-----|-----
TCP     |net.|source-only||      |          |            |       |     |
TCP     |net.|dest-only  ||      |          |            |       |     |
TCP     |sub.|source-only||      |          |            |       |     |
TCP     |sub.|dest-only  ||      |          |            |       |     |
TCP     |net.|source+dest||      |          |            |       |     |
TCP     |sub.|source+dest||      |          |            |       |     |
TCP+port|net.|source-only||Block |          |          |       |     |Yes
TCP+port|net.|dest-only  ||      |          |            |       |     |
TCP+port|sub.|source-only||Block |          |            |       |     |Yes
TCP+port|sub.|dest-only  ||      |          |            |       |     |
TCP+port|net.|source+dest||      |          |            |       |     |
TCP+port|sub.|source+dest||      |          |            |       |     |
UDP     |net.|source-only||      |          |            |       |     |
UDP     |net.|dest-only  ||      |          |            |       |     |
UDP     |sub.|source-only||      |          |            |       |     |
UDP     |sub.|dest-only  ||      |          |            |       |     |
UDP     |net.|source+dest||      |          |            |       |     |
UDP     |sub.|source+dest||      |          |            |       |     |
UDP+port|net.|source-only||      |          |            |       |     |
UDP+port|net.|dest-only  ||      |          |            |       |     |
UDP+port|sub.|source-only||      |          |            |       |     |
UDP+port|sub.|dest-only  ||      |          |            |       |     |
UDP+port|net.|source+dest||      |          |            |       |     |
UDP+port|sub.|source+dest||      |          |            |       |     |
ICMP    |net.|source-only||      |          |            |       |     |
ICMP    |net.|dest-only  ||      |          |            |       |     |
ICMP    |sub.|source-only||      |          |            |       |Yes  |
ICMP    |sub.|dest-only  ||      |          |            |       |     |
other   |net.|source-only||      |          |            |       |     |
other   |net.|dest-only  ||      |          |            |       |     |
other   |sub.|source-only||      |          |            |       |     |
other   |sub.|dest-only  ||      |          |            |       |     |
Empty fields indicate that no value is set and configuration from
the default attack detector is used.
SCE#>
```

How to Display the Default Attack Detector Configuration

From the SCE> prompt, type:

Command	Purpose
show interface linecard 0 attack-detector default	Displays the default attack detector configuration.

Example

```
SCE>show interface LineCard 0 attack-detector default
Default detector:
Protocol|Side|Direction  ||Action|     Thresholds                |Sub- |Alarm
        |    |           ||      |Open flows|Ddos-Suspected flows|notif|
        |    |           ||      |rate      |rate        |ratio  |     |
--------|----|-----------||------|----------|------------|-------|-----|-----
TCP     |net.|source-only||Report|      1000|         500|50     |No   |No
TCP     |net.|dest-only  ||Report|      1000|         500|50     |No   |No
TCP     |sub.|source-only||Report|      1000|         500|50     |No   |No
TCP     |sub.|dest-only  ||Report|      1000|         500|50     |No   |No
TCP     |net.|source+dest||Report|       100|          50|50     |No   |No
TCP     |sub.|source+dest||Report|       100|          50|50     |No   |No
TCP+port|net.|source-only||Report|      1000|         500|50     |No   |No
TCP+port|net.|dest-only  ||Report|      1000|         500|50     |No   |No
TCP+port|sub.|source-only||Report|      1000|         500|50     |No   |No
TCP+port|sub.|dest-only  ||Report|      1000|         500|50     |No   |No
TCP+port|net.|source+dest||Report|       100|          50|50     |No   |No
TCP+port|sub.|source+dest||Report|       100|          50|50     |No   |No
UDP     |net.|source-only||Report|      1000|         500|50     |No   |No
UDP     |net.|dest-only  ||Report|      1000|         500|50     |No   |No
UDP     |sub.|source-only||Report|      1000|         500|50     |No   |No
UDP     |sub.|dest-only  ||Report|      1000|         500|50     |No   |No
UDP     |net.|source+dest||Report|       100|          50|50     |No   |No
UDP     |sub.|source+dest||Report|       100|          50|50     |No   |No
UDP+port|net.|source-only||Report|      1000|         500|50     |No   |No
UDP+port|net.|dest-only  ||Report|      1000|         500|50     |No   |No
UDP+port|sub.|source-only||Report|      1000|         500|50     |No   |No
UDP+port|sub.|dest-only  ||Report|      1000|         500|50     |No   |No
UDP+port|net.|source+dest||Report|       100|          50|50     |No   |No
UDP+port|sub.|source+dest||Report|       100|          50|50     |No   |No
ICMP    |net.|source-only||Report|       500|         250|50     |No   |No
ICMP    |net.|dest-only  ||Report|       500|         250|50     |No   |No
ICMP    |sub.|source-only||Report|       500|         250|50     |No   |No
ICMP    |sub.|dest-only  ||Report|       500|         250|50     |No   |No
other   |net.|source-only||Report|       500|         250|50     |No   |No
other   |net.|dest-only  ||Report|       500|         250|50     |No   |No
other   |sub.|source-only||Report|       500|         250|50     |No   |No
other   |sub.|dest-only  ||Report|       500|         250|50     |No   |No
SCE#>
```

How to Display All Attack Detector Configurations

From the SCE> prompt, type:

Command	Purpose
show interface linecard 0 attack-detector all	Displays all attack detector configurations.

How to Display Filter State (Enabled or Disabled)

From the SCE> prompt, type:

Command	Purpose
show interface linecard 0 attack-filter	Displays filter state (enabled or disabled).

Example

```
SCE>show interface LineCard 0 attack-filter
Enabled state:
------------------
Protocol  |Direction   |State
----------|------------|------------
TCP       |source-only |enabled
TCP       |dest-only   |enabled
TCP       |dest+source |enabled
TCP+port  |source-only |enabled
TCP+port  |dest-only   |enabled
TCP+port  |dest+source |enabled
UDP       |source-only |enabled
UDP       |dest-only   |enabled
UDP       |dest+source |enabled
UDP+port  |source-only |enabled
UDP+port  |dest-only   |enabled
UDP+port  |dest+source |enabled
ICMP      |source-only |enabled
ICMP      |dest-only   |enabled
other     |source-only |enabled
other     |dest-only   |enabled
SCE#>
```

How to Display Configured Threshold Values and Actions

Use this command to display the configured threshold values and actions a specified IP address (and port), taking into account the various specific attack detector access list configurations

Options

In addition to the attack detector options described above, the following options are available:

- **ip-address—**The IP address for which to display information.

  If **attack -direction** is dual-sided, an IP address must be configured for both the source (*source-ip-address*) and the destination (*dest-ip-address*) sides.

- **portnumber—**The port number for which to display information.

From the SCE> prompt, type:

Command	Purpose
show interface linecard 0 attack-filter query ((single-sided ip *ip-address*)**|(dual-sided source-IP** *source-ip-address* **destination-IP** *dest-ip-address*)**) [dest-port** *portnumber*] **configured**	Displays the configured threshold values and actions a specified IP address.

Examples

Example 1

This example shows a query for a single IP address.

```
SCE#>show interface linecard 0 attack-filter query single-sided ip 10.1.1.1 configured
Protocol|Side|Dir.|Action|      Thresholds           |don't- |force-|Sub- |Alarm
        |    |    |      |Open flows|Ddos-Susp. flows|filter|filter|notif|
        |    |    |      |rate      |rate      |ratio|      |      |     |
--------|----|----|------|----------|----------|-----|----- |------|-----|-----
TCP     |net.|src.|Report|      1000|       500|   50|No    |No    |   No|   No
TCP     |net.|dst.|Report|      1000|       500|   50|No    |No    |   No|   No
TCP     |sub.|src.|Report|      1000|       500|   50|No    |No    |   No|   No
TCP     |sub.|dst.|Report|      1000|       500|   50|No    |No    |   No|   No
UDP     |net.|src.|Report|      1000|       500|   50|No    |No    |   No|   No
UDP     |net.|dst.|Report|      1000|       500|   50|No    |No    |   No|   No
UDP     |sub.|src.|Report|      1000|       500|   50|No    |No    |   No|   No
UDP     |sub.|dst.|Report|      1000|       500|   50|No    |No    |   No|   No
ICMP    |net.|src.|Report|       500|       250|   50|No    |No    |   No|   No
ICMP    |net.|dst.|Report|       500|       250|   50|No    |No    |   No|   No
ICMP    |sub.|src.|Report|       500|       250|   50|No    |No    |  Yes|   No
|    |    |    |      |          |          |     |      |      |  (1)|
ICMP    |sub.|dst.|Report|       500|       250|   50|No    |No    |   No|   No
other   |net.|src.|Report|       500|       250|   50|No    |No    |   No|   No
other   |net.|dst.|Report|       500|       250|   50|No    |No    |   No|   No
other   |sub.|src.|Report|       500|       250|   50|No    |No    |   No|   No
other   |sub.|dst.|Report|       500|       250|   50|No    |No    |   No|   No
(N) below a value means that the value is set through attack-detector #N.
SCE#>
```

Example 2

This example shows a query for a single IP address, with a specified port.

```
SCE#>show interface linecard 0 attack-filter query single-sided ip 10.1.1.1 dest-port 21
configured
Protocol|Side|Dir.|Action|      Thresholds           |don't- |force-|Sub- |Alarm
        |    |    |      |Open flows|Ddos-Susp. flows|filter|filter|notif|
        |    |    |      |rate      |rate      |ratio|      |      |     |
--------|----|----|------|----------|----------|-----|----- |------|-----|-----
TCP+port|net.|src.|Block |      1000|       500|   50|No    |No    |   No|  Yes
|    |    |(1)   |      |          |          |     |      |      |     |  (1)
TCP+port|net.|dst.|Report|      1000|       500|   50|No    |No    |   No|   No
TCP+port|sub.|src.|Block |      1000|       500|   50|No    |No    |   No|  Yes
|    |    |(1)   |      |          |          |     |      |      |     |  (1)
TCP+port|sub.|dst.|Report|      1000|       500|   50|No    |No    |   No|   No
UDP+port|net.|src.|Report|      1000|       500|   50|No    |No    |   No|   No
UDP+port|net.|dst.|Report|      1000|       500|   50|No    |No    |   No|   No
UDP+port|sub.|src.|Report|      1000|       500|   50|No    |No    |   No|   No
UDP+port|sub.|dst.|Report|      1000|       500|   50|No    |No    |   No|   No
(N) below a value means that the value is set through attack-detector #N.
SCE#>
```

How to Display the Current Counters

Use this command to display the current counters for the specified attack detector for attack types for a specified IP address.

From the SCE> prompt, type:

Command	Purpose
show interface linecard 0 attack-filter query ((single-sided ip *ip-address*)**|(dual-sided source-IP** *source-ip-address* **destination-IP** *dest-ip-address*)) **[dest-port** *portnumber*] **current**	Displays the current counters.

How to Display all Currently Handled Attacks

From the SCE> prompt, type;

Command	Purpose
show interface linecard 0 attack-filter current-attacks	Displays all currently handled attacks.

How to Display all Existing Force-Filter Settings

From the SCE> prompt, type:

Command	Purpose
show interface linecard 0 attack-filter force-filter	Displays all existing force-filter settings.

How to Display all Existing Don't-Filter Settings

From the SCE> prompt, type:

Command	Purpose
show interface linecard 0 attack-filter dont-filter	Displays all existing dont-filter settings.

How to Display the List of Ports Selected for Subscriber Notification

From the SCE> prompt, type:

Command	Purpose
show interface linecard 0 attack-filter subscriber-notification ports	Displays the list of ports selected for subscriber notification.

How to Find out Whether Hardware Attack Filtering has been Activated

Step 1 From the SCE> prompt, type **show interface linecard 0 attack-filter current-attacks** and press **Enter**.

In the output from this command, look for the "HW-filter" field. If this field is "yes", the user must take into account the probable inaccuracies in the attack reporting.

Note that this information also appears in the attack log file.

```
---|---------------|-----------|------------|----------|------|------|------
---|Source IP -----|Side /     |Open rate / |Handled   |Action|HW-   |force-
---|        Dest IP|Protocol   |Susp. rate  |  flows / |      |filter|filter
---|               |Duration   |            |          |
---|---------------|-----------|------------|----------|------|------|------
    |10.1.1.1       |Subscriber|         523|      4045|Report|No    |No
    |              *|TCP       |           0| 9|       |      |
---|---------------|-----------|------------|------------|------|------|-------
```

Viewing the Attack Log

- The Attack Log, page 12-32
- How to View the Attack Log, page 12-33
- How to Copy the Attack Log to a File, page 12-33

The Attack Log

The attack-log contains a message for each specific-IP detection of attack beginning and attack end. Messages are in CSV format.

The message for detecting attack beginning contains the following data:

- IP address (Pair of addresses, if detected)
- Protocol Port number (If detected)
- Attack-direction (Attack-source or Attack-destination)
- Interface of IP address (subscriber or network)
- Open-flows-rate, suspected-flows-rate and suspected-flows-ratio at the time of attack detection
- Threshold values for the detection
- Action taken

The message for detecting attack end contains the following data:

- IP address (Pair of addresses, if detected)
- Protocol Port number (If detected)
- Attack-direction (Attack-source or Attack-destination)
- Interface of IP address
- Number of attack flows reported/blocked
- Action taken

As with other log files, there are two attack log files. Attack events are written to one of these files until it reaches maximum capacity, at which point the events logged in that file are then temporarily archived. New attack events are then automatically logged to the alternate log file. When the second log file reaches maximum capacity, the system then reverts to logging events to the first log file, thus overwriting the temporarily archived information stored in that file.

The following SNMP trap indicates that the attack log is full and a new log file has been opened

ST_LINE_ATTACK_LOG_IS_FULL



Note When the attack log is large, it is not recommended to display it. Copy a large log to a file to view it.

How to View the Attack Log

From the SCE# prompt, type:

Command	Purpose
more line-attack-log	Displays the attack log.

How to Copy the Attack Log to a File

From the SCE# prompt, type:

Command	Purposes
more line-attack-log redirect *filename*	Writes the log information to the specified file.



Managing the SCMP

Revised: September 23, 2015,

Introduction

This module provides an overview of the Service Control Management Protocol (SCMP) capabilities. It also explains the various procedures for configuring and monitoring SCMP.

- About SCMP, page 13-2
- Configuring the SCMP, page 13-9
- Monitoring the SCMP Environment, page 13-17

About SCMP

The Service Control Management Protocol (SCMP) is a protocol that integrates the Cisco SCE platform and the ISG (Intelligent Service Gateway) functionality of the Cisco routers, thereby providing a mechanism that allows the ISG and the Cisco SCE platform to manage subscriber sessions together without requiring coordination and orchestration by additional components.

- SCMP Terminology, page 13-3
- Deployment Scenarios, page 13-3
- SCMP Peer Devices, page 13-7
- SCMP Subscriber Management, page 13-8

The SCMP is a Cisco proprietary protocol that uses the RADIUS protocol with CoA (Change of Authorization) support as a transport layer. The SCMP provides connection management messages, subscriber management and subscriber accounting messages. Each subscriber in the Cisco SCE platform represents a session in the SCMP peer (as defined by the ISG terminology).

Connection management

The Cisco SCE platform initiates the connection to the peer device. On SCMP connection establishment, the Cisco SCE platform and ISG negotiate the following details:

- Introduction mode – whether the SCMP peer must send a session-provisioning message on session creation.
- Keep-alive message interval
- Protocol version

Subscriber Management

The SCMP peers can work in either of two introduction modes. These introduction modes affect only how and when a session is created on the Cisco SCE platform:

- The SCMP peer provisions the session to the Cisco SCE platform when it is created in the peer device (push)
- The Cisco SCE platform queries the SCMP peer regarding unmapped IP traffic (pull).

The SCMP uses queries as a backup to the push introduction mode, to be robust to issues such as networking problems and Cisco SCE platform reboot.

In addition to session creation, the SCMP supports the following operations:

- Change of session policy and network IDs using the update-session message
- Removal of the session when the user logs-out
- Activate-policy, which changes the session policy
- Deactivate-policy, which sets the policy value of the related anonymous-group template (based on the session manager)

Subscriber Accounting

On session creation, the Cisco SCE platform sends an accounting start message for the session and on logout, it sends an accounting stop message for the session. In addition, for each SCA BB service-counter an accounting-session is maintained (start, interim and stop messages), which provides information regarding the relevant volume, flow-count and duration.

The accounting messages are based on the new Subscriber-Accounting RDR and are sent according to the interval defined in the PQB configuration.

SCMP Terminology

SCMP terminology is similar to, but not identical to, existing Cisco SCE platform terminology. It is derived from the ISG terminology, since every Cisco SCE subscriber is actually an ISG session.

- subscriber – The client who is purchasing service from the Service Provider and is receiving the bill.
- User – A member, employee or guest at the subscriber household or business using the service.
- Session – A logically identifiable entity on the service gateway that represents communication with a peer. It is based on a unique combination of one or more Identity Keys such as an IP address, a subnet, a MAC address, a tunnel termination interface (PPP) or a port.

  Each session is assigned a unique identifier.
- Flow – Characterized by several parameters identifiable from the traffic such as source IP address, destination IP address, source port, destination port, protocol and in some cases direction.
- SCMP Peer – A Cisco device running IOS with the ISG module enabled.
- Identity Key – One of the keys that help identify a Session. The identity keys that are relevant to the Cisco SCE-ISG control-bus are:
  - IP Address/Subnet
  - IP Subnet
- Policy – Defines all aspects of subscriber session processing. A policy consists of conditions and actions. Traffic conditions will classify traffic and allow policing actions to be applied to the traffic. Policies may be provisioned, updated and removed. Policies may also be activated for a session or deactivated for a session. A policy may be referred to by name.

Deployment Scenarios

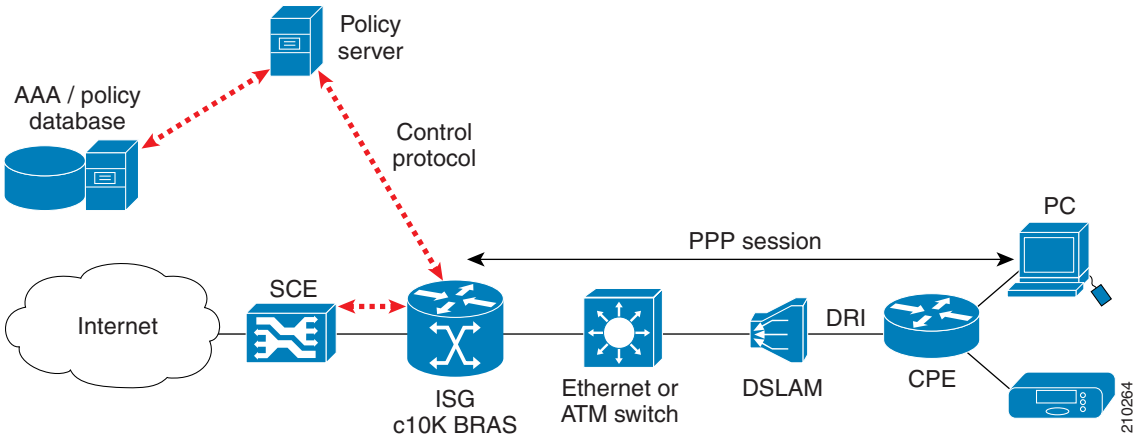
The following sections illustrate the basic types of SCMP deployment scenarios.

- 1xISG – 1xCisco SCE
- 1xISG – 2xCisco SCE (Cisco SCE cascade)
- NxISG – 2xCisco SCE (Cisco SCE cascade)
- NxISG – MxCisco SCE Via Load Balancing (MGSCP)

Single ISG Router with a Single Cisco SCE Platform (1xISG – 1xCisco SCE)

Figure 13-1 illustrates a deployment using one ISG router with a single Cisco SCE platform.

Figure 13-1 Single ISG Router with a Single Cisco SCE Platform



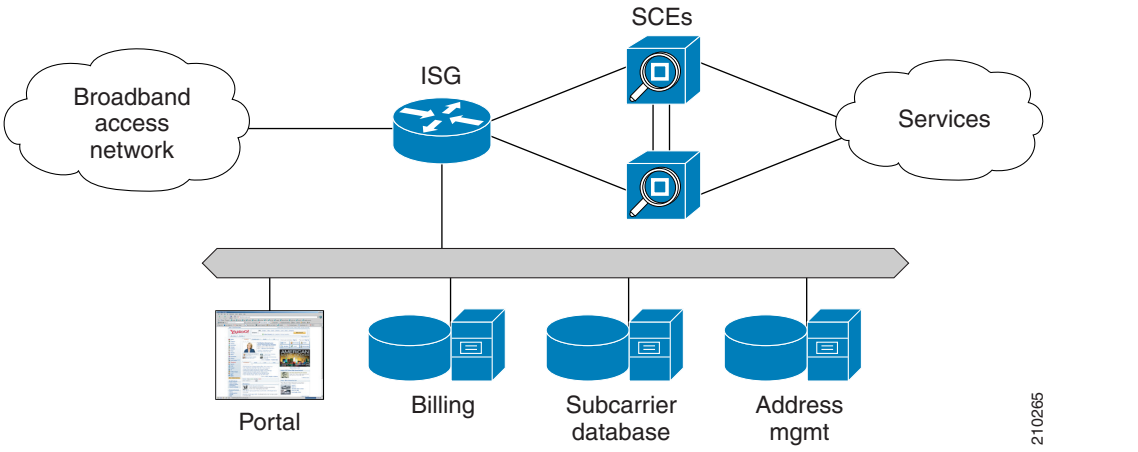
Note the following:

- The red dotted lines depict the control path communication.
- A deployment of this type might be used with ISG running on a service gateway or BRAS terminating a large number of subscribers. However, note that deploying only one Cisco SCE platform results in a single point of failure, which is not generally acceptable in an actual deployment.

Single ISG Router with Two Cascaded Cisco SCE Platforms (1xISG – 2xCisco SCE)

Figure 13-2 illustrates a deployment using one ISG router with two cascaded Cisco SCE platforms.

Figure 13-2 Single ISG Router with Two Cascaded Cisco SCE Platforms



This scenario is similar to the previous one, with ISG running on a service gateway or BRAS terminating a large number of subscribers, however a second Cisco SCE platform has been added to provide redundancy. This redundancy scheme assumes that Cisco SCE platforms are connected in a cascade, with one active Cisco SCE platform and one backup.

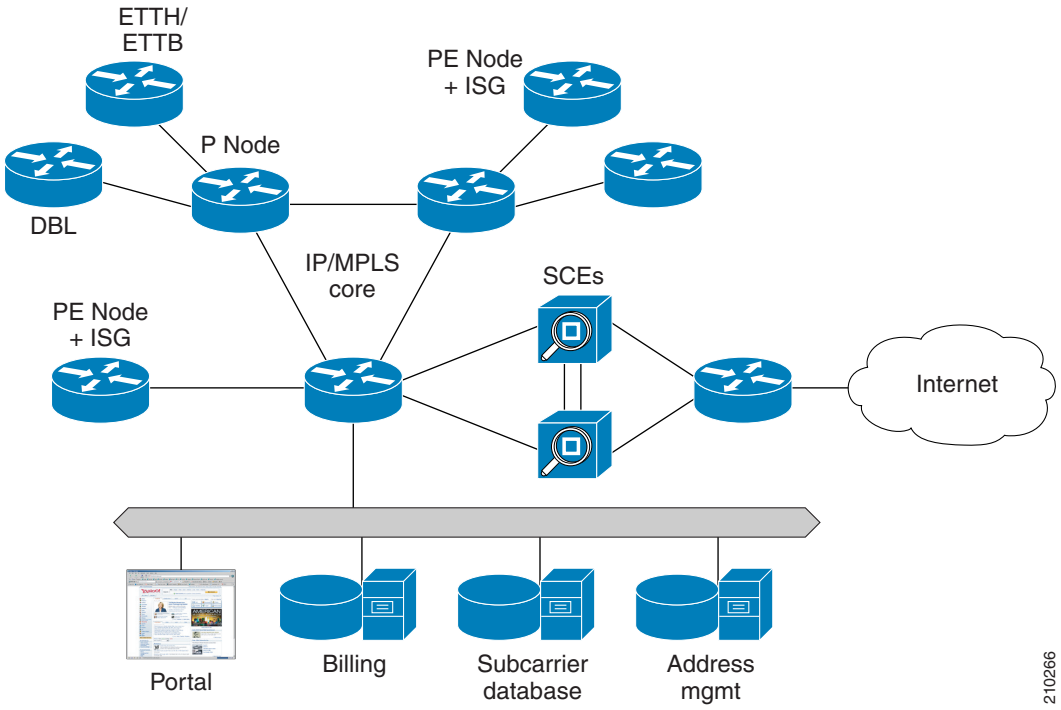
Please note the following:

- When cascaded Cisco SCE platforms are connected to one or more ISG devices, only the active Cisco SCE platform maintains a connection to the ISG devices.
- You can configure the cascaded Cisco SCE platforms to receive session info from the SCMP peer on session creation or pull the session info when the subscribers traffic traverses the Cisco SCE platform.
- An ISG device cannot push sessions to two Cisco SCE platforms at the same time

Multiple ISG Routers with Two Cascaded Cisco SCE Platforms (NxISG – 2xCisco SCE)

Figure 13-3 illustrates a deployment using multiple ISG routers with two cascaded Cisco SCE platforms.

Figure 13-3 Multiple ISG Routers with Two Cascaded Cisco SCE Platforms



Many SPs require an edge platform with MPLS functionality to support L2 and L3 VPN services for business customers, with the possibility of running subscriber management functions for residential and business subscribers terminating on the same platform. If advanced services requiring deep packet inspection are offered, we recommend locating the Cisco SCE platforms centrally, just before traffic requiring such services exits the SP network, since not all traffic needs to be processed by Cisco SCE platforms. Please note the following:

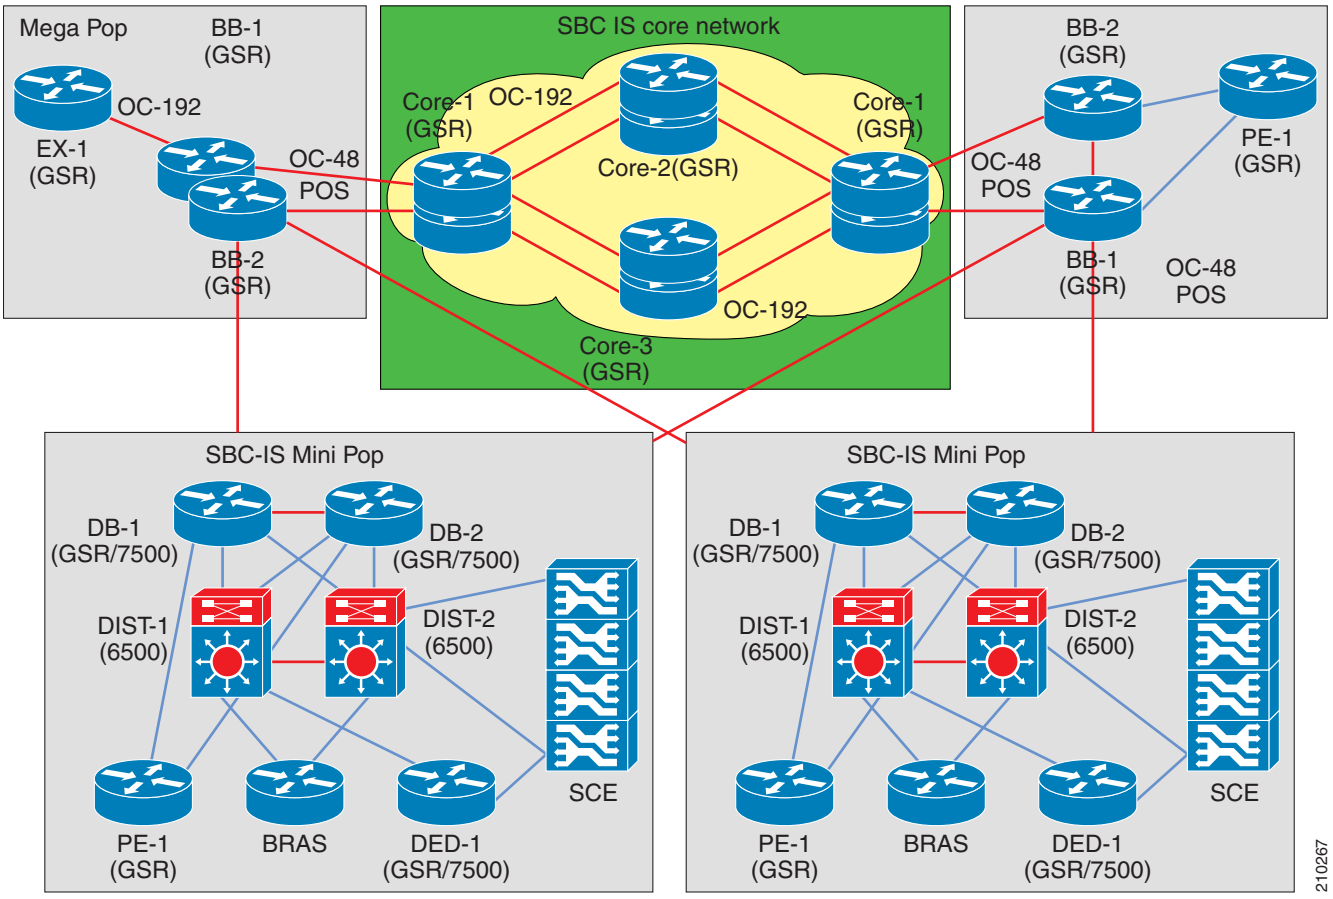
- When cascaded Cisco SCE platforms are connected to one or more ISG devices, only the active Cisco SCE platform maintains a connection to the ISG devices.

- You can configure the cascaded Cisco SCE platforms to receive session info from the SCMP peer on session creation or pull the session info when the subscribers traffic traverses the Cisco SCE platform.
- An ISG device cannot push sessions to two Cisco SCE platforms at the same time.

Multiple ISG Routers with Multiple Cisco SCE Platforms via Load Balancing (NxISG – MxCisco SCE)

Figure 13-4 illustrates a deployment using multiple ISG routers with multiple Cisco SCE platforms via load balancing. This is the scenario required for a MGSCP deployment.

Figure 13-4 Multiple ISG Routers with Multiple Cisco SCE Platforms via Load Balancing



This scenario includes several Cisco SCE platforms connected to a Cisco 7600 router. For efficient control of subscriber flows, the same Cisco SCE platform must process both directions of each subscriber flow, since the Cisco SCE platform keeps the subscriber context. The Cisco 7600 router to which the Cisco SCE platforms are connected acts as a dispatching element, distributing subscriber flows between Cisco SCE platforms and guaranteeing that all flows of a specific subscriber will pass through the same Cisco SCE platform.

This scenario assumes that one (or sometimes more) of the devices in the cluster is redundant.

Please note the following:

- An ISG device cannot push sessions to two Cisco SCE platforms at the same time.

- You must configure multiple Cisco SCE platforms with load-balancing (MGSCP) to work in pull integration mode.

SCMP Peer Devices

An SCMP peer device is a Cisco device running IOS with the ISG module enabled. The Cisco SCE platform supports the ability to communicate with several SCMP peer devices at the same time. However, each peer device manages its own subscribers and the corresponding subscriber network IDs. The Cisco SCE platform recognizes which subscribers belong to which peer device. There are two mechanisms that accomplish this:

- Login operation

  Each SCMP peer device is assigned a unique ID called the Manager-Id. This ID is attached to each subscriber from the moment it is created in the subscriber database, based on the SCMP peer that logged-in the subscriber.

- Anonymous groups

  An anonymous group is a specified IP range, possibly assigned a subscriber template (see "Anonymous Groups and Subscriber Templates" section on page 10-7.)

  SCMP associates each SCMP peer device with at least one anonymous group. SCMP generates subscribers for this anonymous group when it detects traffic from the SCMP peer device that is not mapped to any subscriber. SCMP assigns the SCMP peer manager-Id to this generated anonymous-subscriber. If you have assigned a subscriber template to the group, the anonymous subscribers generated have properties as defined by that template. If you have not assigned a subscriber template, the default template is used.

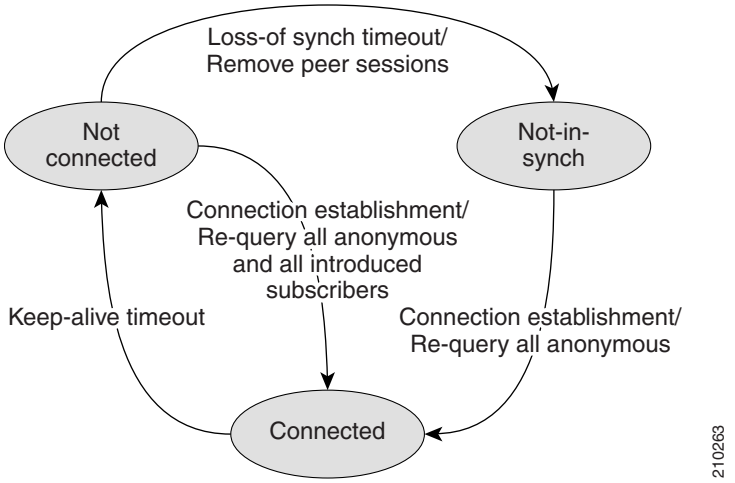
One Cisco SCE platform supports a maximum of 20 SCMP peer devices.

Connection Management

The SCMP attempts to maintain an open connection to each peer device.

Figure 13-5 illustrates the SCMP connection state functionality.

Figure 13-5 SCMP Connection State Functionality



The loss-of-sync timeout prevents the Cisco SCE platform from retaining sessions that are obsolete and whose identity-keys have been replaced or moved to other sessions thus miss-classification risk is limited.

SCMP Subscriber Management

Subscriber virtualization allows multiple SCMP peer devices to simultaneously manage subscribers in the Cisco SCE platform without interfering with each other. (Note that each device must handle a distinct set of subscribers and network IDs.)

The following mechanisms support subscriber virtualization:

- SCMP adds the Manager-Id field to each subscriber record in the database.
- All SCMP subscriber provisioning operations include the Manager-Id parameter for each subscriber.
- SCMP performs synchronizations in the context of the Manager-Id.
- SCMP dispatches queries according to the configuration of the anonymous subscriber groups.

GUID and Subscriber ID

The SCMP requires the use of a globally unique identifier (GUID) that is created by and identifies each SCMP peer device. The GUID is a 16-character-long ASCII string. The Cisco SCE platform uses the GUID for all communication with the SCMP peer.

SCMP creates the Cisco SCE subscriber ID from the concatenation of any or all the following user-related RADIUS attributes, with the GUID as the suffix.

- Calling-Station-Id
- NAS-port-Id
- User-Name

The user defines this subscriber ID structure via CLI.

Configuring the SCMP

- Configuring SCMP Parameters, page 13-9
- Adding an SCMP Peer Device, page 13-12
- Deleting Subscribers Managed by an SCMP Peer Device, page 13-14
- Deleting an SCMP Peer Device, page 13-14
- Defining the Subscriber ID, page 13-14
- Configuring the RADIUS Client, page 13-15

Configuring SCMP Parameters

- How to Enable the SCMP, page 13-9
- How to Disable the SCMP, page 13-10
- How to Configure the SCMP Peer Device to Push Sessions, page 13-10
- Configuring the SCMP Peer Device to Force Each Subscriber to Single Cisco SCE Platform, page 13-10
- Defining the Keep-alive Interval Parameter, page 13-11
- Defining the Reconnect Interval Parameter, page 13-11
- Defining the Loss-of-Sync Timeout Parameter, page 13-12

You can configure the following options for the SCMP:

- Enable the SCMP
- Configure the SCMP peer device to push sessions to the Cisco SCE platform
- Allow the SCMP peer device to provision each subscriber to only one Cisco SCE platform.
- Define the SCMP keep-alive interval
- Define the SCMP reconnect interval
- Define the loss-of-sync timeout
- Define the subscriber ID structure

How to Enable the SCMP

By default, the SCMP is disabled.

From the SCE(config)# prompt, type:

Command	Purpose
scmp	Enables SCMP.

How to Disable the SCMP

From the SCE(config)# prompt, type:

Command	Purpose
no scmp	Disables SCMP.

How to Configure the SCMP Peer Device to Push Sessions

When SCMP establishes a connection with an SCMP peer device, it informs the device whether the SCMP is configured to push sessions or to wait till the sessions are pulled by the Cisco SCE platform.

Use this command to specify push mode. Use the no form of the command to specify pull mode. This configuration takes effect only after the connection is re-established.

Default is disabled (pull mode).

From the SCE(config)# prompt, type:

Command	Purpose
scmp subscriber send-session-start	Configures the SCMP peer device to push sessions.

How to Disable Pushing Sessions

Use this command to disable pushing sessions to the Cisco SCE platform. This means that the Cisco SCE platform will pull all sessions from the SCMP peer.

From the SCE(config)# prompt, type:

Command	Purpose
no scmp subscriber send-session-start	Disables pushing sessions.

Configuring the SCMP Peer Device to Force Each Subscriber to Single Cisco SCE Platform

When SCMP establishes a connection with an SCMP peer device, it informs the device whether the SCMP is configured to allow each subscriber to be provisioned to only one Cisco SCE platform.

Use this command to configure the SCMP peer device to verify that each subscriber is provisioned to only one Cisco SCE platform. If a subscriber was provisioned to a different Cisco SCE platform, the SCMP removes it from the previous Cisco SCE platform and provisions it to the new Cisco SCE platform. This configuration is required in MGSCP topology where, if a failover between Cisco SCE platforms, subscribers might move from one Cisco SCE platform to another. If transferred subscribers are not cleared from the previous Cisco SCE platform, it can cause capacity issues.

Use the **no** form of the command to allow SCMP to provision subscribers to more than one Cisco SCE platform.

This configuration takes effect only after the connection is re-established.

Default is disabled (subscribers can be provisioned to more than one Cisco SCE platform).

From the SCE(config)# prompt, type:

Command	Purpose
scmp subscriber force-single-sce	Configures the SCMP peer device to verify that each subscriber is provisioned to only one Cisco SCE platform.

How to Disable Forcing Each Subscriber to Single Cisco SCE Platform

Use this command to disable forcing each subscriber to only one Cisco SCE platform. This allow subscribers to be provisioned to more than one Cisco SCE platform.

From the SCE(config)# prompt, type:

Command	Purpose
no scmp subscriber force-single-sce	Disables forcing each subscriber to a single Cisco SCE platform.

Defining the Keep-alive Interval Parameter

The keep-alive interval is the amount of time between keep-alive messages to the SCMP peer device. If the SCMP does not receive a response from the SCMP peer device within the defined interval, the connection is assumed to be down; and the SCMP changes the connection state to false and begins attempts to reconnect.

Options

The following options are available:

- **interval—**Interval between keep-alive messages from the Cisco SCE platform to the SCMP peer device in seconds
  - Default = 5 seconds

From the SCE(config)# prompt, type:

Command	Purpose
scmp keepalive-interval *interval*	Defines the keep-alive interval parameter.

Defining the Reconnect Interval Parameter

The reconnect interval is the amount of time between attempts by the Cisco SCE platform to reconnect with an SCMP peer. The Cisco SCE platform attempts to reconnect to the SCMP peer device at the defined intervals by sending an establish-peering-request message.

Options

The following options are available:

- **interval—**Interval between attempts by the Cisco SCE platform to reconnect with an SCMP peer, in seconds
  - Default = 30 seconds

From the SCE(config)# prompt, type:

Command	Purpose
scmp reconnect-interval *interval*	Defines the reconnect interval parameter.

Defining the Loss-of-Sync Timeout Parameter

The loss of sync timeout interval is the amount of time between loss of connection between the Cisco SCE platform and an SCMP peer device and the loss-of-sync event. (To prevent miss-classification, loss-of-sync event removes all subscribers that were provisioned by the relevant SCMP peer device.)

Options

The following options are available:

- **interval—**Loss of sync timeout interval in seconds
  - Default = 90 seconds

From the SCE(config)# prompt, type:

Command	Purpose
scmp loss-of-sync-timeout *interval*	Defines the loss-of-sync timeout parameter.

Adding an SCMP Peer Device

- How to Define an SCMP Peer Device, page 13-12
- How to Assign the SCMP Peer Device to an Anonymous Group, page 13-13

Adding an SCMP peer device is a two-step process:

1. Define the device, configuring the following parameters:
   - device name
   - RADIUS host
   - RADIUS shared secret authorization
   - port number (optional)
   - accounting port number (optional)
2. Associate the device with one or more unmapped anonymous groups.

How to Define an SCMP Peer Device

Options

The following options are available:

- **peer_device_name—**User-assigned name of the SCMP peer device
- radius_hostname—IP address or host-name of the RADIUS host (if a host-name is used, it must be valid at time of the configuration)
- **shared_secret—**RADIUS shared secret

- **auth-portnumber** (optional)—Authorization port number
- **acct-portnumber** (optional)—Accounting port number

Defaults:

- auth-port#—1812
- acct-port#—1813

From the SCE(config)# prompt, type:

Command	Purpose
scmp name *peer_device_name* **radius** *radius_hostname* **secret** *shared_secret* **[auth-port** *auth-portnumber* **acct-port** *acct-portnumber***]**	Defines an SCMP peer device.

How to Assign the SCMP Peer Device to an Anonymous Group

This command defines the specified anonymous group to be the IP range of the SCMP peer device. You must define the specified SCMP peer device before assigning the anonymous group.

Options

The following options are available:

- **group-name—**Name of the anonymous subscriber group to be associated with the specified SCMP peer device.
- **range** (optional)—IP range defined for the anonymous group
- **template** (optional)—Group template assigned to the anonymous group
- **peer-device-name—**User-assigned name of the SCMP peer device

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

Command	Purpose
subscriber anonymous-group name *group-name* **IP-range** *range* **[template** *template***]** **scmp name** *peer-device-name*	Assigns the SCMP peer device to an anonymous group.

How to Remove an Anonymous Group from the SCMP Peer Device

This command removes the specified anonymous group from the SCMP peer device.

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

Command	Purpose
no subscriber anonymous-group name *group-name*	Removes an anonymous group from the SCMP peer device.

Deleting Subscribers Managed by an SCMP Peer Device

Use this command to clear all the subscribers that are managed by a specified SCMP peer device.

Options

The following options are available:

- **peer_device_name—**User-assigned name of the SCMP peer device

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

Command	Purpose
no subscriber scmp name *peer-device-name* **all**	Clears all the subscribers that are managed by a specified SCMP peer device.

Deleting an SCMP Peer Device

You cannot delete an SCMP peer device that has anonymous groups assigned to it. You must remove all associated anonymous groups before deleting the device.

Step 1 First remove all anonymous groups assigned to the device:

```
SCE(config if)# no subscriber anonymous-group name group-name [IP-range range][template template] scmp name peer-device-name
```

Step 2 Repeat this step for all anonymous groups assigned to the SCMP peer device.

Step 3 When all anonymous groups have been removed from the device, exit LineCard Interface Configuration mode

```
SCE(config if)# exit
```

Step 4 Delete the device

```
SCE(config)#no scmp name peer_device_name
```

Defining the Subscriber ID

You can define the structure of the subscriber ID via this command by specifying which of the following elements to include and in which order:

- Calling-Station-Id
- NAS-port-Id
- User-Name

The GUID is always appended at the end of the subscriber ID as defined by this command.

Note You must disable the SCMP interface before executing this command.

Options

The following options are available:

- 1st element—Any one of the following:
  - Calling-Station-Id
  - NAS-Port-Id
  - User-Name
- 2nd element (optional)—Any one of the following (if specified, usually not the option specified as the first element):
  - Calling-Station-Id
  - NAS-Port-Id
  - User-Name
- 3rd element (optional)—Any one of the following (if specified, usually the remaining option not specified as either of the first two elements):
  - Calling-Station-Id
  - NAS-Port-Id
  - User-Name

Default = no elements concatenated with the GUID

Step 1 Disable the SCMP.

```
SCE(config)#no scmp
```

Step 2 Define the subscriber ID:

```
SCE(config)#scmp subscriber id append-to-guid radius-attributes Calling-Station-Id |
NAS-Port-Id | User-Name [Calling-Station-Id | NAS-Port-Id | User-Name] [Calling-Station-Id
| NAS-Port-Id | User-Name]
```

Step 3 Enable the SCMP.

```
SCE(config)#scmp
```

Configuring the RADIUS Client

You can configure the following options for the RADIUS client

- Define the parameters for retransmitting unacknowledged messages.

The RADIUS client polls the sockets to receive the next message and calls the SCMP engine to handle it, based on the type of the received message. Messages that were not acknowledged can be retransmitted up to the configured maximum number of retries.

Options

The following options are available:

- **times—**The maximum number of times the RADIUS client can try unsuccessfully to send a message.
  - Default = 3
- **timeout** (optional)—Timeout interval for retransmitting a message, in seconds
  - Default = 1 second

From the SCE(config)# prompt, type:

Command	Purpose
ip radius-client retry limit *times* **[timeout** *timeout***]**	Configures RADIUS client.

Monitoring the SCMP Environment

- How to Monitor the SCMP, page 13-17
- Monitoring the RADIUS Client, page 13-19

You can monitor the following components of the SCMP environment:

- SCMP
- RADIUS client

How to Monitor the SCMP

- Options, page 13-17
- How to display the general SCMP configuration, page 13-17
- How to display the configuration all currently defined SCMP peer devices, page 13-18
- How to display the configuration for a specified SCMP peer device, page 13-18
- How to display the statistics for all SCMP peer devices, page 13-18
- How to display the statistics for a specified SCMP peer device, page 13-19

Use the following commands to monitor the SCMP. These commands provide the following information:

- General SCMP configuration
- Configuration of all currently defined SCMP peer devices.
- Configuration of a specified SCMP peer device.
- Statistics for either all SCMP peer devices or a specified SCMP peer device.

Options

The following options are available:

- **device-name—**The name of the specific SCMP peer device for which to display the configuration or statistics.

How to display the general SCMP configuration

From the SCE> prompt, type:

Command	Purpose
show scmp	Displays the general SCMP configuration.

Example

```
SCE>show scmp
SCMP enabled:                     yes
Keep-alive interval:              5 seconds
Loss of synchronization timeout:  90 seconds from disconnection
Reconnection interval:            30 seconds
Force subscriber on a single SCE: no
Peer sends subscriber data on session start
Subscriber Id structure: GUID
```

How to display the configuration all currently defined SCMP peer devices

From the SCE> prompt, type:

Command	Purpose
show scmp all	Displays the configuration of all currently defined SCMP peer devices.

How to display the configuration for a specified SCMP peer device

From the SCE> prompt, type:

Command	Purpose
show scmp name *device-name*	Displays the configuration for a specified SCMP peer device.

Example

```
SCE>show scmp name isg
SCMP Connection 'isg' status:
10.56.208.91 auth-port 1812 acct-port 1813
Connection state:      Connected
Peer protocol-version: 1.0
Keep-alive interval:   5 seconds
Force single SCE:      No
Send session start:    Yes
Time connected:        9 seconds
```

How to display the statistics for all SCMP peer devices

From the SCE> prompt, type:

Command	Purpose
show scmp all counters	Displays the statistics for all SCMP peer devices.

How to display the statistics for a specified SCMP peer device

From the SCE> prompt, type:

Command	Purpose
show scmp name *peer_device_name* **counters**	Displays the statistics for a specified SCMP peer device.

Example

```
SCE>show scmp name isg counters
SCMP Connection 'isg' counters:
Total messages sent:             72
Total messages received:         72
Establish requests sent:         1
Establish replies received:      1
Accounting requests sent:        20
Accounting replies received:     20
Subscriber queries sent:         0
Subscriber query response recv:  0
Request retry exceeded:          0
Requests replied with errors:    0
Subscriber requests received:    50
Subscriber responses sent:       50
Failed Requests:                 0
Keep-alive sent:                 1
Keep-alive received:             1
```

Monitoring the RADIUS Client

Use the following command to monitor the SCMP RADIUS client. This command displays the general configuration of the RADIUS client.

From the SCE> prompt, type:

Command	Purpose
show ip radius-client	Monitors the SCMP RADIUS client.



Value-Added Services (VAS) Traffic Forwarding

Revised: September 23, 2015

Introduction

This chapter provides an overview of VAS traffic forwarding, explaining what is it and how it works. It also explains the various procedures for configuring and monitoring the VAS traffic forwarding.

- Information About VAS Traffic Forwarding, page 14-2
- How VAS Traffic Forwarding Works, page 14-3
- VAS Redundancy, page 14-10
- VAS Status and VAS Health Check, page 14-12
- VAS Traffic Forwarding Topologies, page 14-14
- SNMP Support for VAS, page 14-17
- Interactions Between VAS Traffic Forwarding and Other Cisco SCE Platform Features, page 14-18
- Configuring VAS Traffic Forwarding, page 14-20
- Monitoring VAS Traffic Forwarding, page 14-32
- Intelligent Traffic Mirroring, page 14-36

Information About VAS Traffic Forwarding

The VAS feature uses the Cisco SCE platform to access an external "expert system" for classification and control of services not supported by SCA BB. Using the VAS feature, you can forward selected flows to an external, third-party system for per-subscriber processing in addition to the existing services and functions of the SCA BB solution. For example, this feature can be used to forward selected subscriber traffic to third-party servers for intrusion detection or content-filtering.

The VAS feature enables you to divert a specified part of the traffic stream to an individual VAS server or a cluster of servers. The diversion of the traffic stream is based on the subscriber package, flow type, and the availability of the VAS servers. The feature provides load balancing for even distribution of the load on the various VAS servers.

The VAS feature supports multiple VAS service types using different VAS server groups. Several servers of the same type can be deployed in a group to increase the processing capacity and provide redundancy for each VAS service type.

The Cisco SCE platform performs subscriber load sharing between the active servers of the same server group. It is able to identify the active servers among the defined servers through a dedicated health check mechanism.

VAS Service Goals

The VAS traffic forwarding functionality enables the Service Control solution to meet several important service goals:

- Service providers can provide a range of value-added services to their subscribers, thus increasing customer satisfaction.
- The Cisco SCE platform can forward part of the traffic to third-party devices that can provide additional, complementary services.

  The Cisco SCE platform, due to its strong classification capabilities, forwards only the part of the traffic that should get the additional service based on:

  - Subscriber awareness
  - Policy that was configured
- The Service Control solution can include value-added servers that cannot be deployed inline for various reasons (for example, they cannot support throughput or are not carrier grade for inline insertion).
- Easy interoperability and flexibility for setting different services.

  Because the VAS feature emulates a regular IP network for the third-party devices, no special support is required on the part of the third-party entity.

How VAS Traffic Forwarding Works

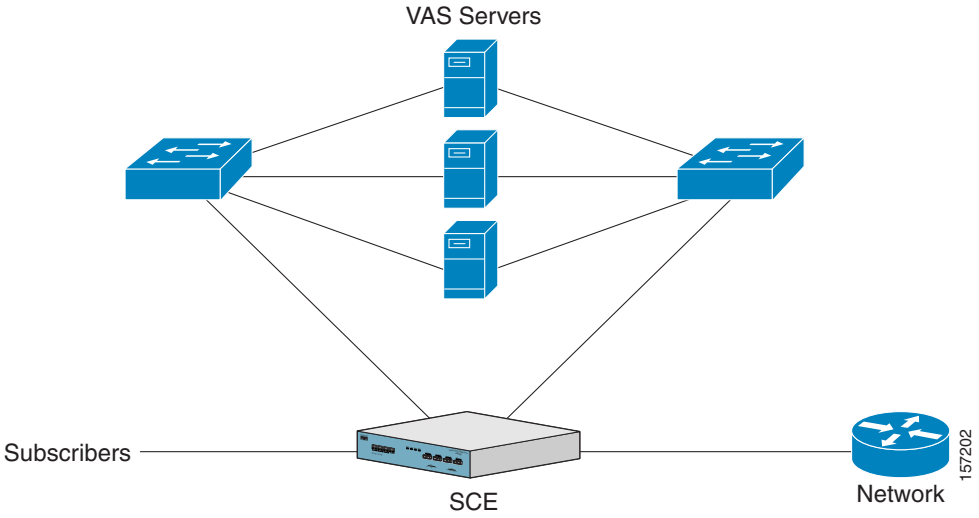
Subscribers are provisioned to the VAS services as part of the normal provisioning process of new subscribers to SCA BB.

When VAS traffic forwarding is enabled (Figure 14-1), in addition to all its basic functions, the SCA BB application classifies each flow as either a VAS flow or as a standard flow (non-VAS flow).

Flows that are classified to a VAS service get the usual SCA BB service in addition to being forwarded to the VAS servers for additional service. Traffic is processed first by the SCA BB application and then forwarded to the VAS servers.

Traffic is routed to the VAS servers using VLAN tags to identify the traffic flows.

Figure 14-1 Typical VAS Traffic Forwarding Installation



VAS traffic forwarding guidelines:

- A single Cisco SCE 8000 platform can support up to 64 VAS servers.
- A maximum of 64 Cisco SCE 8000 platforms can be connected.
- A maximum of eight VAS server groups is supported.
- The same VAS server may be used by more than one Cisco SCE platform.

Note In VAS mode, the Cisco SCE performance envelope might be up to 50 percent lower than in the normal operation mode. The exact performance envelope is specific to the traffic mix in the customer network and should be sized in advance.

The following sections provide a more detailed description of how VAS traffic forwarding works.

- Requirements for VAS Servers, page 14-4
- VAS Traffic Forwarding and SCA BB, page 14-5
- VLAN Tags for VAS Traffic Forwarding, page 14-5
- Service Flow, page 14-5
- Data Flow, page 14-6
- Load Balancing, page 14-8

Requirements for VAS Servers

Because the VAS devices are installed behind the Cisco SCE platform, they should follow the network behavior of the Cisco SCE platform. Therefore, VAS devices must meet the following two requirements:

- VAS devices must be equipped with separate interfaces for the subscriber side and separate interfaces for the network side.
- Traffic towards the subscribers should be sent from the subscriber interface and from the network interface for the internet.
- VAS devices must be transparent in Layer 2. The VAS servers must act like Layer 2 switches, in that, they are not allowed to change traffic headers or to generate new traffic..

Layer 2 Transparency

To handle non-management traffic of VAS services, follow these guidelines:

- The VAS services should work in promiscuous-mode in Layer 2 and accept packets with any destination MAC address.
- When forwarding traffic back to the network after processing, the VAS devices must preserve the original Layer 2 headers containing the MAC addresses and the VLAN tag. The VAS devices must not change the MAC addresses (destination or source) or the VLAN tags. The following restrictions apply to the injected traffic:
  - The VAS device is not permitted to initiate new flows.
  - New traffic can be injected only in the context of an existing flow.
  - When injecting traffic, the Layer 2 information (MAC addresses, VLAN tags, and the TCP/IP parameters) must be taken from the flow into which the traffic is being injected.
- A VAS device must not generate its own network transactions or relay such transactions. Network transactions such as ARP requests or pings are not permitted.

VAS Management Traffic

VAS devices that are managed inband (through the traffic interface) must meet the following requirements:

- Management traffic should either be carried over a dedicated VLAN or without any VLAN header.
- The switches that are connected to the VAS devices should be directly connected to the POP router.
- The switches that are connected to the VAS devices should be configured so that management traffic is sent directly to the router and not through the Cisco SCE platform.

VAS Traffic Forwarding and SCA BB

When VAS traffic forwarding is enabled, in addition to all its basic functions, the SCA BB application classifies each flow as either a VAS flow or as a standard flow (non-VAS flow). This classification is made on the first packet of the flow (for example, TCP SYN packet). The classification must be performed on the very first packet because the classification is used to select the routing of the packet to a VAS server or to the subscriber or network.

The VAS traffic forwarding rules table is configured using the SCA BB console. These rules map certain traffic to the VAS server groups. When a flow is classified as a VAS flow, the VAS server group for this flow is selected. If the group includes more than one VAS server, traffic is forwarded so that the subscriber load is shared between the servers on the same group.

The mapping of traffic portions per package to VAS server groups is also done using the SCA BB console.

VLAN Tags for VAS Traffic Forwarding

The traffic is routed between the Cisco SCE platform and the VAS servers by VLANs. There is a unique VLAN tag for each Cisco SCE platform and VAS server combination.

Before the traffic is forwarded to the VAS servers, the Cisco SCE platform adds the VLAN tags to the original traffic. When the traffic returns to the Cisco SCE platform, the Cisco SCE platform removes the VLAN tag it previously added, and then forwards the traffic on its original link.

The VLAN tag for each VAS server is user-configured. To preserve consistency of the traffic flow, the VAS feature requires a unique VLAN tag be configured for each Cisco SCE platform and VAS server combination.

The VLAN tag has 12 bits, divided as follows:

- The lower six bits identify the VAS server.
- The higher six bits identify the Cisco SCE platform.

For example, 0x171 = 1011 10001 = SCE 11, VAS 17

Observe the following for the six bits that identify the Cisco SCE platform:

- These six bits must be the same for all VAS servers attached to a specific Cisco SCE platform.
- These six bits must be different for VAS servers attached to different Cisco SCE platforms.

The Cisco SCE platform enforces that the user-configured VLAN tags retain this format, that is, the lower bits match the VAS server number for which the VLAN tag is configured and the higher bits match the higher bits previously configured for other VAS servers on this Cisco SCE platform. However, the Cisco SCE platform cannot determine the configuration of other Cisco SCE platforms, and therefore, it is important that the configured Cisco SCE ID (higher bits) be unique for each Cisco SCE platform.

The use of VLAN tags is an integral part of the VAS feature, and therefore, requires that the VAS device be able to work in 802.1q trunk while preserving the VLAN information.

Service Flow

The Cisco SCE platform classifies a flow to a VAS server group based on the subscriber package and the TCP/UDP ports of the flow. It then selects one server within this group to handle the flow.

The Cisco SCE platform performs load sharing between multiple VAS servers belonging to the same server group; the balance is based on the subscriber load. In other words, the Cisco SCE platform ensures that the subscribers are evenly distributed between the VAS servers in the same group.The mapping of subscriber to a VAS server (per group) is maintained even when servers are added or removed from the group either due to configuration changes or changes in the operational status of the servers in the group. The mapping changes only if the same server changes its status.

The following sections explain in more detail when and how the mapping is changed:

- Non-VAS Data Flow, page 14-7
- VAS Data Flow, page 14-7

Data Flow

In a deployment using VAS traffic forwarding, there are two types of data flows:

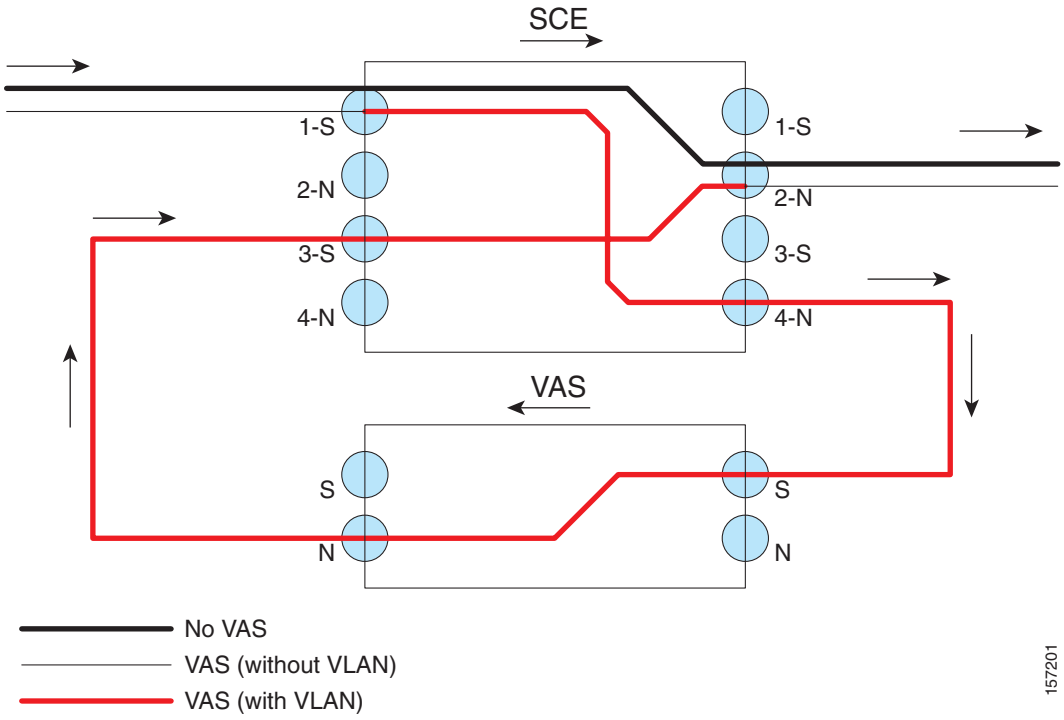
- Non-VAS flow
- VAS flow

Figure 14-2 depicts the two types of data flows running through a single Cisco SCE platform and a single VAS server.

- Ports are illustrated as two unidirectional half ports, RX (on the left side) and TX (on the right side):
  - The Cisco SCE platform has four ports.
  - The VAS server has two ports.
- For the sake of illustration, the Cisco SCE platform traffic flow direction is from left to right while the VAS traffic flow is from right to left. The arrow below the name of the element indicates the traffic flow direction.
- The Ethernet switches are omitted.
- Each line represents a flow:
  - Thick line is a non-VAS flow.
  - Thin line is a VAS flow.
  - Black line indicates part of a flow that does not have VLAN tag
  - Red line indicates part of a flow that has a VLAN tag

Figure 14-2 illustrates the data flow from the subscriber to the network. Data flow from the network to the subscriber works in the same way, but is received on the network port (N) and transmitted on the subscriber port (S).

Figure 14-2 Data Flow in a VAS System



Non-VAS Data Flow

The data flow steps for a non-VAS flow are:

1. A subscriber packet is received at the Cisco SCE platform Port 1 (S).
2. The Cisco SCE platform classifies the flow as non-VAS flow.
3. The packet is sent to the network on Port 2 (N).

VAS Data Flow

A VAS data flow is slightly more complex than the basic data flow. It is received and transmitted in the same manner as the basic non-VAS Cisco SCE platform flow, but before it is transmitted to its original destination, it flows through the VAS server.

The data flow steps for a VAS flow are:

1. A subscriber packet is received at the Cisco SCE platform Port 1 (S).
2. The Cisco SCE platform classifies the flow as a VAS flow.
3. The Cisco SCE platform adds a VLAN tag to the packet.
4. The VLAN tag is used by the Ethernet switch to route the packet to the proper VAS server.

   The packet now has a VLAN tag, which is indicated by the red line.

5. The packet is sent to the VAS subscriber port from Cisco SCE platform Port 4 (N).
6. The VAS server processes the packets and either drops the packet or sends it back to the Cisco SCE platform from the VAS network port to the Cisco SCE platform subscribers Port 3 (S).

   The VAS server passes the VLAN tag transparently. This is important to enable the Ethernet switch (not shown) to route the packet back to the proper Cisco SCE platform.
7. The Cisco SCE platform receives the packet on Port 3 (S), drops the VLAN tag, and passes the packet towards the network through Port 2 (N).

Note To use VAS, at least four interfaces should be active on the Cisco SCE device.

Load Balancing

VAS servers can be grouped logically according to their service type. Consider, for example, a system that requires both FTP caching and virus filtering. A single VAS server for each service might not have enough capacity. For example, assume that the system requires five VAS servers, three to provide FTP caching, and two to provide virus filtering. Defining two VAS server groups, for example, FTP caching and virus filtering, permits load sharing across the servers for each server group.

The subscriber package determines the VAS server group to which the flow should be attached. The selection of a specific VAS server from the VAS servers within the group is based on the current load on each VAS server. The system tries to create an equal subscriber load for all the VAS servers belonging to the same group.

In some cases, a single VAS server may be used by more than one Cisco SCE platform. Remember that the Cisco SCE platform performs load balancing only on the traffic that it sends to the VAS server; it receives no information regarding the load the VAS server may be bearing from a different Cisco SCE platform. It is vital to properly allocate available VAS servers to the Cisco SCE platforms to ensure a balanced load on each VAS server.

Load Balancing and Subscribers

The system balances the usage of the VAS servers within a VAS server group, trying to create an equal subscriber load for all the VAS servers in one VAS server group. The load balancing is subscriber based, that is, the subscribers are evenly distributed between the servers.

VAS load sharing is subscriber based rather than bandwidth based to ensure that all the traffic of the subscriber gets to the same server so that the server can make subscriber-based decisions.

The Cisco SCE platform uses the same VAS server for all the traffic of a subscriber (per server group) even if there is a change in the number of active servers in the group. Traffic from a subscriber is assigned to a new server only if the current server becomes inactive. This applies only on new flows. Flows that were already mapped to a server before it became active remain attached to it.

The mapping of subscriber to VAS servers is not saved across subscriber logouts or Cisco SCE platform reload.

Load Balancing and Subscriber Mode

Load balancing is subscriber based, therefore this feature does not work properly in subscriberless mode, because the entire traffic load would be carried by only one VAS server per group.



Tip Use anonymous mode rather than subscriberless mode with VAS traffic forwarding.

In pull mode, the first flow of the subscriber behaves as configured in the anonymous template. If no anonymous template is configured, such first flows are processed as defined by the default template. Therefore, the default template should provide a proper package, so these flows get VAS service.

VAS Redundancy

The VAS servers should be configured with high availability so that the failure of a single VAS server will not degrade total system performance and availability. This requirement must be considered when determining the number of VAS servers necessary for each VAS service.

There are two mechanisms that guarantee the performance and availability of the VAS services:

- Load sharing—The Cisco SCE platform distributes the subscribers between all the active VAS servers within a server group.
- Monitoring—The Cisco SCE platform monitors connectivity with the VAS servers and handles server failure according to the applied configuration.

In addition to failure of an individual VAS server, a complete VAS server group is considered to be failed if a defined minimum number of servers are not active.

The following sections provide more information regarding the possible points of failure in a VAS traffic forwarding deployment.

- VAS Server Failure, page 14-10
- VAS Server Group Failure, page 14-10
- Ethernet Switch Failure, page 14-11
- Disabling a VAS Server, page 14-11

VAS Server Failure

The system monitors the health of a VAS server by periodically checking the connectivity between the Cisco SCE platform and the VAS server. When the Cisco SCE platform fails to establish or maintain a connection to the server within a configurable window of time, the server is considered to be in **Down** state.

When the server is in **Down** state:

- New logged-in subscribers are distributed between the other active servers in the group.
- Subscribers that are mapped to this server are mapped to a new server if they initiate a new flow.
- The server group may move to a **Failure** state if the failure caused the number of active servers in the group to go below the minimum number configured.

If the connectivity to the server resumes, the state of the server is changed to **Up.** The server returns to the list of active servers and continues to serve subscribers that were mapped to it before the failure and have not yet been mapped to a new server during the failure time, as well as new subscribers.

VAS Server Group Failure

For each VAS server group, you can configure:

- The minimum number of active servers necessary.
- The action to take in case the actual number of active servers goes below the configured minimum.

If the minimum number of active servers equals the total number of configured servers, it means there is no redundancy and failure of one server causes the failure of the whole server group.

When the Cisco SCE platform detects that the number of active servers within a group is below the configured minimum, it changes the state of the group to **Failure**. The configured action-on-failure is then applied to all new flows mapped for that VAS server group (existing flows are not affected.)

There are two possible actions when the VAS server group has failed:

- **Block**—All new flows assigned to the failed VAS server group are blocked by the Cisco SCE platform.
- **Pass**—All new flows assigned to the failed VAS server group are considered as regular non-VAS flows, and are processed without VAS service (that is, they receive SCA BB service but not VAS service).

When the number of active servers is above the minimum and the state of the group is changed to Active again, the configured action-on-failure is no longer applied to the new flows. However, to maintain the coherency of the network, flows that were blocked or passed are not affected by the change in the state of the server group.

Ethernet Switch Failure

The Ethernet switches are a single point of failure in a VAS topology. A complete failure of an Ethernet switch causes all the VAS services to be declared as failed and the configured action (on-failure) is taken for all new VAS flows.

Disabling a VAS Server

A VAS server can be disabled for maintenance via the CLI.

No errors are reported on a disabled VAS server. However, if disabling the server reduces the number of active servers to below the minimum number configured for the group, it brings down the VAS server group because a disabled VAS server is equal to a VAS server in **Down** state.

Health check is not performed on disabled VAS servers.

VAS Status and VAS Health Check

To manage the VAS redundancy, the Cisco SCE platform needs to know the state of each VAS server. The Cisco SCE platform performs periodic health checks for all the configured VAS servers. These checks are the basis for VAS redundancy control; they enable the Cisco SCE platform to identify and react to VAS server failure, and to check the connectivity between the Cisco SCE platform and the VAS server before enabling the server to handle traffic.

The health check is performed over the VAS link, that is, the link that connects the Cisco SCE platform with the VAS servers. It validates the traffic flow between the Cisco SCE platform and the VAS server in both directions through special health check packets generated by the Cisco SCE platform.

The health check mechanism does not require special interaction with the VAS device. This is because the VAS server does not have to answer health check packets; it only passes them as they are, back to the Cisco SCE platform. As long as the packets are received by the Cisco SCE platform, the VAS server is considered to be alive. Failing to receive the packets back from the VAS server within a predefined time window is considered by the Cisco SCE platform as a failure of the VAS server and the server status is changed to **Down**.

Health check packets are:

- Carried over UDP flows.
- Contain source and destination IP addresses that can be user-configured.

  IP addresses should be:

  - Unique to the Cisco SCE platform.
  - Addresses that are not used by the network traffic (such as private IPs).

The Cisco SCE platform uses default UDP ports beginning with 63140 and 63141 for VAS Server 0, unless you configure different ports for the health check.

The Cisco SCE platform adds its own Layer 7 data on top of the UDP transport layer. This data is used by the Cisco SCE platform to validate the correctness of the packet upon retrieval.

The health check is performed under the following conditions:

- VAS mode is enabled.
- VAS server is enabled.
- Health check for the VAS server is enabled.
- Server has a VLAN tag.
- Pseudo IPs are configured for the traffic interfaces.

If the check is enabled, but any one of the conditions is not met, the server state will be **Down** (the same as if the server did not pass the health check).

Check the connectivity between the Cisco SCE platform and the VAS server before you assign the server to a server group.

The health check procedure does not require a special interface with the VAS server; the health check traffic goes through the same network channels as any other VAS traffic. However, there are two assumptions the VAS servers should fulfill:

- The VAS server should not drop traffic unless it is specifically configured to do so. Therefore, if the connectivity between the VAS server and the Cisco SCE platform is operative, the health check packets should reach the Cisco SCE platform safely.

  Alternatively, it should be possible to configure the VAS server to pass traffic on specific ports (the health check ports).

- In case of a failure, the VAS server should drop and not bypass, the traffic (cut the link), so that the Cisco SCE platform is able to identify the failure.

VAS Server States

When determining whether a VAS server is active, the system considers the following two parameters:

- User-configured Admin mode—Enabled or disabled
- VAS server state as reported by the health check

VAS Traffic Forwarding Topologies

The following sections describe the following VAS traffic forwarding topologies:

- Single Cisco SCE Platform, Multiple VAS Servers, page 14-14
- Multiple Cisco SCE Platforms, Multiple VAS Servers, page 14-15

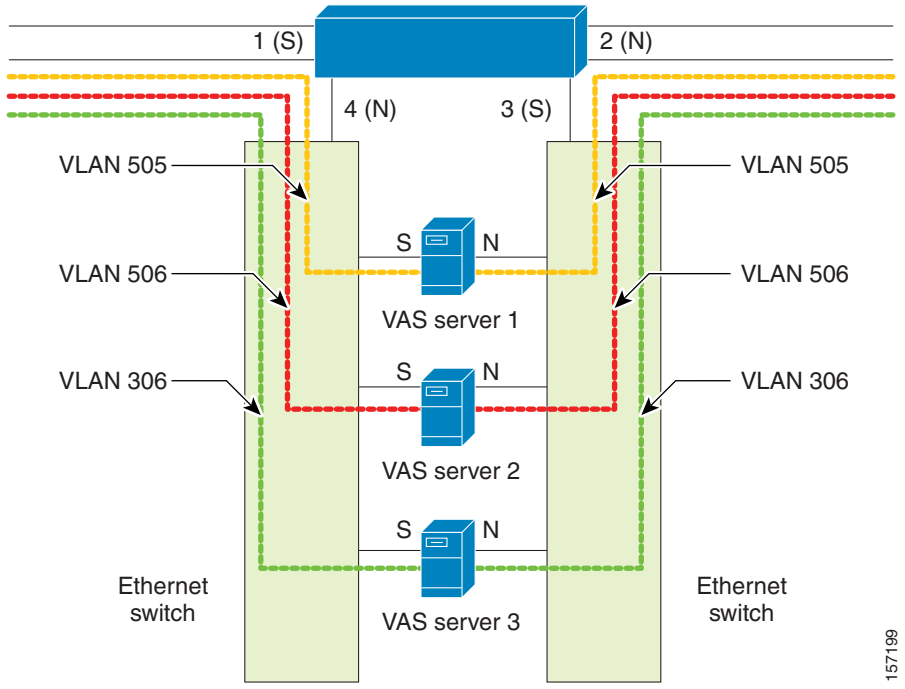
Note A topology in which a VAS server is directly connected to the Cisco SCE platform is not supported. If you want a topology of a single Cisco SCE platform connected to a single VAS server, use a switch between the Cisco SCE platform and the VAS server.

Single Cisco SCE Platform, Multiple VAS Servers

In this topology, a single Cisco SCE platform forwards VAS traffic to one or more VAS servers through two Ethernet switches (Figure 14-3).

The two Ethernet switches are necessary to avoid a situation in which a single MAC address has two ports or a single VLAN tag has two destinations. Each Ethernet switch should be configured to trunk mode with MAC learning disabled.

Figure 14-3 Single Cisco SCE Platform, Multiple VAS Servers



Data Flow

In a data flow:

1. A subscriber packet is received at Port #1 (Subscriber).
2. The Cisco SCE platform opens a flow and classifies the flow as either a non-VAS (blue) flow or as a VAS flow (red).
3. If the flow is non-VAS (blue), the Cisco SCE platform passes the packet to the network. The VAS server is not involved in this case.
4. If the flow is a VAS flow (red), the Cisco SCE platform selects the VAS server to which the packet should be sent, adds the server VLAN tag to the packet, and transmits the packet on Port #4 (Network).
5. The packet is routed by the Ethernet switch to the VAS server according to its VLAN tag (the port towards the VAS server should be the only port with this VLAN tag allowed).
6. The VAS server processes the packet and either drops or forwards it without changing the VLAN tag.
7. The packet is forwarded by the Ethernet switch to the Cisco SCE platform according to its VLAN tag (the port towards the Cisco SCE platform should be the only port with this VLAN tag allowed).
8. The Cisco SCE platform receives the packet on port #3 (Subscriber), strips the VLAN tag, and forwards the packet to the network via Port #2 (Network).

Multiple Cisco SCE Platforms, Multiple VAS Servers

n this topology, multiple Cisco SCE platforms are connected to multiple VAS servers. At least one VAS server receives traffic from more than one Cisco SCE platform; if the VAS servers are each in an exclusive relationship to a particular Cisco SCE platform, it would simply be several single Cisco SCE platform to multiple VAS server topologies grouped together.

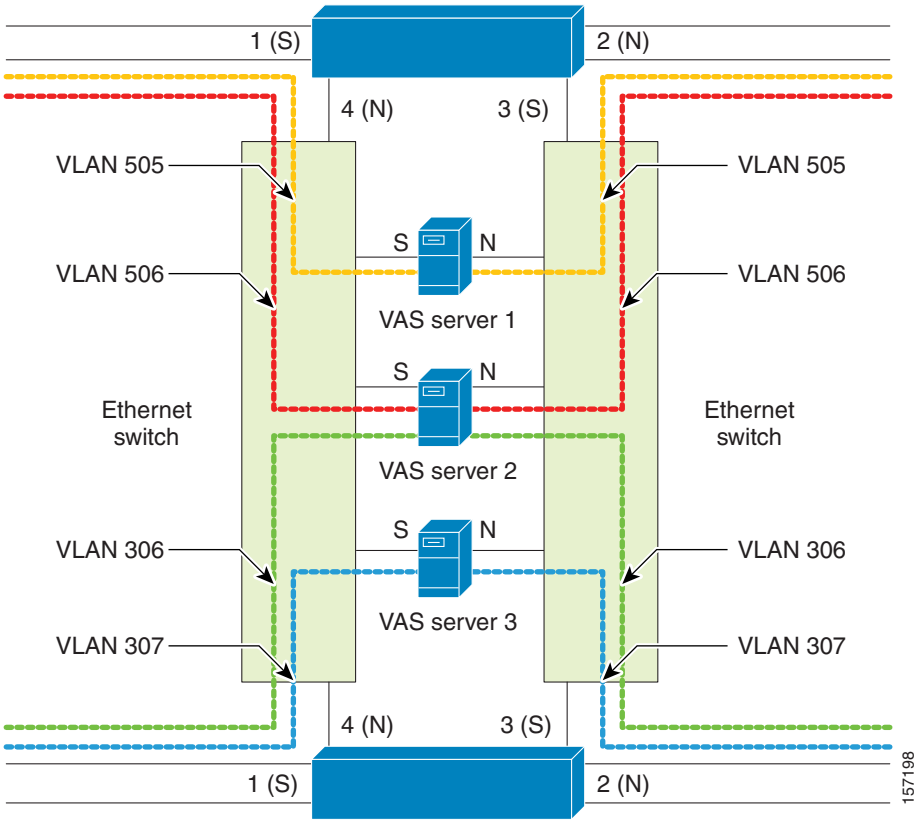
In Figure 14-4, the top Cisco SCE platform forwards traffic to VAS Server 1 and Server 2, while the bottom Cisco SCE platform forwards to VAS Server 2 and Server 3. A unique VLAN tag must designate each Cisco SCE-platform-to-VAS-server path. This topology is illustrated with two Cisco SCE platforms, but a maximum of 512 Cisco SCE platforms is supported (limited by the VLAN tag size).

The two Ethernet switches route the traffic to the VAS servers. The routing is VLAN based. The Ethernet switch should be configured to trunk mode with learning disabled.

The data flow is the same as that for the single Cisco SCE platform to multiple VAS servers topology (see "Data Flow" section on page 14-15).

Note The multiple Cisco SCE platforms to multiple VAS servers topology does not support Cisco SCE platform redundancy on the cascade ports.

Figure 14-4 Multiple Cisco SCE Platforms, Multiple VAS Servers



SNMP Support for VAS

The following items in the "PCUBE-SE-MIB" proprietary MIB support VAS traffic forwarding:

- Cisco SCE-MIB object—vasTrafficForwardingGrp Cisco SCE-MIB
- Object type—vasServersTable provides information on each VAS server operational status.
- SNMP Trap—vasServerOperationalStatusChangeTrap signifies that the agent entity has detected a change in the operational status of a VAS server.

Interactions Between VAS Traffic Forwarding and Other Cisco SCE Platform Features

- Incompatible Cisco SCE Platform Features, page 14-18
- VAS Traffic Forwarding and DDoS Processing, page 14-18
- VAS Traffic Forwarding and Bandwidth Management, page 14-19

Incompatible Cisco SCE Platform Features

There are certain Cisco SCE platform features that are incompatible with VAS traffic forwarding. Before enabling VAS traffic forwarding, it is the responsibility of the user to make sure that no incompatible features or modes are configured.

There are certain Cisco SCE platform features that are incompatible with VAS traffic forwarding. Before you enable VAS traffic forwarding, you must ensure that no incompatible features or modes are configured.

The features and modes listed below cannot coexist with VAS mode:

- Line-card connection modes—receive-only, receive-only-cascade, inline-cascade
- Link mode other than forwarding
- All link encapsulation protocols, including VLAN, MPLS, and L2TP
- Traffic mirroring (see "Intelligent Traffic Mirroring" section on page 14-36)

Note If VAS forwarding is enabled, Cisco SCE devices do not forward VLAN-tagged subscriber traffic received from subscriber side and network side traffic ports to the VAS interface for processing.

VAS Traffic Forwarding and DDoS Processing

VAS traffic forwarding has minor effects on the distributed denial of service (DDoS) mechanisms.

- Specific IP DDoS Attack Detection, page 14-18
- Specific IP Attack Filter, page 14-18

Specific IP DDoS Attack Detection

The specific IP DDoS mechanism uses software counters. The second pass VAS packets do not reach the Service Control Operating System (SCOS), so they are not counted twice.

Network-side packets are handled by the attack-detector in the first pass, when they open a flow, so they are also not counted twice.

Specific IP Attack Filter

The behavior depends on the action configured.

- Report only—VAS is not affected.

- Block—Flow is blocked, no VAS service is provided.
- Bypass—Traffic is bypassed and NO SCA BB or VAS services are provided.

VAS Traffic Forwarding and Bandwidth Management

The complexity of the VAS traffic forwarding results in the modification of some Cisco SCE platform bandwidth management capabilities:

- VAS flows are not subject to global bandwidth control.
- The number of global controllers available to regular flows is decreased from 64 to 48.

Global Controllers and VAS Flows

When VAS traffic forwarding is enabled, the global controllers function slightly differently.

- Only 48 global controllers are available.
- Global controllers 49 to 63 are used to count VAS traffic.
- Reserved global controllers cannot be configured.
- VAS flows do not get the global controller from the traffic controller to which they belong. Rather, the global controller is set according to VAS rules.

Configuring VAS Traffic Forwarding

There are three broad aspects to VAS traffic forwarding configuration in the Cisco SCE platform:

- Configuring global VAS traffic forwarding options, such as enabling or disabling VAS traffic forwarding, or specifying the VAS traffic link.
- Configuring a VAS server, such as enabling or disabling a specific VAS server, or enabling or disabling the VAS health check for a specified VAS server.
- Configuring a VAS server group, such as adding or removing a specific VAS server, configuring the minimum number of active servers per group, or configuring VAS server group failure behavior.

Note Additional VAS traffic forwarding configuration and monitoring options are available from the SCA BB Console. See "Managing VAS Settings" section in the *Cisco Service Control Application for Broadband User Guide*.

Following is a high-level description of the steps in configuring VAS traffic forwarding.

1. Configure the Cisco SCE platform— define the servers and the server groups, configure pseudo IP for the traffic interfaces used for VAS traffic, and enable VAS mode.
2. Verify the state of the individual VAS servers as well as that of the VAS server groups to make sure all are **Up** (see "Monitoring VAS Traffic Forwarding" section on page 14-32).
3. Configure which traffic goes to which server group using the SCA BB console (see "Configuring VAS Traffic Forwarding from the SCA BB Console" section on page 14-21).

This section contains the following information on configuring VAS traffic forwarding:

- Configuring VAS Traffic Forwarding from the SCA BB Console, page 14-21
- Global Options, page 14-21
- Enabling VAS Traffic Forwarding, page 14-21
- Disabling VAS Traffic Forwarding, page 14-22
- Configuring the VAS Traffic Link, page 14-22
- Configuring a VAS Server, page 14-23
- Assigning a VLAN ID to a VAS Server, page 14-24
- Configuring the Health Check, page 14-25
- Configuring a VAS Server Group, page 14-28
- VAS Configuration Example, page 14-31

Configuring VAS Traffic Forwarding from the SCA BB Console

Configuration of the VAS Traffic Forwarding solution is distributed between the SCA BB console and the Cisco SCE platform CLI:

- The Cisco SCE platform CLI configuration:
  - Physical VAS server parameters—VLAN tag, Admin status and health check parameters
  - VAS server groups parameters—the VAS servers that belong to the group and the action to take if the group enters a failure state
- SCA BB console configuration—the traffic forwarding rules, meaning which portion of the subscriber traffic should be forwarded to the VAS servers.

  This configuration is defined per package so different subscribers can receive different VAS service, based on the package they bought.

Global Options

There are two global VAS traffic forwarding options:

- Enable or disable VAS traffic forwarding
- Configure the link number on which to transmit VAS traffic (necessary only if the VAS servers are connected on Link 0, rather than Link 1, which is the default VAS traffic link))

Enabling VAS Traffic Forwarding

By default, VAS traffic forwarding is disabled. If VAS traffic forwarding is required, you must enable it both from your Cisco SCE device and the SCABB console.

For instructions on how to disable VAS traffic forwarding, see "Disabling VAS Traffic Forwarding" section on page 14-22.

There are certain other Cisco SCE platform features that are incompatible with VAS traffic forwarding. Before enabling VAS traffic forwarding, make sure that no incompatible features or modes are configured.

The features and modes listed below cannot coexist with VAS mode:

- Line-card connection modes—receive-only, receive-only-cascade, inline-cascade
- Link mode other than forwarding
- All link encapsulation protocols, including VLAN, MPLS, L2TP
- Enhanced open flow mode

Options

The following options are available:

- **Enable/disable—**Enable or disable VAS traffic forwarding
  - Default—Disable

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

Command	Purpose
VAS-traffic-forwarding	Enables VAS traffic forwarding.

Disabling VAS Traffic Forwarding

Disabling the VAS Traffic Forwarding feature in runtime must be done with special care. There are two points to consider:

- You cannot disable VAS mode in the Cisco SCE platform while the applied SCA BB policy instructs the Cisco SCE platform to forward traffic to the VAS servers.

  Therefore, you must dismiss all VAS Traffic forwarding rules in the applied SCA BB policy before disabling the VAS traffic forwarding in the Cisco SCE platform.

- After the SCA BB has been reconfigured, there may still be some open flows that have already been forwarded to the VAS servers. If the VAS feature is stopped while there are still such flows open, their packets coming back from the VAS servers may be routed to their original destination with the VLAN tag of the VAS server on it.

  Therefore, it is also highly recommended to shutdown the line card before you disable the VAS traffic forwarding in the Cisco SCE platform to avoid inconsistency with flows that were already forwarded to the VAS servers.

Step 1 From the SCA BB console, remove all the VAS table associations to packages and apply the changed policy.

Step 2 From the SCE(config if)# prompt, type **shutdown** and press **Enter**.

Shuts down the line card.

Step 3 From the SCE(config if)# prompt, type **no VAS-traffic-forwarding** and press **Enter**.

Disables VAS traffic forwarding.

Step 4 From the SCE(config if)# prompt, type **no shutdown** and press **Enter**.

Re-enables the line card.

Configuring the VAS Traffic Link

By default, the VAS traffic is transmitted on Link 1. If the VAS servers are connected on Link 0, you must configure the VAS traffic link to Link 0.

Note The VAS traffic link should be in Forwarding mode.

Options

The following option is available:

- **VAS traffic-link {link-0|link-1}—**The link number on which to transmit VAS traffic
  - Default—Link 1

How to Select the Link for VAS Traffic

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

Command	Purpose
VAS-traffic-forwarding traffic-link {link-0|link-1}	Selects the link for VAS traffic.

How to Revert to the Default Link for VAS Traffic

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

Command	Purpose
no VAS-traffic-forwarding traffic-link	Reverts the default link for VAS traffic.

Configuring a VAS Server

You must define the VAS servers. Each VAS server has the following parameters:

- Admin-mode—Enabled or disabled.
- Health check mode—Enabled or disabled
- Health check ports
- VLAN tag

Use the commands in this section to perform these operations for individual VAS servers:

- Enable a specified VAS server
- Disable a specified VAS server
- Define the VLAN tag for a specified VAS server
- Enable or disable the health check for a VAS server
- Define the source and destination ports to use for the health check.
- Delete all properties for a specified VAS server. The server returns to the default state, which is enabled. However, it is not operational since it does not have VLAN.

Note A VAS server is not operational until the VLAN tag is defined, even if the server itself is enabled.

Options

The following option is available:

- **number**—The number of the VAS server.

How to Enable a VAS Server

Note The server is not operational until a VLAN tag has also been defined

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

Command	Purpose
VAS-traffic-forwarding VAS server-id *number* **enable**	Enables VAS server.

How to Disable a VAS Server

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

Command	Purpose
VAS-traffic-forwarding VAS server-id *number* **disable**	Disables VAS server.

How to Restore all VAS Server Properties to Default

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

Command	Purpose
no VAS-traffic-forwarding VAS server-id *number*	Restores all VAS server properties to default.

Assigning a VLAN ID to a VAS Server

This section contains the following topics:

- How to Configure the VLAN Tag Number for a Specified VAS Server, page 14-25
- How to Remove the VLAN Tag Number from a Specified VAS Server, page 14-25

Options

The following options are available:

- **number**—The number of the VAS server.
- **vlan-id**—The VLAN tag to use for the specified VAS server

  The VLAN tag can be redefined as necessary.
  - Default—No VLAN.

Note the following important points:

- The VAS server is not operational until the VLAN tag is defined.
- Disabling the server does not remove the VLAN tag number configured to the server.
- The **no** form of the command (same as the **default** form of the command), removes the previously configured VLAN tag (no VLAN is the default configuration).

How to Configure the VLAN Tag Number for a Specified VAS Server

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

Command	Purpose
VAS-traffic-forwarding VAS server-id *number* **VLAN** *vlan-id*	Configures the VLAN tag number for a specified VAS server.

How to Remove the VLAN Tag Number from a Specified VAS Server

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

Commands	Purpose
no VAS-traffic-forwarding VAS server-id *number* **VLAN**	Removes the VLAN tag number from a specified VAS server.
default VAS-traffic-forwarding VAS server-id *number* **VLAN**	You can also use the default form of the command to remove the VLAN tag configuration.

Configuring the Health Check

This section explains how to to enable and disable the Health Check, and how to define the ports it should use.

By default, the VAS server health check is enabled, however you may disable it.

The health check will be activated only if all the following conditions are true. If the health check is enabled, the server state will be **Down** if one or more conditions are not met:

- VAS traffic forwarding mode is enabled.
- Pseudo IPs are configured for the Cisco SCE platform traffic ports on the VAS traffic link.
- VAS server is enabled.
- Server has a VLAN tag.
- Health check for the server is enabled.

If the health check of the server is disabled, its operational status depends on the following (requirements for **Up** state are in parentheses):

- admin status (enable)
- VLAN tag configuration (VLAN tag defined)
- group mapping (assigned to group)

This section contains the following topics:

- How to Enable VAS Server Health Check, page 14-26
- How to Disable VAS Server Health Check, page 14-27
- How to Define the UDP Ports to be Used for Health Check, page 14-27
- How to Remove the UDP Ports Configuration, page 14-27
- Configuring Pseudo IP Addresses for the Health Check Packets, page 14-27

Options

The following options are available:

- **number—**The ID number of the VAS server for which to enable or disable the health check
- **Enable/disable—**Enable or disable VAS server health check
  - Default—Enable
- **UDP ports—**Specify the UDP ports to be used for the health check:
  - **source portnumber—**Health check source port number
  - **destination portnumber—**Health check destination port number
  - Default—<63140,63141>used for server #0 through <63154,63155>used for server #7.

How to Enable VAS Server Health Check

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

Command	Purpose
VAS-traffic-forwarding VAS server-id *number* **health-check**	Enables VAS server health check.

How to Disable VAS Server Health Check

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

Command	Purpose
no VAS-traffic-forwarding VAS server-id *number* **health-check**	Disables the VAS server health check.

How to Define the UDP Ports to be Used for Health Check

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

Command	Purpose
VAS-traffic-forwarding VAS server-id *number* **health-check UDP ports source** *portnumber* **destination** *portnumber*	Defines the UDP ports to be used for health check.

How to Remove the UDP Ports Configuration

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

Commands	Purpose
no VAS-traffic-forwarding VAS server-id *number* **health-check UDP ports**	Removes UDP ports configuration.
VAS-traffic-forwarding VAS server-id *number* **health-check UDP ports**	You can also use the default form of the command to remove the UDP port configuration.

Configuring Pseudo IP Addresses for the Health Check Packets

You should configure source and destination pseudo IP address for the health check packets. The **pseudo-ip** command allows you to specify a unique IP address to be used by the health check packets.

The pseudo IP address is configured on the interfaces that connect the Cisco SCE platform with the VAS servers.

The Cisco SCE platform uses the pseudo IP as follows:

- Pseudo IP configured for the subscriber side interface:
  - Source IP address for health check packets going in the Upstream direction
  - Destination IP address for health check packets going in the Downstream direction
- Pseudo IP configured for the network side interface:
  - Source IP address for health check packets going in the Downstream direction
  - Destination IP address for health check packets going in the Upstream direction

Note This command is a ROOT level command in the Gigabit Interface Configuration mode.

Options

The following options are available:

- **ip-address**—IP address to be used (any IP address as long as it is not possible to be found in the network traffic, such as a private IP)
  - Default—no IP address
- **mask** (optional)—Defines the range of IP addresses that can be used by the Cisco SCE platform. Note that the Cisco SCE platform is not required to reside in this subnet.
  - Default—255.255.255.255 (The subnet mask can be set to 255.255.255.255, as the health check mechanism requires only one IP address per interface.)

How to Define the Pseudo IP Address

Use this command to define the pseudo IP address to be used for the health check.

From the SCE(config if)#>prompt, type:

Command	Purpose
pseudo-ip *ip-address [mask]*	Defines the pseudo IP address.

How to Delete the Pseudo IP Address

From the SCE(config if)#>prompt, type:

Command	Purpose
no pseudo-ip *ip-address [mask]*	Deletes the pseudo IP address.

Configuring a VAS Server Group

You may define up to eight VAS server groups. Each VAS server group has the following parameters:

- Server Group ID
- A list of VAS servers attached to this group.
- Failure detection—minimum number of active servers required for this group so it will be considered to be Active. If the number of active servers goes below this minimum, the group will be in Failure state.
- Failure action—action performed on all new data flows that should be mapped to this Server Group while it is in Failure state.

  Options:
  - block
  - pass

You can perform these operations for a VAS server group:

- Add a VAS server to or remove a VAS server from a specified group.
- Configure the minimum number of active servers for a specified group.
- Configure failure behavior for a specified group.

Adding and Removing Servers

Options

The following options are available:

- **group-number—**The ID number of the VAS server group
- **id-number—**The ID number of the VAS server

How to Add a VAS Server to a Specified VAS Server Group

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

Command	Purpose
VAS-traffic-forwarding VAS server-group *group-number* **server-id** *id-number*	Adds a VAS server to a specified VAS server group.

How to Remove a VAS Server from a Specified VAS Server Group

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

Command	Purpose
no VAS-traffic-forwarding VAS server-group *group-number* **server-id** *id-number*	Removes a VAS server from a specified VAS server group.

How to Remove all VAS Servers from a Specified VAS Server Group

Use this command to remove all VAS servers from a specified VAS server group and set all group parameters to the default value.

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

Command	Purpose
no VAS-traffic-forwarding VAS server-group *group-number*	Removes all VAS servers from a specified VAS server group.

Configuring VAS Server Group Failure Parameters

You can configure these failure parameters for a specified VAS server group:

- Minimum number of active servers—If the number of active servers in the server group goes below this number, the group will be in Failure state
- Failure action—The action to be applied to all new flows mapped to this server group while it is Failure state:
  - Block—all new flows assigned to the failed VAS server group will be blocked by the Cisco SCE platform.
  - Pass—all new flows assigned to the failed VAS server group will be considered as regular non-VAS flows, and will be processed without VAS service.

Options

The following options are available:

- **group-number**—The ID number of the VAS server group
- minimum-active-servers **min-number**—The minimum number of active servers required for the specified server group
  - Default—1
- **failure action**—Which of the following actions will be applied to all new flows for the specified server group:
  - **block**
  - **pass** (default)

How to Configure the Minimum Number of Active Servers for a Specified VAS Server Group

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

Command	Purpose
VAS-traffic-forwarding VAS server-group *group-number* **failure minimum-active-servers** *min-number*	Configures the minimum number of active servers for a specified VAS server group.

How to Reset the Minimum Number of Active Servers for a Specified VAS Server Group to the Default

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

Command	Purpose
default VAS-traffic-forwarding VAS server-group *group-number* **failure minimum-active-servers** *min-number*	Resets the minimum number of active servers for a specified VAS server group to the default.

How to Configure the Failure Action for a Specified VAS Server Group

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

Command	Purpose
VAS-traffic-forwarding VAS server-group *group-number* **failure action {block | pass}**	Configures the failure action for a specified VAS server group.

How to Configure the Failure Action for a Specified VAS Server Group to the Default

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

Command	Purpose
default VAS-traffic-forwarding VAS server-group *group-number* **failure action**	Configures the failure action for a specified VAS server group to the default.

VAS Configuration Example

	Command	Purpose
Step 1	enable 15	Access root level to configure the pseudo IP address
Step 2	**configure** **interface range TenGigabitEthernet 3/0-1/0**	Enter Ten Gigabit Ethernet Interface configuration mode for the relevant range of 10GBE interfaces.
Step 3	**pseudo-ip** *1.1.1.1 255.255.255.252*	Configure the pseudo IP address for the health check.
Step 4	exit **interface linecard 0**	Enter Interface Linecard configuration mode
Step 5	**shutdown**	You must shutdown the linecard when configuring VAS servers and groups.
Step 6	**VAS-traffic-forwarding**	Set the Cisco SCE platform to forward VAS traffic (enable VAS traffic forwarding).
Step 7	**VAS-traffic-forwarding traffic-link** *link-0*	Set the VAS traffic forwarding link to Link 0.
Step 8	**VAS-traffic-forwarding VAS server-id** *0* **VLAN** *600* **VAS-traffic-forwarding VAS server-id** *1* **VLAN** *601* **VAS-traffic-forwarding VAS server-id** *2* **VLAN** *602*	Assign VAS servers 0-2 to VLAN 600- 602 respectively.
Step 9	**VAS-traffic-forwarding VAS server-group** *0* **server-id** *0* **VAS-traffic-forwarding VAS server-group** *0* **server-id** *1* **VAS-traffic-forwarding VAS server-group** *0* **server-id** *2*	Map VAS servers to server group 0, allowing server redundancy within the group.
Step 10	**VAS-traffic-forwarding VAS server-id** *0* **health-check UDP ports source** *63154* **destination** *63155* **VAS-traffic-forwarding VAS server-id** *1* **health-check UDP ports source** *63156* **destination** *63157* **VAS-traffic-forwarding VAS server-id** *2* **health-check UDP ports source** *63158* **destination** *63159*	Define UDP ports for health check on VAS servers.
Step 11	**VAS-traffic-forwarding VAS server-group** *0* **failure minimum-active-servers** *2*	Configure the minimum number of servers required (2).
Step 12	**VAS-traffic-forwarding VAS server-group** *0* **failure action block**	Configure the failure action to ‘block’.
Step 13	**no shutdown**	Restart the linecard.

Monitoring VAS Traffic Forwarding

Use the commands in this section to display the following information for VAS configuration and operational status summary.

- Global VAS status summary—VAS mode, the traffic link used
- VAS Server Groups information summary—operational status, number of configured servers, number of current active servers.

This information may be displayed for a specific server group or all server groups

- VAS servers information summary—operational status, Health Check operational status, number of subscribers attached to this server.

This information may be displayed for a specific server or all servers

- VAS health check counters

The following sample outputs are included.

- How to Display Global VAS Status and Configuration, page 14-32
- How to Display Operational and Configuration Information for a Specific VAS Server Group, page 14-33
- How to Display Operational and Configuration Information for All VAS Server Groups, page 14-33
- How to Display Operational and Configuration Information for a Specific VAS Server, page 14-33
- How to Display Operational and Configuration Information for All VAS Servers, page 14-34
- How to Display the VAS Servers Used by a Specified Subscriber, page 14-34
- How to Display Health Check Counters for All VAS Servers, page 14-35
- How to Clear the Health Check Counters for a Specified VAS Server, page 14-35
- How to Clear the Health Check Counters for All VAS Servers, page 14-35

How to Display Global VAS Status and Configuration

From the SCE> prompt, type:

Command	Purpose
show interface linecard 0 VAS-traffic-forwarding	Displays the global VAS status and configuration.

Example

```
SCE>show interface linecard 0 VAS-traffic-forwarding
VAS traffic forwarding is enabled
VAS traffic link configured: Link-1  actual: Link-1
```

How to Display Operational and Configuration Information for a Specific VAS Server Group

From the SCE> prompt, type:

Command	Purpose
show interface linecard 0 VAS-traffic-forwarding VAS server-group *id-number*	Displays operational and configuration information for a specific VAS server group.

Example

```
SCE>show interface linecard 0 VAS-traffic-forwarding VAS server-group 0
VAS server group 0:
State: Failure  configured servers: 0   active servers: 0
minimum active servers required for Active state: 1   failure action: Pass
```

How to Display Operational and Configuration Information for All VAS Server Groups

From the SCE> prompt, type:

Command	Purpose
show interface linecard 0 VAS-traffic-forwarding VAS server-group all	Displays operational and configuration information for all VAS server groups.

How to Display Operational and Configuration Information for a Specific VAS Server

From the SCE> prompt, type:

Command	Purpose
show interface linecard 0 VAS-traffic-forwarding VAS server-id *id-number*	Displays operational and configuration information for a specific VAS server.

Example

```
SCE>show interface linecard 0 VAS-traffic-forwarding VAS server-id 0
VAS server 0:
Configured mode: enable   actual mode: enable   VLAN:  520  server group:    3
State: UP
Health Check configured mode: enable  status: running
Health Check source port: 63140  destination port: 63141
Number of subscribers:       0
```

How to Display Operational and Configuration Information for All VAS Servers

From the SCE> prompt, type:

Command	Purpose
show interface linecard 0 VAS-traffic-forwarding VAS server-id all	Displays operational and configuration information for all VAS servers.

How to Display the VAS Servers Used by a Specified Subscriber

From the SCE> prompt, type:

Command	Purpose
show interface linecard 0 subscriber name *subscriber-name* **VAS-servers**	Displays the VAS servers used by a specified subscriber.

How to Display Health Check Counters for a Specified VAS Server

From the SCE> prompt, type:

Command	Purpose
show interface linecard 0 VAS-traffic-forwarding VAS server-id *id-number* **counters health-check**	Displays health check counters for a specified VAS server.

Example

```
SCE>show interface linecard 0 VAS-traffic-forwarding VAS server-id 0
Health Checks statistics for VAS server '0'    Upstream      Downstream
-------------------------------------------------------------------------
Flow Index '0'
-----------------
Total packets sent                          :       31028  :       31027 :
Total packets received                      :       31028  :       31027 :
Good packets received                       :       31028  :       31027 :
Error packets received                      :           0  :           0 :
Not handled packets                         :           0  :           0 :
Average roundtrip (in millisecond)          :           0  :           0 :
Error packets details                       :              :             :
--------------------                        :              :             :
Reordered packets                           :           0  :           0 :
Bad Length packets                          :           0  :           0 :
IP Checksum error packets                   :           0  :           0 :
L4 Checksum error packets                   :           0  :           0 :
L7 Checksum error packets                   :           0  :           0 :
Bad VLAN tag packets                        :           0  :           0 :
Bad Device ID packets                       :           0  :           0 :
Bad Server ID packets                       :           0  :           0 :
```

How to Display Health Check Counters for All VAS Servers

From the SCE> prompt, type:

Command	Purpose
show interface linecard 0 VAS-traffic-forwarding VAS server-id all counters health-check	Displays health check counters for all VAS servers.

How to Clear the Health Check Counters for a Specified VAS Server

From the SCE> prompt, type:

Command	Purpose
clear interface linecard 0 VAS-traffic-forwarding VAS server-id *id-number* **counters health-check**	Clears the health check counters for a specified VAS server.

How to Clear the Health Check Counters for All VAS Servers

From the SCE> prompt, type:

Command	Purpose
clear interface linecard 0 VAS-traffic-forwarding VAS server-id all counters health-check	Clears health check counters for all VAS servers.

Intelligent Traffic Mirroring

- How Traffic Mirroring Works, page 14-37
- Configuring Traffic Mirroring, page 14-41
- Monitoring Traffic Mirroring, page 14-42
- Traffic Mirroring Sample Configuration, page 14-42

Traffic mirroring is a Cisco SCE platform capability that complements the range of services provided by the SCA BB solution. It copies a specified portion of the traffic streams and sends this copy to third-party servers who do offline analysis.

The criteria for traffic mirroring is based on Layer 7 attributes and subscriber awareness. This fine granularity, along with load sharing capability for servers providing the same service, substantially reduces the number of solution components.

The traffic that is copied is also processed by the SCA BB application and forwarded without interruption to its original destination. The copy of the traffic is presumed not to return to the Cisco SCE platform after being processed by the third party servers.

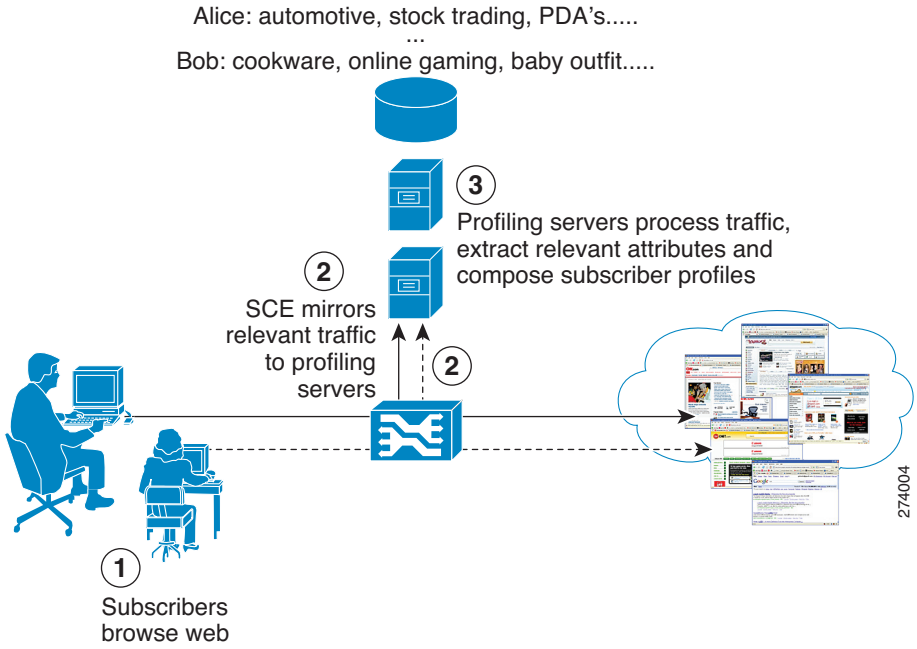
Using Traffic Mirroring for Behavioral Targeting

Today Internet advertising is being executed by content providers (or publishers) in collaboration with ad networks, which actually handle the syndication of ads from advertisers to web sites. The Cisco Service Control behavioral targeting solution provides the means for service providers to participate in the business of the online advertising. This solution allows the SPs to leverage their information about the subscribers and enables highly targeted advertising.

Based on Deep Packet Inspection (DPI) and subscriber integration, the Cisco SCE 8000 platform filters out only what is relevant for the subscriber profiling. This greatly conserves the resources of the advertising servers, by eliminating the irrelevant web traffic before it even reaches them. This filtered traffic is processed as usual by the SCA BB application along with the rest of the traffic, but in addition a copy of it is passed to an external device which can do offline analysis on the subscriber behavior. The data can then be used for the targeted advertising.

Behavioral Targeting is accomplished using several Cisco SCE platform capabilities. This section describes the intelligent traffic mirroring capability, which is one of the features that enable the solution.

Figure 14-5 High Level Overview of an Mirroring-based Behavioral Targeting Solution



For more information regarding targeted advertising, see the following documents:

- *Cisco Service Control Online Advertising Solution Guide: Behavioral Profile Creation Using RDRs*
- *Cisco Service Control Online Advertising Solution Guide: Behavioral Profile Creation Using Traffic Mirroring*

How Traffic Mirroring Works

Traffic mirroring uses the same capabilities as VAS traffic forwarding, but the traffic does not return to the Cisco SCE platform.

The following sections explain how traffic mirroring works:

- Traffic Mirroring and SCA BB, page 14-37
- Mirroring Termination, page 14-38
- Mirroring Exceptions, page 14-38
- Cisco SCE Connectivity, page 14-39
- Traffic Mirroring and Bandwidth Management, page 14-41

Traffic Mirroring and SCA BB

When traffic mirroring is configured for a certain type of traffic, in addition to all its basic functions, the SCA BB application decides whether each flow is to be mirrored or not, based on L7 classification.

Traffic mirroring rules are configured through the SCA BB console. These rules map the traffic to be mirrored and analyzed to the VAS server groups. When a flow is marked for traffic mirroring, the VAS server group for this flow is selected. If the group includes more than one VAS server, traffic will be forwarded in such a way that the subscriber load is shared between the servers on the same group.

The mapping of traffic portions to VAS server groups is done through the standard SCA BB GUI, this definition is given per package.

Mirroring Termination

Mirroring of a flow can continue until the flow is terminated, or be limited to a certain volume passed over the flow. This enables a huge data reduction on the server side, as well as performance saving in the Cisco SCE platform.

An RST packet is sent to the server when the mirroring is stopped due to a stop condition. This is done in order to signal the server that the mirroring has stopped.

The RST packet is sent in the direction of intiator to initatee with the additional VLAN tag.

Mirroring Exceptions

Since the decision to mirror is based on service classification, which can be done on the first payload or after first few packets, the entire TCP handshake is not mirrored.

To save in performance on both sides, zero payload packets are also not mirrored. (note that this type of packets have no real value for offline analysis).

If the VLAN traffic is mirrored, Cisco SCE devices replace the VLAN information from the incoming traffic with the VAS-configured VLAN information before mirroring the traffic on the VAS port.

Mirroring the TCP-Segmented HTTP GET Packets

From Cisco SCE Release 3.7.5, traffic mirroring can be enabled for all the TCP-segmented HTTP GET packets if the HTTP port is 80, 8080, or 8081 even if the number of segmented packets are more than two.

By default, mirroring of all the segmented packets is disabled.

Configuring Cisco SCE to Mirror the TCP-Segmented HTTP GET packets

You can configure Cisco SCE to mirror the TCP-segmented HTTP GET packets using the GT_SEG_GET_MIRROR tunable.

To configure Cisco SCE to mirror the TCP segmented HTTP GET packets, from the SCE(config if)# prompt, enter:

Command	Purpose
tunable GT_SEG_GET_MIRROR value true	Enables the mirroring of all the TCP-segmented HTTP GET packets.
tunable GT_SEG_GET_MIRROR value false	Disables the mirroring of all the TCP-segmented HTTP GET packets.

Cisco SCE Connectivity

Traffic mirroring is implemented by sending the mirrored packets over a designated VLAN through a predefined link of the Cisco SCE platform. The link that has been defined for traffic mirroring can be either used exclusively for this purpose, or it can be one of the traffic ports, in which case the Tx capacity of the link will be shared between the original egress traffic and the mirrored traffic.

The direction of the flow is preserved when mirrored, so traffic that is received on the subscriber interface on either link is sent over a VLAN on the network interface over this predefined link. And traffic that is received on the network interface on either link is sent over a VLAN on the subscriber interface over this predefined link. The mirrored traffic does not return to the Cisco SCE platform.

Note Enabling traffic mirroring is expected to impact the Cisco SCE performance due to the excessive processing associated with it; the actual figure depending on the amount of the mirrored traffic. It is recommended that you monitor Cisco SCE platform performance when enabling this capability.

Figure 14-6 shows a Cisco SCE platform using a dedicated link for mirroring (Link 1).

Figure 14-6 Traffic Mirroring on a Dedicated Link

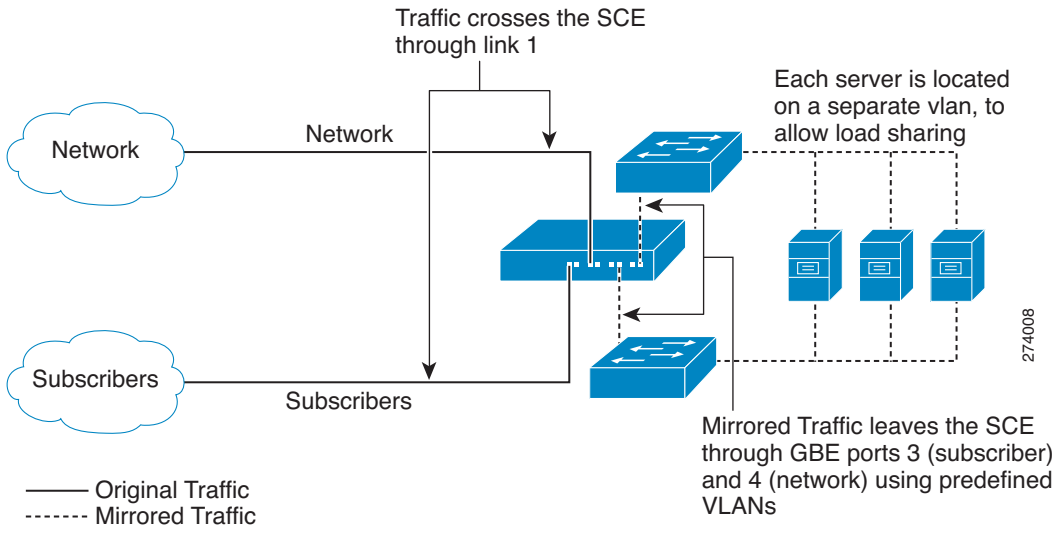
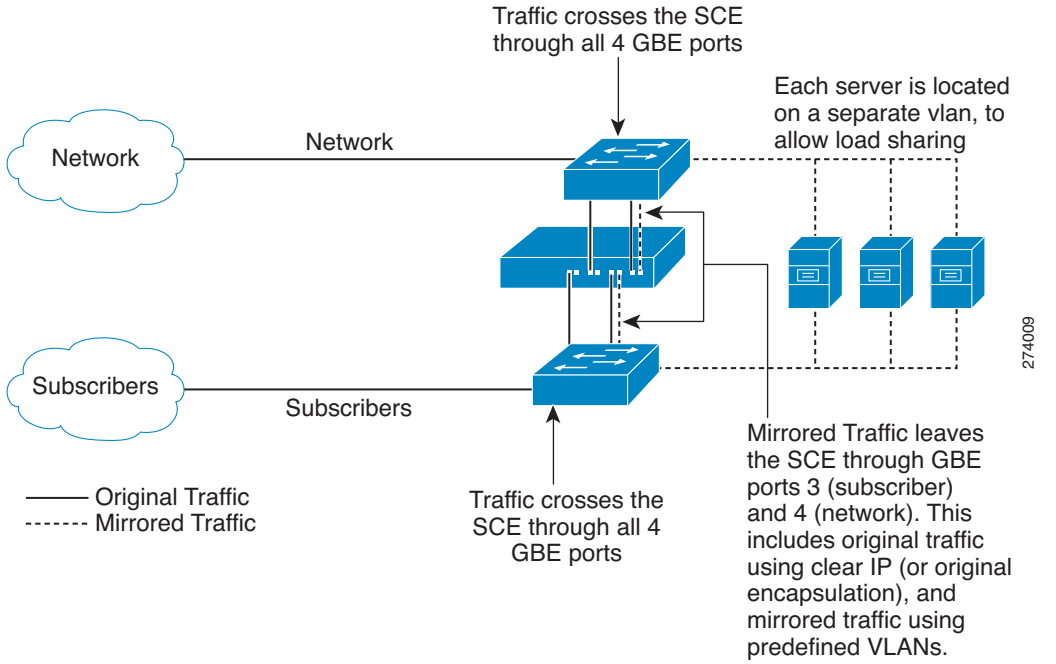


Figure 14-7 shows a Cisco SCE platform using traffic ports for mirroring.

Figure 14-7 Traffic Mirroring over Traffic Ports



Traffic Mirroring and Bandwidth Management

In general, the bandwidth management rules to be applied to flows designated for mirroring are not affected by the mirroring decision. However note the following points:

- Bandwidth enforcement applies on the two copies of the flow as if there was only one copy; that is, if the flow should be limited to 50K, then its original copy which is sent to the original destination is limited to 50K, and the copy that is sent to the VAS server is also limited to 50K. This has no effect on the total volume that should be mirrored.
- The mirrored volume is not counted and accounted in the BWM system, and therefore in cases where the mirrored traffic is congesting the link, the system will not be aware of the link congestion and will not know to shrink the BWC.

Configuring Traffic Mirroring

Following is a high-level description of the steps in configuring traffic mirroring.

1. Configure the Cisco SCE platform— define the servers and the server groups
2. Configure which traffic goes to which server group using the SCA BB console.

Note Additional traffic mirroring configuration and monitoring options are available from the SCA BB Console. See the *Cisco Service Control Application for Broadband User Guide*.

Note Traffic mirroring is not compatible with regular VAS traffic forwarding.

Traffic mirroring configuration is distributed between the SCA BB console and the Cisco SCE platform CLI:

- The Cisco SCE platform CLI configuration:

  There are three broad aspects to traffic mirroring configuration in the Cisco SCE platform:

  - Configuring the VAS traffic link.
  - Configuring the VLAN tag per VAS server.
  - Associating servers with server groups.

  The health check is not relevant to traffic mirroring, so there is no need to configure anything related to the VAS health check.

- SCA BB console configuration—Configure the traffic mirroring rules, specifying which portion of the subscriber traffic should be mirrored to the VAS servers.

  This configuration is defined per package so different subscribers can receive different mirroring service, based on the package they bought.

Monitoring Traffic Mirroring

Use the same commands to monitor traffic mirroring as for regular VAS functionality. (See "Monitoring VAS Traffic Forwarding" section on page 14-32)

Traffic Mirroring Sample Configuration

Following is a sample illustrating the steps in configuring the Cisco SCE 8000 platform for traffic mirroring.

	Command	Purpose
Step 1	**configure** **interface LineCard 0**	Enters LineCard Interface configuration mode
Step 2	**VAS-traffic-forwarding VAS server-id** *0* **VLAN** *640* **VAS-traffic-forwarding VAS server-id** *1* **VLAN** *641* **VAS-traffic-forwarding VAS server-id** *2* **VLAN** *642* **VAS-traffic-forwarding VAS server-id** *3* **VLAN** *643*	Assign VAS servers 0-3 to VLAN 640-643 respectively.
Step 3	**VAS-traffic-forwarding VAS server-group** *0* **server-id** *0* **VAS-traffic-forwarding VAS server-group** *0* **server-id** *1* **VAS-traffic-forwarding VAS server-group** *1* **server-id** *2* **VAS-traffic-forwarding VAS server-group** *1* ***server-id*** 3	Map VAS servers 0-1 and 2-3 to server groups 0 and 1 respectively, allowing server redundancy within each group.



Cisco Service Control MIBs

Revised: September 23, 2015

Introduction

The Cisco SCE platform originally supported the standard MIB-II and a proprietary Service Control Enterprise MIB. The proprietary *pcube* MIB enabled the external management system to perform configuration, performance, troubleshooting and alerting operations specific to the Cisco SCE platform that were not provided by the standard MIB.

The proprietary pcube MIBs has been replaced by a combination of standard and Cisco MIBs and new Cisco Service Control MIBs. The new MIB structure was designed to keep backward compatibility and provide the same information as provided in the past as much as possible.

This appendix explains how to map the proprietary pcube MIB supported in previous releases to the new MIB structure. It points out backward compatible issues and provides mapping guidelines from old MIB or OID group to a new MIB.

Note These MIB updates are supported on the Cisco SCE 8000 platform only. The pcube MIB is backward compatible to 3.1.6 for the Cisco SCE 1000 and Cisco SCE 2000 platforms.

The list of supported MIBs is described at the following URL under the Cisco Service Routing Products section. Select the desired product and then the desired release.

http://www.cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

All Cisco MIBs and some of the common standard MIBs can be obtained at the following URL:

http://tools.cisco.com/Support/SNMP/do/BrowseMIB.do?local=en&step=2

MIB Files

The pcube MIB was grouped into several MIBs, each of which represented a certain aspect or functionality in the Cisco SCE platform (see the tables in the "pcube to Cisco MIB Mapping: Detailed OID Mappings" section on page A-7 section for more details).

The pcube MIB files have been replaced by MIBs from three sources:

- standard MIBs
- existing Cisco MIBs
- new Cisco Service Control MIBs

Table A-1, Table A-2, and Table A-3 describe the former pcube MIB files and the MIB files that replace the pcube MIB.

Table A-1 pcube MIBs

MIB	Description
PCUBE-SMI.my	Defines P-cube enterprise tree structure
PCUBE-PRODUCTS-MIB.my	Defines OIDs of Cisco Service Control products
PCUBE-CONFIG-COPY-MIB.my	Contains a subset of the Cisco Config-Copy-MIB ported to the pcube enterprise subtree
CISCO-SCAS-BB-MIB.my	Contains SCA BB information handlers
PCUBE-SE-MIB.my	Contains information about the Cisco SCE platform

Table A-2 Newly Added CISCO-SERVICE_CONTROL MIBS That Replace pcube MIBs

MIB	Description
CISCO-SERVICE-CONTROL-LINK-MIB.my	Provides information about the status and configuration of links used by service control entities.
CISCO-SERVICE-CONTROL-RDR-MIB.my	Defines objects describing statistics and configuration relating to the Raw Data Record Formatter running on a service control entity.
CISCO-SERVICE-CONTROL-SUBSCRIBERS-MIB.my	Provides global and specific information on subscribers managed by a service control entity
CISCO-SERVICE-CONTROL-TP-STATS-MIB.my	Provides information and statistics on the traffic processor(s) of a service control entity.
CISCO-SERVICE-CONTROL-ATTACK-MIB.my	Provides information related to different types of attacks detected on the network entities and corresponding actions taken in the monitored network
CISCO-SERVICE-CONTROLLER-MIB.my	Provides information about service control traffic controllers.

Table A-3 Standard and Cisco MIBs Used to Replace pcube MIBs

MIBs	Description
CISCO-SMI.my	Defines Cisco enterprise tree structure
CISCO-TC.my	Contains Textual Conventions needed in some MIBs
CISCO-CONFIG-COPY-MIB.my	Facilitates writing of configuration files
CISCO-ENTITY-ALARM-MIB.my	Defines the managed objects that support the monitoring of alarms generated by physical entities contained by the system
CISCO-ENTITY-FRU-CONTROL-MIB.my	Monitor s and configures operational status of Field Replaceable Units
CISCO-ENTITY-REDUNDANCY-MIB.my	Supports configuration, control and monitoring of redundancy protection for various kinds of components on Cisco managed devices.
CISCO-ENTITY-REDUNDANCY-TC-MIB.my	Defines the textual conventions used within Cisco Entity Redundancy MIBs
CISCO-ENTITY-SENSOR-MIB.my	Monitors the values of sensors in the Entity-MIB
CISCO-PROCESS-MIB.my	Provides overall information about the CPU load.
CISCO-QUEUE-MIB.my	Manages interface queuing in Cisco devices.
CISCO-SECURE-SHELL-MIB.my	Displays and configures accounting and Secure Shell (SSH) related features in a device.
CISCO-SYSLOG-EXT-MIB.my	Extends the Cisco Syslog MIB and provides network management support to handle and process Syslog messages as device events
CISCO-SYSLOG-EVENT-EXT-MIB.my	Extends the Cisco Syslog.MIB and provides network management support to handle and process Syslog messages as device events
CISCO-SYSLOG-MIB.my	Describes and stores the system messages generated by the IOS and any other OS which supports syslogs.
CISCO-TELNET-SERVER-MIB.my	Displays and configures Telnet related features in a device.
ENTITY-MIB.my	Represents multiple logical entities supported by a single SNMP agent
ENTITY-STATE-MIB.my	Defines a state extension to the Entity MIB
ENTITY-STATE-TC-MIB.my	Defines state textual conventions.
HOST-RESOURCES-MIB.my	Manages host systems.



Note Only OIDs that are mapped to former pcube MIB OIDs are in use in the standard and Cisco MIBs as listed in this table.

Loading MIBs

It is important to load the MIBs in the proper order.

Before loading any new CISCO-SERVICE-CONTROL MIB, load the following MIBs in this order:

1. SNMPv2-SMI.my
2. SNMPv2-CONF.my
3. SNMPv2-TC.my
4. SNMP-FRAMEWORK-MIB.my
5. ENTITY-MIB.my
6. INET-ADDRESS-MIB.my
7. CISCO-SMI.my
8. CISCO-TC.my

Loading procedure for standard MIBs and other legacy Cisco MIBs is explained here:

http://tools.cisco.com/Support/SNMP/do/BrowseMIB.do?local=en&step=2

pcube to Cisco MIB Mapping

This section is an overview of how the former pcube MIB maps to the current Cisco MIBs. Two P-cube MIBs are mapped; PcubeSeMIB and PcubeEngageMIB (CISCO-SCABB-MIB). Table A-4 lists the overall tree structure and Table A-5 lists the PcubeSeMIB.

Table A-4 Overall tree Structure: Related Objects

Pcube Object Name	New MIB Object Name/MIB Name
pcube	CISCO-SERVICE-CONTROL or cServiceControl
pcubeProducts	ciscoMgmt
pcubeModules	Removed
pcubeSeMIB	See Table A-5 on page A-5
pcubeEngageMIB	See "Pcube Engage MIB (CISCO-SCAS-BB-MIB)" section on page A-6

Table A-5 PcubeSeMIB

PcubeSEMIB Group Name	New MIB Object Name/MIB Name
pcubeSeMIBpcubeModuleGroup	cServiceControlMIB
pcubeSeConformance	cServiceControlNotifs
pcubeSeGroups	cServiceControlObjects
pcubeSystemGroup	ENTITY and ENTITY-EXTENTION MIBs
pcubeChassisGroup	ENTITY and ENTITY-EXTENTION MIBs
pcuebModuleGroup	ENTITY and ENTITY-EXTENTION MIBs
pcubeLinkGroup	CISCO-SERVICE-CONTROL-LINK-MIB
pcubeDiskGroup	HOST-RESOURCES-MIB
pcubeRdrFormatterGroup	CISCO-SERVICE-CONTROL-RDR-MIB
pcubeLoggerGroup	clogHistoryTable from CISCO-SYSLOG-MIB
pcubeSubscribersGroup	CISCO-SERVICE-CONTROL-SUBSCRIBERS-MIB
pcubeTrafficProcessorGroup	CISCO-SERVICE-CONTROL-TP-STATS-MIB
pcubePortGroup	ENTITY-MIB, many objects in the group are no longer necessary
pcubeTxQueuesGroup	CISCO-QUEUE-MIB with a few modifications.
pcubeGlobalControllersGroup	CISCO-SERVICE-CONTROLLER-MIB
pcubeTrafficCountersGroup	CISCO-SERVICE-CONTROL-TP-STATS-MIB
pcubeAttackGroup	CISCO-SERVICE-CONTROL-ATTACK-MIB
pcubeTrapObjectsGroup	Notifications have been mapped to the relevant MIBs

Pcube Engage MIB (CISCO-SCAS-BB-MIB)

The information in the pcubeEnageMIB is available from various RDRs and from tables of the Collection Manager database. Therefore this MIB has not been replaced by a new Cisco Service Control MIB.

For information regarding the mapping of the MIB objects to RDRs and the Collection Manager database, see Table A-21 on page A-24.

pcube to Cisco MIB Mapping: Detailed OID Mappings

The following tables provide the detailed mappings for specific pcubeSeMIB (1.3.6.1.4.1.5655.4.1/0) OIDs to the current standard and Cisco MIBs.

Table A-6 systemGrp (1.3.6.1.4.1.5655.4.1.1)

pcube Object Name	OID	New MIB	New Object Name	OID
sysOperationalStatus	1.3.6.1.4.1.5655.4.1.1.1	ENTITY-STATE-MIB	entStateTable.entStateOper entStateTable.entStateAlarm	1.3.6.1.2.1.131.1.1.1.3 1.3.6.1.2.1.131.1.1.1.5
other(1)			entStateTable.entStateOper = unknown (1) entStateTable.entStateAlarm = indeterminate (2).	
boot(2)			entStateTable.entStateOper = testing (4) entStateTable.entStateAlarm =unknown (0x80)	
operational(3)			entStateTable.entStateOper = enabled (3) entStateTable.entStateAlarm = unknown(0x80)	
warning(4)			entStateTable.entStateOper = enabled (3) entStateTable.entStateAlarm = warning (5)	
failure(5)			entStateTable.entStateOper = disabled (2) entStateTable.entStateAlarm = critical(2) or major(3)	
sysFailureRecovery	1.3.6.1.4.1.5655.4.1.1.2	CISCO-ENTITY-EXT-MIB	not mapped	
sysVersion	1.3.6.1.4.1.5655.4.1.1.2.1	ENTITY-MIB	entPhysicalDescr	1.3.6.1.2.1.47.1.1.1.1.2

Table A-7 pchassisGrp (1.3.6.1.4.1.5655.4.1.2

pcube Object Name	OID	New MIB	New Object Name	OID
pchassisSysType	1.3.6.1.4.1.5655.4.1.2.1		Not mapped. Derived from entPhysicalDescr and entPhysicalClass chassis(3)	
pchassisPowerSupply Alarm	1.3.6.1.4.1.5655.4.1.2.2	CISCO-ENTITY-FRU-CONTROL-MIB	Trap is sent Current status available using the **show environment** CLI command	
pchassisFansAlarm	1.3.6.1.4.1.5655.4.1.2.3	CISCO-ENTITY-FRU-CONTROL-MIB	Trap is sent Current status available using the **show environment** CLI command	
pchassisTempAlarm	1.3.6.1.4.1.5655.4.1.2.4	CISCO-ENTITY-SENSOR-MIB	Trap is sent Current status available using the **show environment** CLI command	
pchassisVoltageAlarm	1.3.6.1.4.1.5655.4.1.2.5	CISCO-ENTITY-SENSOR-MIB	Trap is sent Current status available using the **show environment** CLI command	
pchassisNumSlots	1.3.6.1.4.1.5655.4.1.2.6	ENTITY-MIB	The entity MIB shows the number of slots. According to the Entity MIB mapping and container relationship	
pchassisSlotConfig	1.3.6.1.4.1.5655.4.1.2.7	ENTITY-MIB	The Entity MIB indicates which slots in the chassis have modules inserted.	
pchassisPsuType	1.3.6.1.4.1.5655.4.1.2.8	ENTITY-MIB	**entity physical name** shows the Name	
pchassisLineFeedAlarm	1.3.6.1.4.1.5655.4.1.2.10	ENTITY-MIB	Trap is sent Current status available using the **show environment** CLI command	

Table A-8 pmoduleGrp (1.3.6.1.4.1.5655.4.1.3)

pcube Object Name	OID	New MIB	New Object Name	OID
pmoduleTable	1.3.6.1.4.1.5655.4.1.3.1	ENTITY-MIB	entPhysicalTable	1.3.6.1.2.1.47.1.1.1
pmoduleEntry	1.3.6.1.4.1.5655.4.1.3.1.1	ENTITY-MIB	entPhysicalEntry	1.3.6.1.2.1.47.1.1.1.1
pmoduleIndex	1.3.6.1.4.1.5655.4.1.3.1.1.1	ENTITY-MIB	entPhysicalIndex	1.3.6.1.2.1.47.1.1.1.1.1
pmoduleType	1.3.6.1.4.1.5655.4.1.3.1.1.2	ENTITY-MIB	entPhysicalName entPhysicalClass	1.3.6.1.2.1.47.1.1.1.1.7 1.3.6.1.2.1.47.1.1.1.1.5
pmoduleNumTraffic Processors	1.3.6.1.4.1.5655.4.1.3.1.1.3	CISCO-PROCESS-MIB	cpmCPUTotalTable.cpmCPUTotal PhysicalIndex	1.3.6.1.4.1.9.9.109.1.1.1.1.2
pmoduleSlotNum	1.3.6.1.4.1.5655.4.1.3.1.1.4	ENTITY-MIB	entPhysicalContainedIn	1.3.6.1.2.1.47.1.1.1.1.4
pmoduleHwVersion	1.3.6.1.4.1.5655.4.1.3.1.1.5	ENTITY-MIB	entPhysicalHardwareRev	1.3.6.1.2.1.47.1.1.1.1.8
pmoduleNumPorts	1.3.6.1.4.1.5655.4.1.3.1.1.6	ENTITY-MIB	The number of entries in ENTITY-MIB with entPhysicalClass = port	
pmoduleNumLinks	1.3.6.1.4.1.5655.4.1.3.1.1.7	ENTITY-MIB	The number of entries in ENTITY-MIB with entPhysicalClass = other	
pmoduleConnection Mode	1.3.6.1.4.1.5655.4.1.3.1.1.8	Not mapped	Use CLI command: **show interface linecard connection-mode**	
pmoduleSerialNumber	1.3.6.1.4.1.5655.4.1.3.1.1.9	ENTITY-MIB	entPhysicalSerialNum	1.3.6.1.2.1.47.1.1.1.1.11
pmoduleUpStream AttackFilteringTime	1.3.6.1.4.1.5655.4.1.3.1.1.10	CISCO-SERVICE-CONTROL-ATTACK-MIB	cscaInfoUpStream AttackFilteringTime	1.3.6.1.4.1.9.9.693.1.3.1.1
pmoduleUpStreamLastAttackFilteringTime	1.3.6.1.4.1.5655.4.1.3.1.1.11	CISCO-SERVICE-CONTROL-ATTACK-MIB	cscaInfoUpStreamLast AttackFilteringTime	1.3.6.1.4.1.9.9.693.1.3.1.2
pmoduleDownStream AttackFilteringTime	1.3.6.1.4.1.5655.4.1.3.1.1.12	CISCO-SERVICE-CONTROL-ATTACK-MIB	cscaInfoDownStream AttackFilteringTime	1.3.6.1.4.1.9.9.693.1.3.1.3
pmoduleDownStream LastAttack FilteringTime	1.3.6.1.4.1.5655.4.1.3.1.1.13	CISCO-SERVICE-CONTROL-ATTACK-MIB	cscaInfoDownStreamLast AttackFilteringTime	1.3.6.1.4.1.9.9.693.1.3.1.4
pmoduleAttack ObjectsClearTime	1.3.6.1.4.1.5655.4.1.3.1.1.14		Not mapped	

Table A-8 pmoduleGrp (1.3.6.1.4.1.5655.4.1.3) (continued)

pcube Object Name	OID	New MIB	New Object Name	OID
pmoduleAdminStatus	1.3.6.1.4.1.5655.4.1.3.1.1.15	ENTITY-MIB	entStateAdmin	1.3.6.1.2.1.131.1.1.1.2
pmoduleOperStatus	1.3.6.1.4.1.5655.4.1.3.1.1.16	ENTITY-MIB	entStateOper entStateStandby	1.3.6.1.2.1.131.1.1.1.3 1.3.6.1.2.1.131.1.1.1.6

Table A-9 linkGrp (1.3.6.1.4.1.5655.4.1.4): All Mapped Objects Mapped to CISCO-SERVICE-CONTROL-LINK-MIB

pcube Object Name	OID	New Object Name	OID
linkTable	1.3.6.1.4.1.5655.4.1.4.1	cscLinkStatusTable	1.3.6.1.4.1.9.9.631.1.2
linkEntry	1.3.6.1.4.1.5655.4.1.4.1.1	cscLinkStatusEntry	1.3.6.1.4.1.9.9.631.1.2.1
linkModuleIndex	1.3.6.1.4.1.5655.4.1.4.1.1.1	Not mapped. Not needed in the new structure - Information provided by ENTITY-MIB	
linkIndex	1.3.6.1.4.1.5655.4.1.4.1.1.2	entPhysicalIndex	1.3.6.1.2.1.47.1.1.1.1.1
linkAdminModeOnActive	1.3.6.1.4.1.5655.4.1.4.1.1.3	csclLinkStatusAdminModeOnActive	1.3.6.1.4.1.9.9.631.1.2.1.1
linkAdminModeOnFailure	1.3.6.1.4.1.5655.4.1.4.1.1.4	csclLinkStatusAdminModeOnFailure	1.3.6.1.4.1.9.9.631.1.2.1.2
linkOperMode	1.3.6.1.4.1.5655.4.1.4.1.1.5	cscLinkStatusOperMode	1.3.6.1.4.1.9.9.631.1.2.1.3
linkStatusReflectionEnable	1.3.6.1.4.1.5655.4.1.4.1.1.6	cscLinkStatusStatusReflectionEnable	1.3.6.1.4.1.9.9.631.1.2.1.4
linkSubscriberSidePortIndex	1.3.6.1.4.1.5655.4.1.4.1.1.7	cscLinkStatusSubscriberSidePortIndex	1.3.6.1.4.1.9.9.631.1.2.1.5
linkNetworkSidePortIndex	1.3.6.1.4.1.5655.4.1.4.1.1.8	cscLinkStatusNetworkSidePortIndex	1.3.6.1.4.1.9.9.631.1.2.1.6
		New objects: cscLinkStatusStatusReflectionState = noLinkReflection(1) reflectingFailureToNetwork (2) reflectingFailureToSubscriber(3)	1.3.6.1.4.1.9.9.631.1.2.1.7

Table A-10 diskGrp (1.3.6.1.4.1.5655.4.1.5): All Objects Mapped to HOST-RESOURCE-MIB

pcube Object Name	OID	New Object Name	OID
diskNumUsedBytes	1.3.6.1.4.1.5655.4.1.5.1	hrStorageTable.hrStorageUsed	1.3.6.1.2.1.25.2.3.1.6
diskNumFreeBytes	1.3.6.1.4.1.5655.4.1.5.2	hrStorageTable.hrStorageUsed hrStorageTable.hrStorageSize	1.3.6.1.2.1.25.2.3.1.6 1.3.6.1.2.1.25.2.3.1.5

Table A-11 rdrFormatterGrp (1.3.6.1.4.1.5655.4.1.6): All Mapped Objects Mapped to CISCO-SERVICE-CONTROL-RDR-MIB

pcube Object Name	OID	New Object Name	OID
rdrFormatterEnable	1.3.6.1.4.1.5655.4.1.6.1	cServiceControlRDRFormatterEnable	1.3.6.1.4.1.9.9.637.1.1.1.1
rdrFormatterDestTable	1.3.6.1.4.1.5655.4.1.6.2	cServiceControlRDRFormatterDestTable	1.3.6.1.4.1.9.9.637.1.2
rdrFormatterDestEntry	1.3.6.1.4.1.5655.4.1.6.2.1	cServiceControlRDRFormatterDestEntry	1.3.6.1.4.1.9.9.637.1.2 .1
rdrFormatterDest IPAddr	1.3.6.1.4.1.5655.4.1.6.2.1.1	cServiceControlRDRFormatterDestIPAddr	1.3.6.1.4.1.9.9.637.1.2.1.3
rdrFormatterDestPort	1.3.6.1.4.1.5655.4.1.6.2.1.2	cServiceControlRDRFormatterDestPort	1.3.6.1.4.1.9.9.637.1.2.1.4
rdrFormatterDestPriority	1.3.6.1.4.1.5655.4.1.6.2.1.3	cServiceControlRDRFormatterDestPriority	1.3.6.1.4.1.9.9.637.1.2.1.5
rdrFormatterDestStatus	1.3.6.1.4.1.5655.4.1.6.2.1.4	cServiceControlRDRFormatterDestStatus	1.3.6.1.4.1.9.9.637.1.2.1.6
rdrFormatterDest ConnectionStatus	1.3.6.1.4.1.5655.4.1.6.2.1.5	cServiceControlRDRFormatterDestConnectionStatus	1.3.6.1.4.1.9.9.637.1.2.1.7
rdrFormatterDestNum ReportsSent	1.3.6.1.4.1.5655.4.1.6.2.1.6	cServiceControlRDRFormatterDestNumReportsSent	1.3.6.1.4.1.9.9.637.1.2.1.8
rdrFormatterDestNum ReportsDiscarded	1.3.6.1.4.1.5655.4.1.6.2.1.7	cServiceControlRDRFormatterDestNumReportsDiscarded	1.3.6.1.4.1.9.9.637.1.2.1.9
rdrFormatterDestReport Rate	1.3.6.1.4.1.5655.4.1.6.2.1.8	cServiceControlRDRFormatterDestReportRate	1.3.6.1.4.1.9.9.637.1.2.1.10
rdrFormatterDestReport RatePeak	1.3.6.1.4.1.5655.4.1.6.2.1.9	Not mapped	
rdrFormatterDestReport RatePeakTime	1.3.6.1.4.1.5655.4.1.6.2.1.10	Not mapped	
rdrFormatterNum ReportsSent	1.3.6.1.4.1.5655.4.1.6.3	cServiceControlRDRFormatterNumReportsSent	1.3.6.1.4.1.9.9.637.1.1.1.2
rdrFormatterNum ReportsDiscarded	1.3.6.1.4.1.5655.4.1.6.4	cServiceControlRDRFormatterNumReportsDiscarded	1.3.6.1.4.1.9.9.637.1.1.1.3

Table A-11 rdrFormatterGrp (1.3.6.1.4.1.5655.4.1.6): All Mapped Objects Mapped to CISCO-SERVICE-CONTROL-RDR-MIB (continued)

pcube Object Name	OID	New Object Name	OID
rdrFormatterClear CountersTime	1.3.6.1.4.1.5655.4.1.6.5	Not mapped	
rdrFormatterReportRate	1.3.6.1.4.1.5655.4.1.6.6	cServiceControlRDRFormatterReportRate	1.3.6.1.4.1.9.9.637.1.1.1.4
rdrFormatterReportRate Peak	1.3.6.1.4.1.5655.4.1.6.7	cscRdrFormatterReportRatePeak	1.3.6.1.4.1.9.9.637.1.1.1.5
rdrFormatterReportRate PeakTime	1.3.6.1.4.1.5655.4.1.6.8	cscRdrFormatterReportRatePeakTime	1.3.6.1.4.1.9.9.637.1.1.1.6
rdrFormatterProtocol	1.3.6.1.4.1.5655.4.1.6.9	cServiceControlRDRFormatterProtocol	1.3.6.1.4.1.9.9.637.1.1.1.7
rdrFormatterForwarding Mode	1.3.6.1.4.1.5655.4.1.6.10	cServiceControlRDRFormatterForwardingMode	1.3.6.1.4.1.9.9.637.1.1.1.8
rdrFormatterCategory Table	1.3.6.1.4.1.5655.4.1.6.11	Available through the CLI.	
rdrFormatterCategory Entry	1.3.6.1.4.1.5655.4.1.6.11.1	Available through the CLI.	
rdrFormatterCategory Index	1.3.6.1.4.1.5655.4.1.6.11.1.1	Available through the CLI.	
rdrFormatterCategory Name	1.3.6.1.4.1.5655.4.1.6.11.1.2	Available through the CLI.	
rdrFormatterCategory NumReportsSent	1.3.6.1.4.1.5655.4.1.6.11.1.3	Available through the CLI.	
rdrFormatterCategory NumReportsDiscarded	1.3.6.1.4.1.5655.4.1.6.11.1.4	Available through the CLI.	
rdrFormatterCategory ReportRate	1.3.6.1.4.1.5655.4.1.6.11.1.5	Available through the CLI.	
rdrFormatterCategory ReportRatePeak	1.3.6.1.4.1.5655.4.1.6.11.1.6	Not mapped	
rdrFormatterCategory ReportRatePeakTime	1.3.6.1.4.1.5655.4.1.6.11.1.7	Not mapped	
rdrFormatterCategory NumReportsQueued	1.3.6.1.4.1.5655.4.1.6.11.1.8	Available through the CLI.	
rdrFormatterCategory DestTable	1.3.6.1.4.1.5655.4.1.6.12	Available through the CLI.	
rdrFormatterCategory DestEntry	1.3.6.1.4.1.5655.4.1.6.12.1	Available through the CLI.	

Table A-11 rdrFormatterGrp (1.3.6.1.4.1.5655.4.1.6): All Mapped Objects Mapped to CISCO-SERVICE-CONTROL-RDR-MIB (continued)

pcube Object Name	OID	New Object Name	OID
rdrFormatterCategory DestPriority	1.3.6.1.4.1.5655.4.1.6.12.1.1	Available through the CLI.	
rdrFormatterCategory DestStatus	1.3.6.1.4.1.5655.4.1.6.12.1.2	Available through the CLI.	

Table A-12 loggerGrp (1.3.6.1.4.1.5655.4.1.7): all Mapped Objects Mapped to CISCO-SYSLOG-EVENT-EXT-MIB

pcube Object Name	OID	New Object Name	OID
loggerUserLogEnable	1.3.6.1.4.1.5655.4.1.7.1	Not mapped	
loggerUserLogNumInfo	1.3.6.1.4.1.5655.4.1.7.2	cslogEventDispositionTable	1.3.6.1.4.1.9.9.270.1.1.5.1.3
loggerUserLogNumWarning	1.3.6.1.4.1.5655.4.1.7.3	cslogEventDispositionTable	1.3.6.1.4.1.9.9.270.1.1.5.1.4
loggerUserLogNumError	1.3.6.1.4.1.5655.4.1.7.4	cslogEventDispositionTable	1.3.6.1.4.1.9.9.270.1.1.5.1.5
loggerUserLogNumFatal	1.3.6.1.4.1.5655.4.1.7.5	cslogEventDispositionTable	1.3.6.1.4.1.9.9.270.1.1.5.1.6
loggerUserLogClearCountersTime	1.3.6.1.4.1.5655.4.1.7.6	Not mapped	

Table A-13 subscribersGrp (1.3.6.1.4.1.5655.4.1.8): All Mapped Objects Mapped to CISCO-SERVICE-CONTROL-SUBSCRIBERS-MIB

pcube Object Name	OID	New Object Name	OID
subscribersInfoTable	1.3.6.1.4.1.5655.4.1.8.1.1	cServiceControlSubscribersInfoTable	1.3.6.1.4.1.9.9.628.1.2
subscribersInfoEntry	1.3.6.1.4.1.5655.4.1.8.1.1.1	cServiceControlSubscribersInfoEntry	1.3.6.1.4.1.9.9.628.1.2.1
subscribersNumIntroduced	1.3.6.1.4.1.5655.4.1.8.1.11.1	cServiceControlSubscribersNumIntroduced	1.3.6.1.4.1.9.9.628.1.2.1.1
subscribersNumFree	1.3.6.1.4.1.5655.4.1.8.1.1.2	cServiceControlSubscribersNumFree	1.3.6.1.4.1.9.9.628.1.2.1.2
subscribersNumIpAddrMappings	1.3.6.1.4.1.5655.4.1.8.1.1.3	cServiceControlSubscribersNumIpAddrMappings	1.3.6.1.4.1.9.9.628.1.2.1.3
subscribersNumIpAddrMappings Free	1.3.6.1.4.1.5655.4.1.8.1.1.4	cServiceControlSubscribersNumIpAddrMappingsFree	1.3.6.1.4.1.9.9.628.1.2.1.4
subscribersNumIpRange Mappings	1.3.6.1.4.1.5655.4.1.8.1.1.5	cServiceControlSubscribersNumIpRangeMappings	1.3.6.1.4.1.9.9.628.1.2.1.5
subscribersNumIpRange MappingsFree	1.3.6.1.4.1.5655.4.1.8.1.1.6	cServiceControlSubscribersNumIpRange MappingsFree	1.3.6.1.4.1.9.9.628.1.2.1.6
subscribersNumVlanMappings	1.3.6.1.4.1.5655.4.1.8.1.1.7	cServiceControlSubscribersNumVlanMappings	1.3.6.1.4.1.9.9.628.1.2.1.7
subscribersNumVlan MappingsFree	1.3.6.1.4.1.5655.4.1.8.1.1.8	cServiceControlSubscribersNumVlanMappingsFree	1.3.6.1.4.1.9.9.628.1.2.1.8

Table A-13 subscribersGrp (1.3.6.1.4.1.5655.4.1.8): All Mapped Objects Mapped to CISCO-SERVICE-CONTROL-SUBSCRIBERS-MIB (continued)

pcube Object Name	OID	New Object Name	OID
subscribersNumActive	1.3.6.1.4.1.5655.4.1.8.1.1.9	cServiceControlSubscribersNumActive	1.3.6.1.4.1.9.9.628.1.2.1.9
subscribersNumActivePeak	1.3.6.1.4.1.5655.4.1.8.1.1.10	Not mapped	
subscribersNumActivePeakTime	1.3.6.1.4.1.5655.4.1.8.1.1.11	Not mapped	
subscribersNumUpdates	1.3.6.1.4.1.5655.4.1.8.1.1.12	cServiceControlSubscribersNumUpdates	1.3.6.1.4.1.9.9.628.1.2.1.10
subscribersCountersClearTime	1.3.6.1.4.1.5655.4.1.8.1.1.13	Not mapped	
subscribersNumTpIpRange Mappings	1.3.6.1.4.1.5655.4.1.8.1.1.14	cServiceControlSubscribersNumTpIpRangeMappings	1.3.6.1.4.1.9.9.628.1.2.1.11
subscribersNumTpIpRange MappingsFree	1.3.6.1.4.1.5655.4.1.8.1.1.15	cServiceControlSubscribersNumTpIpRange MappingsFree	1.3.6.1.4.1.9.9.628.1.2.1.12
subscribersNumAnonymous	1.3.6.1.4.1.5655.4.1.8.1.1.16	cServiceControlSubscribersNumAnonymous	1.3.6.1.4.1.9.9.628.1.2.1.13
subscribersNumWithSessions	1.3.6.1.4.1.5655.4.1.8.1.1.17	cServiceControlSubscribersNumWithSessions	1.3.6.1.4.1.9.9.628.1.2.1.14
subscribersPropertiesTable	1.3.6.1.4.1.5655.4.1.8.2	Not mapped.	
subscribersPropertiesEntry	1.3.6.1.4.1.5655.4.1.8.2.1	Not mapped.	
spIndex	1.3.6.1.4.1.5655.4.1.8.2.1.1	Not mapped.	
spName	1.3.6.1.4.1.5655.4.1.8.2.1.2	Not mapped.	
spType	1.3.6.1.4.1.5655.4.1.8.2.1.3	Not mapped.	
subscribersPropertiesValueTable	1.3.6.1.4.1.5655.4.1.8.3	Not mapped.	
subscribersPropertiesValueEntry	1.3.6.1.4.1.5655.4.1.8.3.1	Not mapped.	
spvIndex	1.3.6.1.4.1.5655.4.1.8.3.1.1	Not mapped.	
spvSubName	1.3.6.1.4.1.5655.4.1.8.3.1.2	Not mapped.	
spvPropertyName	1.3.6.1.4.1.5655.4.1.8.3.1.3	Not mapped.	
spvRowStatus	1.3.6.1.4.1.5655.4.1.8.3.1.4	Not mapped.	
spvPropertyStringValue	1.3.6.1.4.1.5655.4.1.8.3.1.5	Not mapped.	
spvPropertyUintValue	1.3.6.1.4.1.5655.4.1.8.3.1.6	Not mapped.	
spvPropertyCounter64Value	1.3.6.1.4.1.5655.4.1.8.3.1.7	Not mapped.	
New MIB Objects			
		cServiceControlSubscriberMappingFailedReason	1.3.6.1.4.1.9.9.628.1.2.1.15
		cServiceControlSubsribersMaxSupported	1.3.6.1.4.1.9.9.628.1.2.1.16
		cServiceControlSubscribersRowStatus	1.3.6.1.4.1.9.9.628.1.1.1.4

Table A-13 subscribersGrp (1.3.6.1.4.1.5655.4.1.8): All Mapped Objects Mapped to CISCO-SERVICE-CONTROL-SUBSCRIBERS-MIB (continued)

pcube Object Name	OID	New Object Name	OID
		cServiceControlSubscribersPackageIndex	1.3.6.1.4.1.9.9.628.1.1.1.5
		cServiceControlSubscribersRealTimeMonitor	1.3.6.1.4.1.9.9.628.1.1.1.6

Table A-14 trafficProcessorGrp (1.3.6.1.4.1.5655.4.1.9)

pcube Object Name	OID	New MIB	New Object Name	OID
tpInfoTable	1.3.6.1.4.1.5655.4.1.9.1	CISCO-SERVICE-CONTROL-TP-STATS-MIB	cscTpTable	1.3.6.1.4.1.9.9.634.1.1
tpInfoEntry	1.3.6.1.4.1.5655.4.1.9.1.1	CISCO-SERVICE-CONTROL-TP-STATS-MIB	cscTpEntry	1.3.6.1.4.1.9.9.634.1.1.1
tpModuleIndex	1.3.6.1.4.1.5655.4.1.9.1.1.1	CISCO-SERVICE-CONTROL-TP-STATS-MIB	entPhysicalIndex	1.3.6.1.2.1.47.1.1.1.1.1
tpIndex	1.3.6.1.4.1.5655.4.1.9.1.1.2	CISCO-SERVICE-CONTROL-TP-STATS-MIB	entPhysicalIndex	1.3.6.1.2.1.47.1.1.1.1.1
tpTotalNumHandled Packets	1.3.6.1.4.1.5655.4.1.9.1.1.3	CISCO-SERVICE-CONTROL-TP-STATS-MIB	cscTpTotalHandledPackets	1.3.6.1.4.1.9.9.634.1.1.1.1
tpTotalNumHandled Flows	1.3.6.1.4.1.5655.4.1.9.1.1.4	CISCO-SERVICE-CONTROL-TP-STATS-MIB	cscTpTotalHandledFlows	1.3.6.1.4.1.9.9.634.1.1.1.2
tpNumActiveFlows	1.3.6.1.4.1.5655.4.1.9.1.1.5	CISCO-SERVICE-CONTROL-TP-STATS-MIB	cscTpActiveFlows	1.3.6.1.4.1.9.9.634.1.1.1.3
tpNumActiveFlowsPeak	1.3.6.1.4.1.5655.4.1.9.1.1.6		Not mapped.	
tpNumActiveFlows PeakTime	1.3.6.1.4.1.5655.4.1.9.1.1.7		Not mapped.	
tpNumTcpActiveFlows	1.3.6.1.4.1.5655.4.1.9.1.1.8	CISCO-SERVICE-CONTROL-TP-STATS-MIB	cscTpTcpActiveFlows	1.3.6.1.4.1.9.9.634.1.1.1.4
TpNumTcpActive FlowsPeak	1.3.6.1.4.1.5655.4.1.9.1.1.9		Not mapped.	
tpNumTcpActive FlowsPeakTime	1.3.6.1.4.1.5655.4.1.9.1.1.10		Not mapped.	
tpNumUdpActiveFlows	1.3.6.1.4.1.5655.4.1.9.1.1.11	CISCO-SERVICE-CONTROL-TP-STATS-MIB	cscTpUdpActiveFlows	1.3.6.1.4.1.9.9.634.1.1.1.5
tpNumUdpActive FlowsPeak	1.3.6.1.4.1.5655.4.1.9.1.1.12		Not mapped.	

Table A-14 trafficProcessorGrp (1.3.6.1.4.1.5655.4.1.9) (continued)

pcube Object Name	OID	New MIB	New Object Name	OID
tpNumUdpActive FlowsPeakTime	1.3.6.1.4.1.5655.4.1.9.1.1.13		Not mapped.	
tpNumNonTcpUdp ActiveFlows	1.3.6.1.4.1.5655.4.1.9.1.1.14	CISCO-SERVICE-CONTROL-TP-STATS-MIB	cscTpUdpActiveFlows	1.3.6.1.4.1.9.9.634.1.1.1.5
tpNumNonTcpUdp ActiveFlowsPeak	1.3.6.1.4.1.5655.4.1.9.1.1.15		Not mapped.	
tpNumNonTcpUdp ActiveFlowsPeakTime	1.3.6.1.4.1.5655.4.1.9.1.1.16		Not mapped.	
tpTotalNum BlockedPackets	1.3.6.1.4.1.5655.4.1.9.1.1.17	CISCO-SERVICE-CONTROL-TP-STATS-MIB	cscTpTotalBlockedPackets	1.3.6.1.4.1.9.9.634.1.1.1.6
tpTotalNum BlockedFlows	1.3.6.1.4.1.5655.4.1.9.1.1.18	CISCO-SERVICE-CONTROL-TP-STATS-MIB	cscTpTotalBlockedFlows	1.3.6.1.4.1.9.9.634.1.1.1.7
tpTotalNumDiscarded PacketsDueToBwLimit	1.3.6.1.4.1.5655.4.1.9.1.1.19	CISCO-SERVICE-CONTROL-TP-STATS-MIB	cscTpTotalBandwidth DiscardedPackets	1.3.6.1.4.1.9.9.634.1.1.1.8
tpTotalNumWred DiscardedPackets	1.3.6.1.4.1.5655.4.1.9.1.1.20	CISCO-SERVICE-CONTROL-TP-STATS-MIB	cscTpTotalWred DiscardedPackets	1.3.6.1.4.1.9.9.634.1.1.1.9
tpTotalNumFragments	1.3.6.1.4.1.5655.4.1.9.1.1.21	CISCO-SERVICE-CONTROL-TP-STATS-MIB	cscTpTotalFragments	1.3.6.1.4.1.9.9.634.1.1.1.10
tpTotalNum NonIpPackets	1.3.6.1.4.1.5655.4.1.9.1.1.22	CISCO-SERVICE-CONTROL-TP-STATS-MIB	cscTpTotalNonIpPackets	1.3.6.1.4.1.9.9.634.1.1.1.11
tpTotalNumIp CrcErrPackets	1.3.6.1.4.1.5655.4.1.9.1.1.23	CISCO-SERVICE-CONTROL-TP-STATS-MIB	cscTpTotalIpChecksum ErrPackets	1.3.6.1.4.1.9.9.634.1.1.1.12
tpTotalNumIpLength ErrPackets	1.3.6.1.4.1.5655.4.1.9.1.1.24	CISCO-SERVICE-CONTROL-TP-STATS-MIB	cscTpTotalIpLengt ErrPackets	1.3.6.1.4.1.9.9.634.1.1.1.13
tpTotalNumIp BroadcastPackets	1.3.6.1.4.1.5655.4.1.9.1.1.25	CISCO-SERVICE-CONTROL-TP-STATS-MIB	cscTpTotalIp BroadcastPackets	1.3.6.1.4.1.9.9.634.1.1.1.14
tpTotalNum TtlErrPackets	1.3.6.1.4.1.5655.4.1.9.1.1.26	CISCO-SERVICE-CONTROL-TP-STATS-MIB	cscTpTotalTTLErrPackets	1.3.6.1.4.1.9.9.634.1.1.1.15
tpTotalNumTcpUdp CrcErrPackets	1.3.6.1.4.1.5655.4.1.9.1.1.27	CISCO-SERVICE-CONTROL-TP-STATS-MIB	cscTpTotalTcpUdp ChksumErrPackets	1.3.6.1.4.1.9.9.634.1.1.1.16
tpClearCountersTimed	1.3.6.1.4.1.5655.4.1.9.1.1.28		Not mapped.	
tpHandledPacketsRate	1.3.6.1.4.1.5655.4.1.9.1.1.29	CISCO-SERVICE-CONTROL-TP-STATS-MIB	cscTpHandledPacketsRate	1.3.6.1.4.1.9.9.634.1.1.1.17

Table A-14 trafficProcessorGrp (1.3.6.1.4.1.5655.4.1.9) (continued)

pcube Object Name	OID	New MIB	New Object Name	OID
tpHandledPackets RatePeak	1.3.6.1.4.1.5655.4.1.9.1.1.30		Not mapped.	
tpHandledPackets RatePeakTime	1.3.6.1.4.1.5655.4.1.9.1.1.31		Not mapped.	
tpHandledFlowsRate	1.3.6.1.4.1.5655.4.1.9.1.1.32	CISCO-SERVICE-CONTROL-TP-STATS-MIB	cscTpHandledFlowsRate	1.3.6.1.4.1.9.9.634.1.1.1.18
tpHandledFlows RatePeak	1.3.6.1.4.1.5655.4.1.9.1.1.33		Not mapped	
tpHandledFlows RatePeakTime	1.3.6.1.4.1.5655.4.1.9.1.1.34		Not mapped	
tpCpuUtilization	1.3.6.1.4.1.5655.4.1.9.1.1.35	CISCO-PROCESS-MIB	cpmCPUTotal1minRev cpmCPUTotal5minRev	1.3.6.1.4.1.9.9.109.1.1.1.1.7 1.3.6.1.4.1.9.9.109.1.1.1.1.8
tpCpuUtilizationPeak	1.3.6.1.4.1.5655.4.1.9.1.1.36		Not mapped	
tpCpuUtilization PeakTime	1.3.6.1.4.1.5655.4.1.9.1.1.37		Not mapped	
tpFlowsCapacity Utilization	1.3.6.1.4.1.5655.4.1.9.1.1.38	CISCO-SERVICE-CONTROL-TP-STATS-MIB	cscTpFlowsCapacity Utilization	1.3.6.1.4.1.9.9.634.1.1.1.19
tpFlowsCapacity UtilizationPeak	1.3.6.1.4.1.5655.4.1.9.1.1.39		Not mapped	
tpFlowsCapacity UtilizationPeakTime	1.3.6.1.4.1.5655.4.1.9.1.1.40		Not mapped	
tpServiceLoss	1.3.6.1.4.1.5655.4.1.9.1.1.41	CISCO-SERVICE-CONTROL-TP-STATS-MIB	cscTpServiceLoss	1.3.6.1.4.1.9.9.634.1.1.1.20

Table A-15 pportGrp (1.3.6.1.4.1.5655.4.1.10)

pcube Object Name	OID	New MIB	New Object Name	OID
pportTable	1.3.6.1.4.1.5655.4.1.10.1	Not mapped	Information provided by ENTITY-MIB	
pportEntry	1.3.6.1.4.1.5655.4.1.10.1.1	Not mapped	Information provided by ENTITY-MIB	
pportModuleIndex	1.3.6.1.4.1.5655.4.1.10.1.1.1	ENTITY-MIB	entPhysicalContainedIn	1.3.6.1.2.1.47.1.1.1.1.4
pportIndex	1.3.6.1.4.1.5655.4.1.10.1.1.2	ENTITY-MIB	Each port is uniquely defined. Index is implied by entPhysicalIndex or ifIndex	1.3.6.1.2.1.47.1.1.1.1.1

Table A-15 pportGrp (1.3.6.1.4.1.5655.4.1.10) (continued)

pcube Object Name	OID	New MIB	New Object Name	OID
pportType	1.3.6.1.4.1.5655.4.1.10.1.1.3	RFC1213-MIB	ifType and cevPort Also entPhysicalVendorType can be queried for this value.	1.3.6.1.2.1.2.2.1.3
pportNumTxQueues	1.3.6.1.4.1.5655.4.1.10.1.1.4		Not mapped Information provided by CISCO-QUEUE-MIB cQIfSubqueues	
pportIfIndex	1.3.6.1.4.1.5655.4.1.10.1.1.5		Not mapped ifIndex mapping information provided by entAliasMappingTable.	
pportAdminSpeed	1.3.6.1.4.1.5655.4.1.10.1.1.6		Not mapped	
pportAdminDuplex	1.3.6.1.4.1.5655.4.1.10.1.1.7		Not mapped	
pportOperDuplex	1.3.6.1.4.1.5655.4.1.10.1.1.8		Not mapped	
pportLinkIndex	1.3.6.1.4.1.5655.4.1.10.1.1.9	ENTITY-MIB	entPhysicalIndex Available from ENTITY-MIB container/containee relationship: Chassis contains Slot contains Link contains Port	1.3.6.1.2.1.47.1.1.1.1.1
pportOperStatus	1.3.6.1.4.1.5655.4.1.10.1.1.10	ENTITY-MIB	entPhysicalIndex Defined in ENTITY-STATE-MIB.	1.3.6.1.2.1.47.1.1.1.1.1

Table A-16 txQueuesGrp (1.3.6.1.4.1.5655.4.1.11)

pcube Object Name	OID	New MIB	New Object Name	OID
txQueuesTable	1.3.6.1.4.1.5655.4.1.11.1	CISCO-QUEUE-MIB	cQIfTable and cQStatsTable	1.3.6.1.4.1.9.9.37.1.2
txQueuesEntry	1.3.6.1.4.1.5655.4.1.11.1.1	CISCO-QUEUE-MIB	cQStatsEntry	1.3.6.1.4.1.9.9.37.1.2.1
txQueuesModuleIndex	1.3.6.1.4.1.5655.4.1.11.1.1.1		Not mapped	
txQueuesPortIndex	1.3.6.1.4.1.5655.4.1.11.1.1.2	RFC1213-MIB	ifIndex The entry is indexed by ifIndex of IF-MIB	1.3.6.1.2.1.2.2.1.1
txQueuesQueueIndex	1.3.6.1.4.1.5655.4.1.11.1.1.3	CISCO-QUEUE-MIB	cQStatsQNumber	1.3.6.1.4.1.9.9.37.1.2.1.1
txQueuesDescription	1.3.6.1.4.1.5655.4.1.11.1.1.4	CISCO-QUEUE-MIB	cQIfTable.cQifQType	1.3.6.1.4.1.9.9.37.1.1.1.1

Table A-16 txQueuesGrp (1.3.6.1.4.1.5655.4.1.11) (continued)

pcube Object Name	OID	New MIB	New Object Name	OID
txQueuesBandwidth	1.3.6.1.4.1.5655.4.1.11.1.1.5	CISCO-QUEUE-MIB	cQStatsBandwidth	1.3.6.1.4.1.9.9.37.1.2.1.5
txQueuesUtilization	1.3.6.1.4.1.5655.4.1.11.1.1.6		Not mapped	
txQueuesUtilizationPeak	1.3.6.1.4.1.5655.4.1.11.1.1.7		Not mapped	
txQueuesUtilizationPeak Time	1.3.6.1.4.1.5655.4.1.11.1.1.8		Not mapped	
txQueuesClearCounters Time	1.3.6.1.4.1.5655.4.1.11.1.1.9		Not mapped	
txQueuesDroppedBytes	1.3.6.1.4.1.5655.4.1.11.1.1.10	CISCO-QUEUE-MIB	cQStatsDiscards This object counts bytes	1.3.6.1.4.1.9.9.37.1.2.1.4

Table A-17 globalControllerssGrp (1.3.6.1.4.1.5655.4.1.12): All Mapped Objects Mapped to CISCO-SERVICE-CONTROLLER-MIB

pcube Object Name	OID	New Object Name	OID
globalControllersTable	1.3.6.1.4.1.5655.4.1.12.1	cscGlobalControllersTable	1.3.6.1.4.1.9.9.667.0.1
globalControllersEntry	1.3.6.1.4.1.5655.4.1.12.1.1	cscGlobalControllersEntry	1.3.6.1.4.1.9.9.667.0.1.1
globalControllersModuleIndex	1.3.6.1.4.1.5655.4.1.12.1.1.1	Not mapped Provided by entPhysicalIndex	
globalControllersPortIndex	1.3.6.1.4.1.5655.4.1.12.1.1.2	Not mapped Provided by entityPhyIndex	
globalControllersIndex	1.3.6.1.4.1.5655.4.1.12.1.1.3	cscGlobalControllersId	1.3.6.1.4.1.9.9.667.0.1.1.2
globalControllersDescription	1.3.6.1.4.1.5655.4.1.12.1.1.4	cscGlobalControllersDescription	1.3.6.1.4.1.9.9.667.0.1.1.3
globalControllersBandwidth	1.3.6.1.4.1.5655.4.1.12.1.1.5	cscGlobalControllersBandwidth	1.3.6.1.4.1.9.9.667.0.1.1.5
globalControllersUtilization	1.3.6.1.4.1.5655.4.1.12.1.1.6	cscGlobalControllersUtilization	1.3.6.1.4.1.9.9.667.0.1.1.6
globalControllers UtilizationPeak	1.3.6.1.4.1.5655.4.1.12.1.1.7	Not mapped	
globalControllers UtilizationPeakTime	1.3.6.1.4.1.5655.4.1.12.1.1.8	Not mapped	
globalControllers ClearCountersTime	1.3.6.1.4.1.5655.4.1.12.1.1.9	Not mapped	
globalControllers DroppedBytes	1.3.6.1.4.1.5655.4.1.12.1.1.10	Not mapped	

Table A-18 trafficCountersGrp (1.3.6.1.4.1.5655.4.1.14): All Objects Mapped to CISCO-SERVICE-CONTROL-TP-STATS-MIB

pcube Object Name	OID	New Object Name	OID
trafficCountersTable	1.3.6.1.4.1.5655.4.1.14.1	cscTpStatsTrafficCountersTable	1.3.6.1.4.1.9.9.634.1.2
trafficCountersEntry	1.3.6.1.4.1.5655.4.1.14.1.1	cscTpStatsTrafficCountersEntry	1.3.6.1.4.1.9.9.634.1.2.1
trafficCounterIndex	1.3.6.1.4.1.5655.4.1.14.1.1.1	cscTpStatsTrafficCounterIndex	1.3.6.1.4.1.9.9.634.1.2.1.1
trafficCounterValue	1.3.6.1.4.1.5655.4.1.14.1.1.2	cscTpStatsTrafficCounterValue	1.3.6.1.4.1.9.9.634.1.2.1.2
trafficCounterName	1.3.6.1.4.1.5655.4.1.14.1.1.3	cscTpStatsTrafficCounterName	1.3.6.1.4.1.9.9.634.1.2.1.3
trafficCounterType	1.3.6.1.4.1.5655.4.1.14.1.1.4	cscTpStatsTrafficCounterType	1.3.6.1.4.1.9.9.634.1.2.1.4

Table A-19 attackGrp (1.3.6.1.4.1.5655.4.1.15): All Objects Mapped to CISCO-SERVICE-CONTROL-ATTACK-MIB

pcube Object Name	OID	New Object Name	OID
attackTypeTable	1.3.6.1.4.1.5655.4.1.15.1	cscaTypeTable	1.3.6.1.4.1.9.9.693.1.2
attackTypeEntry	1.3.6.1.4.1.5655.4.1.15.1.1	cscaTypeEntry	1.3.6.1.4.1.9.9.693.1.2.1
attackTypeIndex	1.3.6.1.4.1.5655.4.1.15.1.1.1	cscaTypeIndex	1.3.6.1.4.1.9.9.693.1.2.1.1
attackTypeName	1.3.6.1.4.1.5655.4.1.15.1.1.2	not mapped	
attackTypeCurrentNumAttacks	1.3.6.1.4.1.5655.4.1.15.1.1.3	cscaTypeCurrentNumAttacks	1.3.6.1.4.1.9.9.693.1.2.1.2
attackTypeTotalNumAttacks	1.3.6.1.4.1.5655.4.1.15.1.1.4	cscaTypeTotalNumAttacks	1.3.6.1.4.1.9.9.693.1.2.1.3
attackTypeTotalNumFlows	1.3.6.1.4.1.5655.4.1.15.1.1.5	cscaTypeTotalNumFlows	1.3.6.1.4.1.9.9.693.1.2.1.4
attackTypeTotalNumSeconds	1.3.6.1.4.1.5655.4.1.15.1.1.6	cscaTypeTotalNumSeconds	1.3.6.1.4.1.9.9.693.1.2.1.5
New Objects at this MIB - Used For TRAPS			
		cscaType	1.3.6.1.4.1.9.9.693.1.1.1
		cscaSourceAddressType	1.3.6.1.4.1.9.9.693.1.1.2
		cscaSourceAddress	1.3.6.1.4.1.9.9.693.1.1.3
		cscaDestinationAddressType	1.3.6.1.4.1.9.9.693.1.1.4
		cscaDestinationAddress	1.3.6.1.4.1.9.9.693.1.1.5
		cscaDestinationPort	1.3.6.1.4.1.9.9.693.1.1.6
		cscaFilterStatus	1.3.6.1.4.1.9.9.693.1.1.7

Table A-20 PCUBE SeEvents (1.3.6.1.4.1.5655.4.0)

pcube Object Name	OID	New MIB	New Object Name	OID
operationalStatus OperationalTrap	1.3.6.1.4.1.5655.4.0.1	CISCO-ENTITY-FRU-CONTROL-MIB	cefcModuleOperStatus = oper_ok	1.3.6.1.4.1.9.9.117.2.0.1
operationalStatus WarningTrap	1.3.6.1.4.1.5655.4.0.2	CISCO-ENTITY-FRU-CONTROL-MIB	cefcModuleOperStatus = ok_but_diag_failed	1.3.6.1.4.1.9.9.117.2.0.1
operationalStatusFailure Trap	1.3.6.1.4.1.5655.4.0.3	CISCO-ENTITY-FRU-CONTROL-MIB	cefcModuleOperStatus = failed	1.3.6.1.4.1.9.9.117.2.0.1
systemResetTrap	1.3.6.1.4.1.5655.4.0.4	CISCO-ENTITY-FRU-CONTROL-MIB	cefcPowerStatusChange cefcFRUPowerOperStatus =off_env_other(1) cefcFRUPowerAdminStatus =power_on(1)	1.3.6.1.4.1.9.9.117.2.0.2
chassisTempAlarmOn Trap	1.3.6.1.4.1.5655.4.0.5	CISCO-ENTITY-SENSOR-MIB	entSensorThresholdNotification entSensorThresholdValue =1 entSensorValue =1	1.3.6.1.4.1.9.9.91.2.0.1
chassisTempAlarmOff Trap	1.3.6.1.4.1.5655.4.0.6	CISCO-ENTITY-SENSOR-MIB	entSensorThresholdNotification entSensorThresholdValue =0 entSensorValue =0	1.3.6.1.4.1.9.9.91.2.0.1
chassisVoltageAlarmOn Trap	1.3.6.1.4.1.5655.4.0.7	CISCO-ENTITY-SENSOR-MIB	entSensorThresholdNotification entSensorThresholdValue =1 entSensorValue =1	1.3.6.1.4.1.9.9.91.2.0.1
chassisFansAlarmOn Trap	1.3.6.1.4.1.5655.4.0.8	CISCO-ENTITY-FRU-CONTROL-MIB	cefcFanTrayStatusChange This trap is sent only when the fan tray is removed.	1.3.6.1.4.1.9.9.117.2.0.6
chassisPowerSupply AlarmOnTrap	1.3.6.1.4.1.5655.4.0.9	CISCO-ENTITY-FRU-CONTROL-MIB	cefcPowerSupplyOutputChange Trap functions as follows: • Unplug power cord from Cisco SCE platform—trap sent • Plug power cord into Cisco SCE platform—trap not sent • Remove a PSU—trap sent • Insert a PSU—trap not sent	1.3.6.1.4.1.9.9.117.2.0.7

Table A-20 PCUBE SeEvents (1.3.6.1.4.1.5655.4.0) (continued)

pcube Object Name	OID	New MIB	New Object Name	OID
rdrActiveConnection Trap	1.3.6.1.4.1.5655.4.0.10	CISCO-SERVICE-CONTROL-RDR-MIB	cServiceControlRdrActiveConnectionTrap	1.3.6.1.4.1.9.9.637.0.5
rdrNoActiveConnection Trap	1.3.6.1.4.1.5655.4.0.11	CISCO-SERVICE-CONTROL-RDR-MIB	cServiceControlRdrNoActiveConnection Trap	1.3.6.1.4.1.9.9.637.0.3
rdrConnectionUpTrap	1.3.6.1.4.1.5655.4.0.12	CISCO-SERVICE-CONTROL-RDR-MIB	cServiceControlRdrConnectionStatus UpTrap	1.3.6.1.4.1.9.9.637.0.6
rdrConnectionDown Trap	1.3.6.1.4.1.5655.4.0.13	CISCO-SERVICE-CONTROL-RDR-MIB	cServiceControlRdrConnectionStatusDow nTrap	1.3.6.1.4.1.9.9.637.0.4
telnetSessionStartedTrap	1.3.6.1.4.1.5655.4.0.14	CISCO-TELNET-SERVER-MIB	ctsSessionStarted	1.3.6.1.4.1.9.9.630.0.2
telnetSessionEndedTrap	1.3.6.1.4.1.5655.4.0.15	CISCO-TELNET-SERVER-MIB	ctsSessionEnded	1.3.6.1.4.1.9.9.630.0.1
telnetSessionDenied AccessTrap	1.3.6.1.4.1.5655.4.0.16	CISCO-TELNET-SERVER-MIB	ctsSessionDenied	1.3.6.1.4.1.9.9.630.0.3
telnetSessionBadLogin Trap	1.3.6.1.4.1.5655.4.0.17	CISCO-TELNET-SERVER-MIB	ctsSessionLoginFailure	1.3.6.1.4.1.9.9.630.0.4
loggerUserLogIsFull Trap	1.3.6.1.4.1.5655.4.0.18	CISCO-ENTITY-ALARM-MIB	ceAlarmAsserted/ ceAlarmCleared	1.3.6.1.4.1.9.9.138.2.0.1/ 1.3.6.1.4.1.9.9.138.2.0.2
sntpClockDriftWarnTrap	1.3.6.1.4.1.5655.4.0.19	CISCO-ENTITY-ALARM-MIB	ceAlarmAsserted/ ceAlarmCleared	1.3.6.1.4.1.9.9.138.2.0.1/ 1.3.6.1.4.1.9.9.138.2.0.2
linkModeBypassTrap	1.3.6.1.4.1.5655.4.0.20	CISCO-SERVICE-CONTROL-LINK-MIB	cServiceControlLinkModeChangeTrap	1.3.6.1.4.1.9.9.631.0.1
linkModeForwarding Trap	1.3.6.1.4.1.5655.4.0.21	CISCO-SERVICE-CONTROL-LINK-MIB	cServiceControlLinkModeChangeTrap	1.3.6.1.4.1.9.9.631.0.1
linkModeCutoffTrap	1.3.6.1.4.1.5655.4.0.22	CISCO-SERVICE-CONTROL-LINK-MIB	cServiceControlLinkModeChangeTrap	1.3.6.1.4.1.9.9.631.0.1
moduleAttackFilter ActivatedTrap	1.3.6.1.4.1.5655.4.0.25	CISCO-SERVICE-CONTROL-ATTACK-MIB	cscaFilterChange	1.3.6.1.4.1.9.9.693.0.1
moduleAttackFilter DeactivatedTrap	1.3.6.1.4.1.5655.4.0.26	CISCO-SERVICE-CONTROL-ATTACK-MIB	cscaFilterChange	1.3.6.1.4.1.9.9.693.0.1
moduleEmAgentGeneric Trap	1.3.6.1.4.1.5655.4.0.27		Not mapped	

Table A-20 PCUBE SeEvents (1.3.6.1.4.1.5655.4.0) (continued)

pcube Object Name	OID	New MIB	New Object Name	OID
linkModeSniffingTrap	1.3.6.1.4.1.5655.4.0.28	CISCO-SERVICE-CONTROL-LINK-MIB	cServiceControlLinkModeChangeTrap	1.3.6.1.4.1.9.9.631.0.1
moduleRedundancy ReadyTrap	1.3.6.1.4.1.5655.4.0.29	CISCO-ENTITY-REDUNDANCY-MIB	ceRedunProtectStatusChange ceRedunMbrStatusCurrent=protection_provided(0x10)	1.3.6.1.4.1.9.9.498.0.2
moduleRedundant ConfigurationMismatch Trap	1.3.6.1.4.1.5655.4.0.30	CISCO-ENTITY-REDUNDANCY-MIB	ceRedunProtectStatusChange ceRedunMbrStatusCurrent=failure(0x40)	1.3.6.1.4.1.9.9.498.0.2
moduleLostRedundancy Trap	1.3.6.1.4.1.5655.4.0.31	CISCO-ENTITY-REDUNDANCY-MIB	ceRedunProtectStatusChange ceRedunMbrStatusCurrent=protection_lock_out(0x0)	1.3.6.1.4.1.9.9.498.0.2
moduleSmConnection DownTrap	1.3.6.1.4.1.5655.4.0.32	CISCO-ENTITY-ALARM-MIB	ceAlarmAsserted/ ceAlarmCleared	1.3.6.1.4.1.9.9.138.2.0.1/ 1.3.6.1.4.1.9.9.138.2.0.2
moduleSmConnection UpTrap	1.3.6.1.4.1.5655.4.0.33	CISCO-ENTITY-ALARM-MIB	ceAlarmAsserted/ ceAlarmCleared	1.3.6.1.4.1.9.9.138.2.0.1/ 1.3.6.1.4.1.9.9.138.2.0.2
moduleOperStatus ChangeTrap	1.3.6.1.4.1.5655.4.0.34	CISCO-ENTITY-FRU-CONTROL-MIB	Not mapped	1.3.6.1.4.1.9.9.117.2.0.1
portOperStatusChange Trap	1.3.6.1.4.1.5655.4.0.35	CISCO-ENTITY-FRU-CONTROL-MIB	entStateOperEnabled or entStateOperDisabled	1.3.6.1.4.1.9.9.117.2.0.1
chassisLineFeedAlarm OnTrap	1.3.6.1.4.1.5655.4.0.36	CISCO-ENTITY-FRU-CONTROL-MIB	cefcPowerStatusChange cefcFRUPowerOperStatus =off_env_power(5) cefcFRUPowerAdminStatus =power_on(1)	1.3.6.1.4.1.9.9.117.2.0.2
rdrFormatterCategory DiscardingReportsTrap	1.3.6.1.4.1.5655.4.0.37	CISCO-SERVICE-CONTROL-RDR-MIB	cServiceControlRdrCategoryDiscarding ReportsTrap	1.3.6.1.4.1.9.9.637.0.2
rdrFormatterCategory StoppedDiscarding ReportsTrap	1.3.6.1.4.1.5655.4.0.38	CISCO-SERVICE-CONTROL-RDR-MIB	cServiceControlRdrCategoryStopped DiscardingReportsTrap	1.3.6.1.4.1.9.9.637.0.1
sessionStartedTrap	1.3.6.1.4.1.5655.4.0.39	CISCO-SECURE-SHELL-MIB	cssSessionStartedTrap	1.3.6.1.4.1.9.9.339.0.3
sessionEndedTrap	1.3.6.1.4.1.5655.4.0.40	CISCO-SECURE-SHELL-MIB	cssSessionEndedTrap	1.3.6.1.4.1.9.9.339.0.1

Table A-20 PCUBE SeEvents (1.3.6.1.4.1.5655.4.0) (continued)

pcube Object Name	OID	New MIB	New Object Name	OID
sessionDeniedAccess Trap	1.3.6.1.4.1.5655.4.0.41	CISCO-SECURE-SHELL-MIB	cssSessionDeniedTrap	1.3.6.1.4.1.9.9.339.0.2
sessionBadLoginTrap	1.3.6.1.4.1.5655.4.0.42	CISCO-SECURE-SHELL-MIB	cssSessionDeniedTrap cssSessionDeniedReason	1.3.6.1.4.1.9.9.339.0.2 1.3.6.1.4.1.9.9.339.1.3.2.1.4
illegalSubscriber MappingTrap	1.3.6.1.4.1.5655.4.0.43	CISCO-SERVICE-CONTROL-SUBSCRIBER-MIB	cServiceControlSubscriberMappingTrap	1.3.6.1.4.1.9.9.628.0.1
loggerLineAttackLog FullTrap	1.3.6.1.4.1.5655.4.0.44	CISCO-ENTITY-ALARM-MIB	ceAlarmAsserted/ ceAlarmCleared	1.3.6.1.4.1.9.9.138.2.0.1/ 1.3.6.1.4.1.9.9.138.2.0.2
pullRequestNumber	1.3.6.1.4.1.5655.4.0.46		Not mapped	
pullRequestRetryFailed Trap	1.3.6.1.4.1.5655.4.0.47	Not mapped	For SM connection issues the following notification can be used: ceAlarmAsserted ceAlarmCleared	For SM connection issues the following OIDs can be used: 1.3.6.1.4.1.9.9.138.2.0.1 1.3.6.1.4.1.9.9.138.2.0.2
mplsVpnTotalHW MappingsThreshold ExceededTrap	1.3.6.1.4.1.5655.4.0.48	CISCO-ENTITY-ALARM-MIB	ceAlarmAsserted/ ceAlarmCleared	1.3.6.1.4.1.9.9.138.2.0.1/ 1.3.6.1.4.1.9.9.138.2.0.2

Note applicationGrp (1.3.6.1.4.1.5655.4.1.13) is not mapped to a current MIB.

Table A-21 pcubeEnageMIB 1.3.6.1.4.1.5655.4.2

pcube Object Name	OID	Corresponding RDR	Objects not mapped	OID
linkGrp	1.3.6.1.4.1.5655.4.2.2	Link Usage RDRs or the RPT_LUR DB table	linkServiceUpDroppedPackets linkServiceDownDroppedPackets linkServiceUpDroppedBytes	1.3.6.1.4.1.5655.4.2.2
packageGrp	1.3.6.1.4.1.5655.4.2.3	Package Usage RDRs or the RPT_PUR DB table	packageServiceUpDroppedPackets packageServiceDownDroppedPackets packageServiceUpDroppedBytes	1.3.6.1.4.1.5655.4.2.3

Table A-21 pcubeEnageMIB 1.3.6.1.4.1.5655.4.2 (continued)

pcube Object Name	OID	Corresponding RDR	Objects not mapped	OID
subscriberGrp	1.3.6.1.4.1.5655.4.2.4	Subscriber Usage RDRs	none	1.3.6.1.4.1.5655.4.2.4
serviceCounterGrp	1.3.6.1.4.1.5655.4.2.5	Service Configuration API or the INI_VALUES DB table	none	1.3.6.1.4.1.5655.4.2.5

Cisco SCE Platform-Specific MIB Information

This section contains definitions that are specific to the Cisco SCE platforms for certain standard and Cisco MIB objects.

CISCO-ENTITY-ALARM-MIB

ceAlarmDescrSeverity (integer)

ceAlarmDescrSeverity.1.1—3

ceAlarmDescrSeverity.1.2—3

ceAlarmDescrSeverity.1.3—2

ceAlarmDescrSeverity.1.4—4

ceAlarmDescrSeverity.1.5—3

ceAlarmDescrSeverity.1.6—3

ceAlarmDescrSeverity.1.7—3

ceAlarmDescrText (octet string)

ceAlarmDescrText.1.1—Logger user log is full

ceAlarmDescrText.1.2—Sntp clock drift warn

ceAlarmDescrText.1.3— Module sm connection down

ceAlarmDescrText.1.4—Module sm connection up

ceAlarmDescrText.1.5—Logger line attack log full

ceAlarmDescrText.1.6—External bypass device isn't connected

ceAlarmDescrText.1.7—External bypass device is connected

MIB Updates

Release 3.5.5 MIB Updates

The definitions of the following MIB objects have been updated, but the updated definitions may not yet appear in the online MIB.

CISCO-SERVICE-CONTROL-TP-STATS-MIB

cscTpServiceLoss

The average service loss **per second in the last minute** (or since the last time the TP counters were cleared) in a traffic Processor, in units of 0.001%.

The service loss is computed as the relative amount of traffic which was bypassed by the Cisco SCE from one side to another without being serviced due to lack of resources (either CPU or memory).

In previous versions, it indicated the service loss **since last reboot** or last time the counters were cleared.

Release 3.6.0 MIB Updates

CISCO-PROCESSOR-MIB

cpmCPUTotalTable

Previous to SCOS Release 3.6.0, the cpmCPUTotalTable provided information only on the traffic processors. Starting in SCOS Release 3.6.0, the cpmCPUTotalTable provides statistics about the control processor as well as the traffic processors.

ENTITY-MIB

In SCOS Release 3.6.0, the following hardware entities were added to the ENTITY-MIB:

- Second Cisco SCE 8000-SCM module in slot 2.
- Control processor

Index Changes

In SCOS Release 3.6.0, for the following objects, the index was changed from blade index to chassis index:

MIB tables:

- cisco-service-control-attack-mib
  - cscaTypeTable
  - cscaInfoTable
- cisco-service-control-rdr-mib
  - rdrDestTable
  - rdrCategoryTable
  - rdrFormatterTable
  - rdrCategoryDest
- cisco-service-control-subscriber-mib
  - subscriberTable
  - subscriberInfoTable
- cisco-service-control-tp-stats-mib
  - tpStatsTrafficCounterTable
- entity-state-mib
  - entityStateTable

Traps

- rdr-active-connection
- rdr-no-active-connection
- rdr-connection-up
- rdr-connection-down
- rdr-formatterCategoryDisacard
- rdrCategoryStoppedDiscard
- userlogFull
- userLogNotFull
- lineAttackLog
- SmConnectionDown
- SmConnectionUp
- illegalSubscriberMapping
- PowerStatusChange
- warningStatusChange
- failureStatusChange
- SipAttack

Release 3.6.5 MIB Updates

Temperature Sensor Traps Updated

In SCOS Release 3.6.5, the following changes were made in the traps related to the CISCO-ENTITY-SENSOR MIB, entSensorThresholdNotification (OID 1.3.6.1.4.1.9.9.91.2.0.1):

- Traps will be sent for the entity corresponding to the FRU that actually violated the threshold.

  Previously all temperature sensor traps were from the FAN entity.

- Only one trap will be sent that indicates both the current value as well as the threshold value of the component within the FRU that is responsible for the threshold violation or conformance.

  Previously there were two traps (on and off), and there was no indication of the actual current value or the threshold.

- The trap will also identify the component within the FRU that is associated with the threshold violation or conformance.

Note CISCO-ENTITY-SENSOR MIB is read-only. Thresholds are internally defined and cannot be changed.

Note Temperature reported for the entities (FRUs) are a normalized temperature since there is no single temperature reading for an entire FRU.

Release 3.7.0 MIB Updates

SNMP Support for Aggregative Global Controllers

In Release 3.7.0, the CISCO-SERVICE-CONTROLLER-MIB has been expanded to provide information regarding the Aggregative Global Controllers. The following tables have been added to the CISCO-SERVICE-CONTROLLER-MIB:

- cscAggregativeGlobalControllersTable
- cscAggregativeGlobalControllersLinkTable

You can display the contents of the Aggregative Global Controllers MIB using the following command:

show interface LineCard 0 aggregative-global-controllers side (subscriber | network) agc *agc-id*

Note For a complete description of the CISCO-SERVICE-CONTROLLER-MIB, see CISCO-SERVICE-CONTROLLER-MIB.

linkUp and linkDown Notification Traps

In SCOS Release 3.7.0, the ifDescr varbind has been added to the linkUp and linkDown traps, which provide the description of the link that is either in the down state or some other state (as indicated by ifOperStatus).

The linkUp and linkDown notification traps related to IF-MIB have the following variable bindings (varbinds):

- ifIndex--Represents the index value of the interface.
- ifAdminStatus--Represents the desired state of the interface.
- ifOperStatus--Represents the current operational state of the interface.

Cisco SCE 8000 supports the linkUp/linkDown trap only on management ports. If there is a change to the management-port state, Cisco SCE 8000 sends two traps—one linkUp/linkDown trap and one entStateOperEnabled/entStateOperDisabled trap. But if there is a change to the traffic-port state, Cisco SCE 8000 sends only the entStateOperEnabled/entStateOperDisabled trap.

Release 4.1.0 MIB Updates

This section provides MIB updates in Release 4.1.0:

SNMP TRAP for Global Attacks

The CISCO-SERVICE-CONTROL-ATTACK-MIB is updated with a new SNMP trap cscaGlobalAttackFilterChange (1.3.6.1.4.1.9.9.693.0.2) to notify global attacks. If a global attack occurs, Cisco SCE sends this trap at the start and stop of the attack.

The trap includes following components:

- entPhysicalName (1.3.6.1.2.1.47.1.1.1.1.7)—Indicates the name of the originating physical entity.
- cscaGlobalAttackType (1.3.6.1.4.1.9.9.693.1.1.10)—Indicates the type of the global attack.
- cscaFilterStatus (1.3.6.1.4.1.9.9.693.1.1.7)—Indicates whether the global attack has started or ended. This also indicates whether the attack filter status is active or not.
- cscaTypeOriginatedByNetworkSide (1.3.6.1.4.1.9.9.693.1.2.1.6)—Indicates the origin or source of the attack, whether it originated from network or from subscriber.

Possible values of cscaGlobalAttackType indicating the type of global attack:

- icmpAttack (1)
- udpAttack (2)
- udpFragmentAttack (3)
- tcpSynAttack (4)
- tcpRstAttack (5)
- tcpFragmentAttack (6)
- tcpNonSynAttack (7)

SNMP Walk Functionality for Temperature MIBs

- Supports the SNMP walk on the entSensorValueTable to get the temperature values from the sensors module. The temperature values for each temperature sensor on SCE 8000 is captured.



Monitoring Cisco SCE Platform Utilization

Revised: September 23, 2015

Introduction

This appendix explains how to monitor Cisco SCE platforms that are installed in real traffic.

- "Cisco SCE Platform Utilization Indicators" section on page B-2
- Service Loss, page B-3

As with any network device, the Cisco SCE platform has its performance and capacity envelopes. As the network evolves, the utilization of the Cisco SCE platform can increase and these envelopes might be reached. It is, therefore, advisable to monitor Cisco SCE platform to be sure that utilization remains at a level that supports reliable and consistent service.

When the Cisco SCE platform reaches its performance envelopes, it activates certain mechanisms that insure that no traffic will be dropped while in this state. These mechanisms will prioritize packet handling over service related actions. As a result, symptoms of service loss might be experienced. Following are several examples:

- Broken reports during the congestion period (sometimes appears as saw-tooth pattern).
- Bandwidth enforcement levels are not met.
- No UDP traffic is being reported (this is because the Cisco SCE platform will automatically filter all UDP traffic in certain cases as a last resort).

Monitoring the Cisco SCE platform can be divided onto two main areas:

- Monitoring Cisco SCE platform utilization
- Monitoring service loss

Cisco SCE Platform Utilization Indicators

The Cisco SCE platform exposes several indicators to allow the network operators to easily monitor whether it is working within its performance and capacity specifications:

- CPU Utilization, page B-2
- Flows Capacity, page B-2
- Subscribers Capacity, page B-2

CPU Utilization

- **SNMP**

  cpmCPUTotal1minRev and *cpmCPUTotal5minRev*, available for each Traffic Processor.

  Refer to the Cisco process MIB for more information.

- **CLI commands**

  show snmp MIB cisco-process | include cpmCPUTotal1minRev

  show snmp MIB cisco-process | include cpmCPUTotal5minRev

It is advisable to consider sizing of the solution when the CPU utilization exceeds 85% regularly at peak hours.

Flows Capacity

- **SNMP**

  cscTpFlowsCapacityUtilization available for each Traffic Processor.

  Refer to the cisco-service-control-tp-stats MIB for more information.

- **CLI command**

  show snmp MIB cisco-service-control-tp-stats | include cscTpFlowsCapacityUtilization

It is advisable to consider sizing of the solution when the flows capacity utilization exceeds 90% regularly at peak hours.

Subscribers Capacity

- **SNMP**

  cServiceControlSubscribersInfoEntry

  Refer to the cisco-service-control-subscriber MIB for more information

- **CLI command**

  show snmp MIB cisco-service-control-subscriber

The Cisco SCE 8000 platform supports up to 1M subscribers. You should make sure that the number of introduced subscribers plus the number of anonymous subscribers stays below this figure.

It is advisable that when subscribers utilization exceeds 90%, special attention should be given and sizing should be reconsidered.

Service Loss

Service Loss is an event which occurs when the Cisco SCE platform does not provide the processing it was expected to perform for any transaction in the network. This can occur due to either CPU or Flows shortage.

There are two different situations which can result with service loss in the Cisco SCE platform:

- Temporary – This might occur when some network pattern which is short in its nature occurred and caused the Cisco SCE platform to exhaust some of it resources temporarily. An example could be a DDoS attack that the Cisco SCE platform could not detect and filter.

  This is usually measured in seconds.

- Permanent – In cases where the Cisco SCE platform is installed in locations where the network traffic does not match its capacity and performance envelopes, permanent service loss can occur.

  This is measured in hours.

Service loss is defined as the ratio of the number of packets that did not receive service as expected to the total number of packets that were processed by the Cisco SCE platform.

Monitoring Service Loss

- **SNMP**

  cscTpServiceLoss MIB available for each traffic processor.

  Refer to the cisco-service-control-tp-stats MIB for more information.

It is expected that the Cisco SCE platform user will define timeslots in which this variable is monitored (reset it between timeslots).

Note that the units for this variable are 0.001% and the information is rounded down.



Cisco SCE 8000 Licensing Information

Revised: September 23, 2015

OpenSSH License

This file is part of the OpenSSH software.

The licences which components of this software fall under are as follows. First, we will summarize and say that all components are under a BSD licence, or a licence more free than that.

OpenSSH contains no GPL code.

1. * Copyright (c) 1995 Tatu Ylonen <ylo@cs.hut.fi>, Espoo, Finland

   * All rights reserved

   *

   * As far as I am concerned, the code I have written for this software

   * can be used freely for any purpose. Any derived versions of this

   * software must be clearly marked as such, and if the derived work is

   * incompatible with the protocol description in the RFC file, it must be

   * called by a name other than "ssh" or "Secure Shell".

[Tatu continues]

   * However, I am not implying to give any licenses to any patents or

   * copyrights held by third parties, and the software includes parts that

   * are not under my direct control. As far as I know, all included

   * source code is used in accordance with the relevant license agreements

   * and can be used freely for any purpose (the GNU license being the most

   * restrictive); see below for details.

   [However, none of that term is relevant at this point in time. All of these restrictively licenced software components which he talks about have been removed from OpenSSH, i.e.,

   - RSA is no longer included, found in the OpenSSL library
   - IDEA is no longer included, its use is deprecated
   - DES is now external, in the OpenSSL library

- GMP is no longer used, and instead we call BN code from OpenSSL
- Zlib is now external, in a library
- The make-ssh-known-hosts script is no longer included
- TSS has been removed
- MD5 is now external, in the OpenSSL library
- RC4 support has been replaced with ARC4 support from OpenSSL
- Blowfish is now external, in the OpenSSL library

[The licence continues]

Note that any information and cryptographic algorithms used in this software are publicly available on the Internet and at any major bookstore, scientific library, and patent office worldwide. More information can be found e.g. at "http://www.cs.hut.fi/crypto".

The legal status of this program is some combination of all these permissions and restrictions. Use only at your own responsibility. You will be responsible for any legal consequences yourself; I am not making any claims whether possessing or using this is legal or not in your country, and I am not taking any responsibility on your behalf.

NO WARRANTY

BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

2. The 32-bit CRC compensation attack detector in deattack.c was contributed by CORE SDI S.A. under a BSD-style license.

```
* Cryptographic attack detector for ssh - source code
*
* Copyright (c) 1998 CORE SDI S.A., Buenos Aires, Argentina.
*
* All rights reserved. Redistribution and use in source and binary
* forms, with or without modification, are permitted provided that
* this copyright notice is retained.
*
```

* THIS SOFTWARE IS PROVIDED ``AS IS'' AND ANY EXPRESS OR IMPLIED
* WARRANTIES ARE DISCLAIMED. IN NO EVENT SHALL CORE SDI S.A. BE
* LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY OR
* CONSEQUENTIAL DAMAGES RESULTING FROM THE USE OR MISUSE OF THIS
* SOFTWARE.
*
* Ariel Futoransky <futo@core-sdi.com>
* <http://www.core-sdi.com>

3. ssh-keyscan was contributed by David Mazieres under a BSD-style license.

* Copyright 1995, 1996 by David Mazieres <dm@lcs.mit.edu>.
*
* Modification and redistribution in source and binary forms is
* permitted provided that due credit is given to the author and the
* OpenBSD project by leaving this copyright notice intact.

4. The Rijndael implementation by Vincent Rijmen, Antoon Bosselaers and Paulo Barreto is in the public domain and distributed with the following license:

* @version 3.0 (December 2000)
*
* Optimised ANSI C code for the Rijndael cipher (now AES)
*
* @author Vincent Rijmen <vincent.rijmen@esat.kuleuven.ac.be>
* @author Antoon Bosselaers <antoon.bosselaers@esat.kuleuven.ac.be>
* @author Paulo Barreto <paulo.barreto@terra.com.br>
*
* This code is hereby placed in the public domain.
*
* THIS SOFTWARE IS PROVIDED BY THE AUTHORS "AS IS" AND ANY EXPRESS
* OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED
* WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE
* ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHORS OR CONTRIBUTORS BE
* LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR
* CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF
* SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR
* BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY,
* WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE
* OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE,
* EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

5. One component of the ssh source code is under a 3-clause BSD license, held by the University of California, since we pulled these parts from original Berkeley code.

```
* Copyright (c) 1983, 1990, 1992, 1993, 1995
*      The Regents of the University of California.  All rights reserved.
*
* Redistribution and use in source and binary forms, with or without
* modification, are permitted provided that the following conditions
* are met:
* 1. Redistributions of source code must retain the above copyright
*    notice, this list of conditions and the following disclaimer.
* 2. Redistributions in binary form must reproduce the above copyright
*    notice, this list of conditions and the following disclaimer in the
*    documentation and/or other materials provided with the distribution.
* 3. Neither the name of the University nor the names of its contributors
*    may be used to endorse or promote products derived from this software
*    without specific prior written permission.
*
* THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS ``AS IS'' AND
* ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE
* IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR
  PURPOSE
* ARE DISCLAIMED.  IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE
* FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR
  CONSEQUENTIAL
* DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE
  GOODS
* OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
* HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT,
  STRICT
* LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY
* OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF
* SUCH DAMAGE.
```

6. Remaining components of the software are provided under a standard 2-term BSD licence with the following names as copyright holders:

- Markus Friedl
- Theo de Raadt
- Niels Provos
- Dug Song
- Aaron Campbell

- Damien Miller
- Kevin Steves
- Daniel Kouril
- Wesley Griffin
- Per Allansson
- Nils Nordman
- Simon Wilkinson

Portable OpenSSH additionally includes code from the following copyright holders, also under the 2-term BSD license:

- Ben Lindstrom
- Tim Rice
- Andre Lucas
- Chris Adams
- Corinna Vinschen
- Cray Inc.
- Denis Parker
- Gert Doering
- Jakob Schlyter
- Jason Downs
- Juha Yrjölä
- Michael Stone
- Networks Associates Technology, Inc.
- Solar Designer
- Todd C. Miller
- Wayne Schroeder
- William Jones
- Darren Tucker
- Sun Microsystems
- The SCO Group
- Daniel Walsh

7.

```
* Redistribution and use in source and binary forms, with or without
* modification, are permitted provided that the following conditions
* are met:
* 1. Redistributions of source code must retain the above copyright
*    notice, this list of conditions and the following disclaimer.
* 2. Redistributions in binary form must reproduce the above copyright
*    notice, this list of conditions and the following disclaimer in the
```

* documentation and/or other materials provided with the distribution.
*
* THIS SOFTWARE IS PROVIDED BY THE AUTHOR ``AS IS'' AND ANY EXPRESS OR
* IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES
* OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED.
* IN NO EVENT SHALL THE AUTHOR BE LIABLE FOR ANY DIRECT, INDIRECT,
* INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT
* NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE,
* DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY
* THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT
* (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF
* THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

8. Portable OpenSSH contains the following additional licenses:

a) md5crypt.c, md5crypt.h

* "THE BEER-WARE LICENSE" (Revision 42):
* <phk@login.dknet.dk> wrote this file. As long as you retain this
* notice you can do whatever you want with this stuff. If we meet
* some day, and you think this stuff is worth it, you can buy me a
* beer in return. Poul-Henning Kamp

b) snprintf replacement

* Copyright Patrick Powell 1995
* This code is based on code written by Patrick Powell
* (papowell@astart.com) It may be used for any purpose as long as this
* notice remains intact on all source code distributions

c) Compatibility code (openbsd-compat)

Apart from the previously mentioned licenses, various pieces of code in the openbsd-compat/ subdirectory are licensed as follows:

Some code is licensed under a 3-term BSD license, to the following copyright holders:

Todd C. Miller

Theo de Raadt

Damien Miller

Eric P. Allman

The Regents of the University of California

Constantin S. Svintsoff

* Redistribution and use in source and binary forms, with or without
* modification, are permitted provided that the following conditions

* are met:

* 1. Redistributions of source code must retain the above copyright

* notice, this list of conditions and the following disclaimer.

* 2. Redistributions in binary form must reproduce the above copyright

* notice, this list of conditions and the following disclaimer in the

* documentation and/or other materials provided with the distribution.

* 3. Neither the name of the University nor the names of its contributors

* may be used to endorse or promote products derived from this software

* without specific prior written permission.

*

* THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS ``AS IS'' AND

* ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE

* IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE

* ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE

* FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL

* DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS

* OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)

* HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT

* LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY

* OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF

* SUCH DAMAGE.

Some code is licensed under an ISC-style license, to the following copyright holders:

Internet Software Consortium.

Todd C. Miller

Reyk Floeter

Chad Mynhier

* Permission to use, copy, modify, and distribute this software for any

* purpose with or without fee is hereby granted, provided that the above

* copyright notice and this permission notice appear in all copies.

*

* THE SOFTWARE IS PROVIDED "AS IS" AND TODD C. MILLER DISCLAIMS ALL

* WARRANTIES WITH REGARD TO THIS SOFTWARE INCLUDING ALL IMPLIED WARRANTIES

* OF MERCHANTABILITY AND FITNESS. IN NO EVENT SHALL TODD C. MILLER BE LIABLE

* FOR ANY SPECIAL, DIRECT, INDIRECT, OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES

* WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION

* OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN

* CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

Some code is licensed under a MIT-style license to the following copyright holders:

Free Software Foundation, Inc.

* Permission is hereby granted, free of charge, to any person obtaining a *

* copy of this software and associated documentation files (the *

* "Software"), to deal in the Software without restriction, including *

* without limitation the rights to use, copy, modify, merge, publish, *

* distribute, distribute with modifications, sublicense, and/or sell *

* copies of the Software, and to permit persons to whom the Software is *

* furnished to do so, subject to the following conditions: *

*

* The above copyright notice and this permission notice shall be included *

* in all copies or substantial portions of the Software. *

*

* THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS

* OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF *

* MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. *

* IN NO EVENT SHALL THE ABOVE COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM,

* DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR

* OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR

* THE USE OR OTHER DEALINGS IN THE SOFTWARE. *

*

* Except as contained in this notice, the name(s) of the above copyright *

* holders shall not be used in advertising or otherwise to promote the *

* sale, use or other dealings in this Software without prior written *

* authorization. *

****

$OpenBSD: LICENCE,v 1.19 2004/08/30 09:18:08 markus Exp $

NetSNMP License

Various copyrights apply to this package, listed in various separate parts below. Be sure to read all parts.

---- Part 1: CMU/UCD copyright notice: (BSD like) -----

Copyright 1989, 1991, 1992 by Carnegie Mellon University

Derivative Work - 1996, 1998-2000

Copyright 1996, 1998-2000 The Regents of the University of California

All Rights Reserved

Permission to use, copy, modify and distribute this software and its documentation for any purpose and without fee is hereby granted, provided that the above copyright notice appears in all copies and that both that copyright notice and this permission notice appear in supporting documentation, and that the name of CMU and The Regents of the University of California not be used in advertising or publicity pertaining to distribution of the software without specific written permission.

CMU AND THE REGENTS OF THE UNIVERSITY OF CALIFORNIA DISCLAIM ALL WARRANTIES WITH REGARD TO THIS SOFTWARE, INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS. IN NO EVENT SHALL CMU OR THE REGENTS OF THE UNIVERSITY OF CALIFORNIA BE LIABLE FOR ANY SPECIAL, INDIRECT OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM THE LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

---- Part 2: Networks Associates Technology, Inc copyright notice (BSD) -----

Copyright (c) 2001-2003, Networks Associates Technology, Inc All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of the Networks Associates Technology, Inc nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS ``AS IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDERS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

---- Part 3: Cambridge Broadband Ltd. copyright notice (BSD) -----

Portions of this code are copyright (c) 2001-2003, Cambridge Broadband Ltd. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- The name of Cambridge Broadband Ltd. may not be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDER ``AS IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

---- Part 4: Sun Microsystems, Inc. copyright notice (BSD) -----

Copyright © 2003 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, U.S.A. All rights reserved.

Use is subject to license terms below.

This distribution may include materials developed by third parties.

Sun, Sun Microsystems, the Sun logo and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of the Sun Microsystems, Inc. nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS ``AS IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDERS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY,

WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

---- Part 5: Sparta, Inc copyright notice (BSD) -----

Copyright (c) 2003-2008, Sparta, Inc

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of Sparta, Inc nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS ``AS IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDERS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

---- Part 6: Cisco/BUPTNIC copyright notice (BSD) -----

Copyright (c) 2004, Cisco, Inc and Information Network

Center of Beijing University of Posts and Telecommunications.

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of Cisco, Inc, Beijing University of Posts and Telecommunications, nor the names of their contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS ``AS IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDERS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS;

OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

---- Part 7: Fabasoft R&D Software GmbH & Co KG copyright notice (BSD) -----

Copyright (c) Fabasoft R&D Software GmbH & Co KG, 2003

oss@fabasoft.com

Author: Bernhard Penz <bernhard.penz@fabasoft.com>

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentatio n and/or other materials provided with the distribution.
- The name of Fabasoft R&D Software GmbH & Co KG or any of its subsidiaries, brand or product names may not be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDER ``AS IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.