



Cisco Service Control Engine (SCE) Software Configuration Guide

Release 3.1.5
August 2007

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
<http://www.cisco.com>
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

Text Part Number: OL-7827-07

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.

CCDE, CCENT, Cisco Eos, Cisco HealthPresence, the Cisco logo, Cisco Lumin, Cisco Nexus, Cisco StadiumVision, Cisco TelePresence, Cisco WebEx, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn and Cisco Store are service marks; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQuick Study, IronPort, the IronPort logo, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website 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. (0812R)

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 Service Control Engine (SCE) Software Configuration Guide
© 2007 Cisco Systems, Inc. All rights reserved.



CONTENTS

About this Guide xxxi

Document Revision History xxxi

Organization xxxiv

Related Publications xxxvi

Conventions xxxvii

Obtaining Documentation and Submitting a Service Request xxxviii

CHAPTER 1

General Overview 1-1

Information About the Cisco Service Control Concept 1-1

The Cisco Service Control Solution 1-1

Service Control for Broadband Service Providers 1-2

Cisco Service Control Capabilities 1-2

The SCE Platform 1-3

Information About Management and Collection 1-4

Network Management 1-5

Subscriber Management 1-5

Service Configuration Management 1-6

Data Collection 1-6

CHAPTER 2

Command Line Interface 2-1

Information About Getting Help 2-2

Information About Authorization and Command Levels (Hierarchy) 2-3

CLI Command Hierarchy 2-3

Prompt Indications 2-6

Information About CLI Help Features 2-7

Partial Help 2-7

Argument Help 2-7

The [no] Prefix 2-8

Information About Navigational and Shortcut Features 2-8

Command History 2-9

Keyboard Shortcuts 2-9

Tab Completion 2-10

FTP User Name and Password 2-11

- Information About Managing Command Output 2-11
 - Scrolling the Screen Display 2-11
 - Filtering Command Output 2-12
 - Redirecting Command Output to a File 2-12
- CLI Authorization Levels 2-13
 - How to change from User to Viewer level authorization 2-14
 - How to log in with Root level authorization 2-15
- Exiting Modes 2-15
 - How to exit from the Privileged Exec mode and revert to the Viewer mode 2-15
 - How to exit from the Global Configuration Mode 2-16
- Information About Navigating Between Configuration Modes 2-16
 - Entering and Exiting Global Configuration Mode 2-16
 - How to enter the Global Configuration Mode 2-16
 - How to exit the Global Configuration Mode 2-16
 - Information About Interface Configuration Modes 2-17
 - Configuring the Physical Ports 2-17
 - Entering Management Interface Configuration Mode 2-18
 - How to enter Linecard Interface Configuration mode 2-19
 - Entering Ethernet Line Interface Configuration Mode 2-19
 - How to navigate from one Interface Configuration Mode to another 2-20
 - The "do" Command: Executing Commands Without Exiting 2-20
- How to create a CLI script 2-21

CHAPTER 3

Operations 3-1

- How to Manage Configurations 3-1
 - Managing Configurations 3-2
 - How to View Configurations 3-2
 - Viewing Configurations 3-2
 - Viewing Configurations: Example 3-3
 - How to Remove a Configuration 3-3
 - Removing the Configuration 3-3
 - How to Save the Configuration Settings 3-4
 - Saving the Configuration Settings 3-4
 - Saving the Configuration Settings: Example 3-5
 - How to Restore a Previous Configuration 3-5
 - Restoring a Previous Configuration 3-6
 - Restoring a Previous Configuration: Example 3-6
 - How to Backup Configuration Files 3-7
 - Backing Up Configuration Files 3-7

| | |
|--|------|
| Options | 3-7 |
| How to Create a Backup Configuration File | 3-7 |
| How to Upload a Backup Configuration File | 3-8 |
| How to Upgrade the SCE Platform Firmware | 3-8 |
| Upgrading SCE Platform Firmware | 3-9 |
| Upgrading SCE Platform Firmware: Example | 3-10 |
| How to Manage Application Files | 3-10 |
| Configuring Applications | 3-10 |
| Managing Application Files | 3-11 |
| How to Display Information about an Application File | 3-11 |
| How to Install an Application | 3-11 |
| How to Uninstall an Application | 3-12 |
| How to Upgrade an Application | 3-12 |
| How to Undo an Upgrade of an Application | 3-12 |
| How to Display the Last pqi File that was Installed | 3-12 |
| How to Monitor the Operational Status of the SCE Platform | 3-13 |
| Monitoring the Operational Status of the SCE Platform | 3-13 |
| How to Display the Current Operational Status of the SCE Platform | 3-15 |
| Displaying the Current Operational Status of the SCE Platform: Example | 3-15 |
| How to Display the SCE Platform Version Information | 3-15 |
| Displaying the SCE Platform Version Information | 3-15 |
| Displaying the SCE Platform Version Information: Example | 3-16 |
| How to Display the SCE Platform Inventory | 3-16 |
| Displaying the SCE Platform Inventory | 3-17 |
| Displaying the SCE Platform Inventory: Example | 3-17 |
| How to Display the System Uptime | 3-17 |
| Displaying the System Uptime | 3-17 |
| Displaying the System Uptime: Example | 3-18 |
| How to Reboot and Shut Down the SCE Platform | 3-18 |
| How to Reboot the SCE Platform | 3-18 |
| Rebooting the SCE Platform | 3-18 |
| Rebooting the SCE Platform: Example | 3-18 |
| How to Shut Down the SCE Platform | 3-18 |
| Shutting Down the SCE Platform | 3-19 |
| Shutting Down the SCE Platform: Examples | 3-19 |

CHAPTER 4**Utilities 4-1**

| | |
|------------------------------|-----|
| How to Use the Setup Command | 4-1 |
| Setup Command Parameters | 4-1 |

- How to Enter the Setup Command 4-4
- How to Define Lists in the Setup Utility 4-5
 - Multiple entry parameters (Lists) 4-5
- Working with SCE Platform Files 4-6
 - Working with Directories 4-6
 - How to Create a Directory 4-6
 - How to Delete a Directory 4-6
 - How to Change Directories 4-7
 - How to Display your Working Directory 4-7
 - How to List the Files in a Directory 4-7
 - Working with Files 4-8
 - How to Rename a File 4-8
 - How to Delete a File 4-8
 - Copying Files 4-8
 - How to Display File Contents 4-9
 - How to Unzip a File 4-9

CHAPTER 5

Configuring the Management Interface and Security 5-1

- About Management Interface and Security 5-2
- How to Configure the Management Ports 5-2
 - Configuring the Management Ports 5-2
- How to Enter Management Interface Configuration Mode 5-3
 - About Management Interface Configuration Mode 5-3
- How to Configure the Management Port Physical Parameters 5-4
 - How to Set the IP Address and Subnet Mask of the Management Interface 5-4
 - Setting the IP Address and Subnet Mask of the Management Interface 5-4
 - Options 5-4
 - Setting the IP Address and Subnet Mask of the Management Interface: Example 5-5
 - How to Configure the Management Interface Speed and Duplex Parameters 5-5
 - Interface State Relationship to Speed and Duplex 5-6
 - How to Configure the Speed of the Management Interface 5-6
 - How to Configure the Duplex Operation of the Management Interface 5-7
 - How to Specify the Active Management Port 5-7
 - Specifying the Active Management Port 5-7
 - Options 5-8
 - Specifying the Active Management Port: Example 5-8
- How to Configure Management Interface Redundancy 5-8
 - About Management Port Redundancy 5-9
 - How to Configure the Management Ports for Redundancy 5-9

| | |
|---|------|
| How to Configure the Fail-Over Mode | 5-10 |
| Configuring the Fail-Over Mode | 5-10 |
| Options | 5-10 |
| How to Enable Automatic Fail-Over Mode | 5-10 |
| How to Disable Automatic Fail-Over Mode | 5-10 |
| How to Configure Management Interface Security | 5-10 |
| About Management Interface Security | 5-11 |
| How to Configure the IP Fragment Filter | 5-11 |
| Options | 5-11 |
| How to Enable the IP Fragment Filter | 5-11 |
| How to Disable the IP Fragment Filter | 5-11 |
| How to Configure the Permitted and Not-permitted IP Address Monitor | 5-12 |
| Options | 5-12 |
| How to Monitor Management Interface IP Filtering | 5-12 |
| Monitoring Management Interface IP Filtering | 5-12 |
| How to Configure the Available Interfaces | 5-13 |
| How to Configure TACACS+ Authentication, Authorization, and Accounting | 5-13 |
| Information About TACACS+ Authentication, Authorization, and Accounting | 5-13 |
| How to Configure the SCE Platform TACACS+ Client | 5-17 |
| How to Manage the User Database | 5-20 |
| How to Configure AAA Login Authentication | 5-24 |
| How to Configure AAA Privilege Level Authorization Methods | 5-25 |
| How to Configure AAA Accounting | 5-26 |
| How to Monitor TACACS+ Servers | 5-27 |
| How to Monitor TACACS+ Users | 5-27 |
| How to Configure Access Control Lists (ACLs) | 5-27 |
| About Access Control Lists | 5-28 |
| Options | 5-29 |
| How to Add Entries to an ACL | 5-29 |
| How to Remove an ACL | 5-29 |
| How to Define a Global ACL | 5-30 |
| How to Configure the Telnet Interface | 5-30 |
| About the Telnet Interface | 5-30 |
| How to Prevent Telnet Access | 5-30 |
| How to Assign an ACL to the Telnet Interface | 5-31 |
| How to Configure the Telnet Timeout | 5-31 |
| How to Configure the SSH Server | 5-31 |
| Information About the SSH Server | 5-32 |
| How to Manage the SSH Server | 5-32 |
| How to Monitor the Status of the SSH Server | 5-34 |

| | |
|--|------|
| Enabling the SNMP Interface | 5-34 |
| How to Enable the SNMP Interface | 5-34 |
| How to Disable the SNMP Interface | 5-34 |
| How to Configure and Manage the SNMP Interface | 5-34 |
| Information About the SNMP Interface | 5-35 |
| The SNMP Interface | 5-35 |
| SNMP Protocol | 5-35 |
| Security Considerations | 5-36 |
| Information About CLI | 5-36 |
| Information About MIBs | 5-37 |
| Configuration via SNMP | 5-43 |
| How to Configure SNMP Community Strings | 5-43 |
| Configuring SNMP Community Strings | 5-43 |
| How to Define a Community String | 5-43 |
| How to Remove a Community String | 5-44 |
| How to Display the Configured Community Strings | 5-45 |
| How to Configure SNMP Notifications | 5-45 |
| About SNMP Notifications | 5-45 |
| How to Define SNMP Hosts | 5-46 |
| How to Manage Passwords | 5-48 |
| About Passwords | 5-49 |
| Changing Your Password | 5-49 |
| How to Change Your Password | 5-49 |
| Verifying that the Password has been Successfully Changed | 5-50 |
| Password Encryption | 5-50 |
| About Password Encryption | 5-51 |
| How to Enable Password Encryption | 5-51 |
| How to Disable Password Encryption | 5-51 |
| Password Recovery | 5-51 |
| How to Recover the Passwords: SCOS versions before 2.5.5 | 5-51 |
| How to Recover the Passwords: SCOS versions 2.5.5 or later | 5-54 |
| IP Configuration | 5-54 |
| How to Configure the IP Routing Table | 5-54 |
| About the IP Routing Table | 5-55 |
| How to Configure the Default Gateway | 5-55 |
| How to Add an Entry to the IP Routing Table | 5-55 |
| How to Display the IP Routing Table | 5-56 |
| IP Advertising | 5-57 |
| About IP Advertising | 5-57 |

| | |
|---|------|
| How to Configure IP Advertising | 5-58 |
| How to Display the Current IP Advertising Configuration | 5-59 |
| How to Configure the IP Address of the Management Interface | 5-59 |
| About the Management Interface IP Address | 5-59 |
| Options | 5-59 |
| Configuring the IP Address of the Management Interface: Example | 5-60 |
| How to Configure Time Clocks and Time Zone | 5-60 |
| About Time Clocks and Time Zone | 5-60 |
| How to Display the System Time | 5-61 |
| Displaying the System Time: Example | 5-61 |
| How to Display the Calendar Time | 5-61 |
| Displaying the Calendar Time: Example | 5-61 |
| How to Set the System Clock | 5-61 |
| Options | 5-61 |
| Setting the System Clock: Example | 5-62 |
| How to Set the Calendar | 5-62 |
| Options | 5-62 |
| Setting the Calendar: Example | 5-62 |
| How to Set the Time Zone | 5-62 |
| Options | 5-63 |
| Setting the Time Zone: Example | 5-63 |
| How to Remove the Current Time Zone Setting | 5-63 |
| How to Configure Daylight Saving Time | 5-63 |
| Options | 5-64 |
| Guidelines | 5-65 |
| How to Define Recurring Daylight Saving Time Transitions | 5-65 |
| How to Define Non-Recurring Daylight Saving Time Transitions | 5-66 |
| How to Cancel the Daylight Saving Time Configuration | 5-66 |
| How to Display the Current Daylight Saving Time Configuration | 5-66 |
| How to Configure SNTP | 5-66 |
| About SNTP | 5-67 |
| How to Enable the SNTP Multicast Client | 5-67 |
| How to Disable the SNTP Multicast Client | 5-67 |
| How to Enable the SNTP Unicast Client | 5-67 |
| Options | 5-68 |
| Enabling SNTP Unicast Client: Example | 5-68 |
| How to Disable the SNTP Unicast Client | 5-68 |
| How to Disable the SNTP Unicast Client and Remove All Servers | 5-68 |
| How to Remove One SNTP Server | 5-68 |
| How to Define the SNTP Unicast Update Interval | 5-68 |

- Options **5-69**
- Defining the SNMP Unicast Update Interval: Example **5-69**
- How to Display SNMP Information **5-69**
 - Displaying SNMP Information: Example **5-69**
- How to Configure Domain Name Server (DNS) Settings **5-69**
 - About Domain Name Servers **5-70**
 - How to Configure DNS Lookup **5-70**
 - How to Enable DNS Lookup **5-70**
 - How to Disable DNS Lookup **5-70**
 - How to Configure Name Servers **5-71**
 - Options **5-71**
 - How to Define Domain Name Servers **5-71**
 - How to Remove a Domain Name Server **5-71**
 - How to Remove All Domain Name Servers **5-72**
 - How to Add a Host to the Host Table **5-72**
 - Options **5-72**
 - Adding Hosts to Removing them from the Host Table: Examples **5-72**
 - How to Display Current DNS Settings **5-72**
 - Displaying Current DNS Settings: Example **5-73**
- How to Configure the Management Port Physical Parameters **5-73**
 - How to Configure the Management Interface Speed and Duplex Parameters **5-73**
 - How to Configure the Duplex Operation of the Management Interface **5-73**
 - How to Configure the Speed of the Management Interface **5-74**
 - How to Monitor the Management Interface **5-75**

CHAPTER 6

Configuring the Line Interface 6-1

- How to Configure the Line Interfaces **6-1**
 - Line Interfaces **6-1**
 - How to Configure the Gigabit Ethernet Line Interfaces **6-2**
 - How to Configure the Fast Ethernet Line Interfaces **6-2**
- How to Configure Tunneling Protocols **6-3**
 - About Tunneling Protocols **6-3**
 - VPNs **6-4**
 - How to Select the Tunneling Mode **6-5**
 - How to Configure IP Tunnels **6-6**
 - How to Configure the VLAN Environment **6-6**
 - How to Configure the MPLS Environment **6-7**
 - How to Configure the L2TP Environment **6-8**
 - How to Display the Tunneling Configuration **6-8**

| | |
|--|------|
| How to Display the Logged-in VPNs | 6-9 |
| Options | 6-9 |
| How to Configure VLAN Translation | 6-9 |
| About VLAN Translation | 6-10 |
| VLAN Translation Features and Limitations | 6-10 |
| How to Set the VLAN Translation Constant | 6-11 |
| Options | 6-11 |
| Setting the VLAN Translation Constant: Example | 6-11 |
| How to Disable VLAN Translation | 6-12 |
| How to Monitor VLAN Translation | 6-12 |
| How to Configure Traffic Rules and Counters | 6-12 |
| Information About Traffic Rules and Counters | 6-12 |
| What are Traffic Rules and Counters? | 6-12 |
| Traffic Rules | 6-13 |
| Traffic Counters | 6-14 |
| How to Configure Traffic Counters | 6-14 |
| How to Create a Traffic Counter | 6-14 |
| How to Delete a Traffic Counter | 6-15 |
| How to Delete all Existing Traffic Counters | 6-15 |
| How to Configure Traffic Rules | 6-15 |
| How to Create a Traffic Rule | 6-15 |
| How to Delete a Traffic Rule | 6-18 |
| How to Delete all Traffic Rules | 6-18 |
| How to Manage Traffic Rules and Counters | 6-18 |
| How to View a Specified Traffic Rule | 6-18 |
| How to View all Traffic Rules | 6-18 |
| How to View a Specified Traffic Counter | 6-19 |
| How to View all Traffic Counters | 6-19 |
| How to Reset a Specified Traffic Counter | 6-19 |
| How to Reset all Traffic Counters | 6-19 |
| TOS Marking | 6-20 |
| How to Display the TOS Marking Configuration | 6-20 |
| How to Count Dropped Packets | 6-20 |
| About Counting Dropped Packets | 6-21 |
| How to Disable the Hardware Packet Drop | 6-21 |

CHAPTER 7**Configuring the Connection 7-1**

| | |
|--------------------------------------|-----|
| How to Configure the Connection Mode | 7-1 |
| Options | 7-2 |

| | |
|---|------|
| Configuring the Connection Mode: Examples | 7-3 |
| Example 1 | 7-3 |
| Example 2 | 7-3 |
| How to Monitor the Connection Mode | 7-3 |
| Monitoring the Connection Mode: Example | 7-3 |
| How to Configure the Link Mode | 7-4 |
| About the Link Mode | 7-4 |
| Options | 7-4 |
| How to Configure Asymmetric Routing Topology | 7-5 |
| About Asymmetric Routing Topology | 7-6 |
| Asymmetric Routing and Other Service Control Capabilities | 7-6 |
| Enabling Asymmetric Routing | 7-6 |
| How to Monitor Asymmetric Routing | 7-7 |
| Monitoring Asymmetric Routing: Example | 7-7 |
| How to Configure a Forced Failure | 7-7 |
| How to Force a Virtual Failure | 7-7 |
| How to Exit from a Virtual Failure | 7-8 |
| How to Configure the Failure Recovery Mode | 7-8 |
| Options | 7-8 |
| Configure the Failure Recovery Mode: Examples | 7-8 |
| Example 1 | 7-8 |
| Example 2 | 7-8 |
| How to Configure the SCE Platform/SM Connection | 7-9 |
| About the SCE Platform/SM Connection | 7-9 |
| How to Configure the Behavior of the SCE Platform in Case of Failure of the SM | 7-9 |
| Options | 7-9 |
| How to Configure the SM-SCE Platform Connection Timeout | 7-10 |
| Options | 7-10 |
| How to Enable and Disable Link Failure Reflection | 7-10 |
| About Link Failure Reflection | 7-10 |
| How to Enable Link Failure Reflection | 7-10 |
| How to Disable Link Failure Reflection | 7-11 |
| How to Enable and Disable Link Failure Reflection on All Ports | 7-11 |
| About Link Failure Reflection on All Ports | 7-11 |
| Options | 7-11 |
| How to Enable Link Failure Reflection on All Ports | 7-12 |
| How to Disable Link Failure Reflection on All Ports | 7-12 |
| How to Configure Link Failure Reflection in Linecard-Aware Mode (SCE 2000 only) | 7-12 |
| About Link Failure Reflection in Linecard-Aware Mode | 7-12 |

- How to Enable Linecard-Aware Mode 7-12
- How to Disable Linecard-Aware Mode 7-13

CHAPTER 8

- Raw Data Formatting: The RDR Formatter and NetFlow Exporting** 8-1
 - Information About the RDR Formatter and NetFlow Exporting Support 8-1
 - The RDR Formatter 8-2
 - NetFlow 8-2
 - NetFlow Terminology 8-2
 - NetFlow Exporting Support 8-3
 - Data Destinations 8-3
 - About Data Destinations 8-4
 - Categories 8-5
 - Priority 8-5
 - Setting DSCP for NetFlow 8-5
 - Forwarding Modes 8-6
 - Protocol 8-6
 - Transport Type 8-6
 - How to Configure Data Destinations and Categories 8-6
 - How to Configure a Data Destination 8-7
 - About Configuring the Data Destinations 8-7
 - Options 8-7
 - Configuring the Data Destinations: Examples 8-7
 - How to Configure the Data Categories 8-8
 - How to Configure a Destination and Assign Categories 8-9
 - How to Configure the Forwarding Mode 8-12
 - Options 8-13
 - Configuring the Forwarding Mode: Example 8-13
 - How to Configure the RDR Formatter 8-13
 - Options 8-13
 - How to Enable the RDR Formatter 8-13
 - How to Disable the RDR Formatter 8-14
 - How to Configure the Size of the RDR Formatter History Buffer 8-14
 - Options 8-14
 - How to Configure the NetFlow Exporting Support 8-14
 - Options 8-14
 - How to Configure a DSCP Value for NetFlow 8-15
 - Options 8-15
 - How to Configure the Template Refresh Interval 8-15
 - Options 8-15

- How to Configure Dynamic Mapping of RDRs to Categories 8-15
 - About Dynamic Mapping of RDRs to Categories 8-16
 - How to Configuring Mappings 8-16
 - Options 8-16
 - How to Add a Mapping to a Category 8-16
 - How to Remove a Mapping from a Category 8-16
 - How to Restore the Default Mapping for a Specified RDR Tag 8-17
- How to Display Data Destination Configuration and Statistics 8-17
 - How to the Display the Current RDR Formatter Configuration 8-17
 - Displaying the RDR Formatter Configuration: Example 8-18
 - How to the Display the Current RDR Formatter Statistics 8-18
 - Displaying the Current RDR Formatter Statistics: Example 8-19
- How to Disable the Linecard from Sending RDRs 8-20

CHAPTER 9

Managing Subscribers 9-1

- Information About Subscribers 9-1
 - What is a Subscriber? 9-2
 - Subscriber Modes in Service Control Solutions 9-3
 - Aging Subscribers 9-4
 - Synchronizing Subscriber Information in a Cascade System 9-4
 - Anonymous Groups and Subscriber Templates 9-5
 - Information About Subscriber Files 9-5
 - Subscriber Files 9-6
 - Subscriber default csv file format 9-6
 - Subscriber anonymous groups csv file format 9-7
- How to Import and Export Subscriber Information 9-7
 - Options 9-8
 - How to Import Subscriber Information 9-8
 - How to Export Subscriber Information 9-8
 - How to Import a Subscriber Template 9-8
 - How to Export a Subscriber Template 9-8
- How to Remove Subscribers and Templates 9-9
 - How to Remove a Specific Subscriber 9-9
 - Options 9-9
 - How to Remove All Introduced Subscribers 9-9
 - How to Remove a Specific Anonymous Subscriber 9-10
 - Options 9-10
 - How to Remove All Anonymous Subscriber Groups 9-10
 - How to Remove All Anonymous Subscribers 9-10

| | |
|--|------|
| How to Remove All Subscriber Templates | 9-10 |
| How to Remove VPN-based Subscribers | 9-11 |
| About VPN-based Subscribers | 9-11 |
| How to Remove Subscribers by Device | 9-11 |
| How to Remove Subscribers from the SM | 9-11 |
| How to Remove Subscribers from a Specified SCMP Peer Device | 9-11 |
| How to Import and Export Anonymous Groups | 9-12 |
| How to Import Anonymous Groups | 9-12 |
| Options | 9-12 |
| How to Export Anonymous Groups | 9-12 |
| Options | 9-12 |
| How to Monitor Subscribers | 9-13 |
| How to Monitor the Subscriber Database | 9-13 |
| How to Display the Subscriber Database Counters | 9-14 |
| Clearing the Subscriber Database Counters | 9-14 |
| How to Display Subscribers | 9-15 |
| Displaying Subscribers: All Current Subscriber Names | 9-15 |
| Displaying Subscribers: By Subscriber Property or Prefix | 9-16 |
| How to Display Subscribers: By Mapping (IP Address, VPN, VLAN ID, or MPLS/VPN) | 9-18 |
| How to Display Subscriber Information | 9-20 |
| How to display a listing of subscriber properties | 9-20 |
| How to display complete information for a specified subscriber | 9-21 |
| How to display values of subscriber properties for a specified subscriber | 9-21 |
| How to display mappings for a specified subscriber | 9-21 |
| How to display OS counters for a specified subscriber | 9-21 |
| How to Display Anonymous Subscriber Information | 9-22 |
| How to display currently configured anonymous groups | 9-22 |
| How to display currently configured templates for anonymous groups | 9-22 |
| How to display current configuration for a specified anonymous group | 9-22 |
| How to display subscribers in a specified anonymous group | 9-23 |
| How to display all subscribers currently in anonymous groups | 9-23 |
| How to display the number of subscribers in a specified anonymous group | 9-23 |
| How to display the total number of subscribers in all anonymous groups | 9-23 |
| Subscriber Traffic Processor IP Ranges | 9-23 |
| Information About Traffic Processor IP Ranges | 9-24 |
| Subscriber Mapping Modes | 9-25 |
| Subscriber Mapping Conflicts | 9-25 |
| Subscriber Rules for TIRs | 9-25 |
| How to Reserve Rules for TIRs | 9-26 |

- Options 9-26
- How to Configure TIRs 9-26
 - Options 9-26
 - How to Create or Update a TIR 9-27
 - How to update a TIR even if subscriber mappings exist 9-27
- How to Remove TIRs and Subscriber Mappings 9-27
 - How to Remove a Specified TIR 9-28
 - How to Remove All TIRs 9-28
 - How to Remove Mappings from a Specified TIR 9-28
 - How to Remove Mappings from a Specified IP Range 9-28
- How to Import and Export TIRs 9-28
 - About TIR csv Files 9-29
 - Options 9-29
 - How to Import TIRs from a csv File 9-29
 - How to Export TIRs to a csv File 9-29
- How to Monitor TIRs 9-30
 - How to Display Traffic Processor Mappings State 9-30
 - How to Display Configuration of a Specified TIR 9-30
 - How to Display Configuration of All TIRs 9-30
 - How to Display Mappings Related to a Specified TIR 9-30
 - How to Display the Number of Subscribers with Mappings Related to a Specified TIR 9-31
 - How to Display Complete Subscriber Information 9-31
 - How to Display All Subscribers Mapped to a Specified IP Range 9-31
 - How to Display the Number of Subscribers Mapped to a Specified IP Range 9-31
- How to Configure Subscriber Aging 9-31
 - How to Enable Aging for Anonymous Group Subscribers 9-31
 - How to Enable Aging for Introduced Subscribers 9-32
 - How to Disable Aging for Anonymous Group Subscribers 9-32
 - How to Disable Aging for Introduced Subscribers 9-32
 - How to Set the Aging Timeout Period for Anonymous Group Subscribers 9-32
 - Options 9-32
 - How to Set the Aging Timeout Period for Introduced Subscribers 9-33
 - Options 9-33
 - How to Display Aging for Anonymous Group Subscribers 9-33
 - How to Display Aging for Introduced Subscribers 9-33
- How to Configure the SCE Platform/SM Connection 9-33
 - Options 9-34
 - How to Configure the Behavior of the SCE Platform in Case of Failure of the SM 9-34
 - How to Configure the SM-SCE Platform Connection Timeout 9-34

CHAPTER 10**Redundancy and Fail-Over 10-1**

| | |
|---|-------|
| Information About Redundancy and Fail-Over | 10-1 |
| Terminology and Definitions | 10-2 |
| Redundant Topologies | 10-2 |
| In-line Dual Link Redundant Topology | 10-3 |
| Failure Detection | 10-4 |
| Link Failure Reflection | 10-4 |
| How to Configure Forced Failure | 10-5 |
| How to Force a Virtual Failure Condition | 10-5 |
| How to Exit a Virtual Failure Condition | 10-5 |
| Hot Standby and Fail-over | 10-5 |
| Hot Standby | 10-6 |
| Fail-over | 10-6 |
| Failure in the Cascade Connection | 10-7 |
| Installing a Cascaded System | 10-8 |
| Recovery | 10-9 |
| Replacing the SCE platform (manual recovery) | 10-9 |
| Manual steps: | 10-9 |
| Automatic steps (in parallel with the manual steps, requires no user intervention): | 10-10 |
| Reboot only (fully automatic recovery) | 10-10 |
| CLI Commands for Cascaded Systems | 10-10 |
| Topology-Related Parameters for Redundant Topologies | 10-11 |
| Configuring the Connection Mode | 10-11 |
| Examples | 10-11 |
| How to Monitor the System | 10-12 |
| How to View the Current Connection Mode | 10-12 |
| How to View the Current Link Mode | 10-12 |
| How to View Current Link Mappings | 10-12 |
| System Upgrades | 10-13 |
| Firmware Upgrade (package installation) | 10-13 |
| Application Upgrade | 10-13 |
| Simultaneous Upgrade of Firmware and Application | 10-14 |

CHAPTER 11**Identifying and Preventing Distributed-Denial-Of-Service Attacks 11-1**

| | |
|---|------|
| Information About Attack Filtering and Attack Detection | 11-1 |
| Attack Filtering | 11-2 |
| Specific Attack Filtering | 11-2 |
| Attack Detection | 11-3 |
| Attack Detection Thresholds | 11-4 |

- Attack Handling **11-5**
 - Subscriber Notification **11-5**
 - Hardware Filtering **11-6**
- How to Configure Attack Detectors **11-7**
 - About Configuring Attack Detectors **11-7**
 - How to Enable Specific-IP Detection **11-9**
 - Options **11-9**
 - How to Enable Specific-IP Detection **11-9**
 - How to Enable Specific-IP Detection for the TCP Protocol Only for all Attack Directions **11-9**
 - How to Enable Specific-IP Detection for the TCP Protocol for Port-based Detections Only for Dual-sided Attacks **11-10**
 - How to Disable Specific-IP Detection for Protocols Other than TCP, UDP, and ICMP for all Attack Directions **11-10**
 - How to Disable Specific-IP Detection for ICMP for Single-sided Attacks Defined by the Source IP **11-10**
 - How to Configure the Default Attack Detector **11-10**
 - Options **11-11**
 - How to Define the Default Action and Optionally the Default Thresholds **11-11**
 - How to Reinstate the System Defaults for a Selected Set of Attack Types **11-12**
 - How to Reinstate the System Defaults for All Attack Types **11-12**
- Specific Attack Detectors **11-13**
 - Options **11-13**
 - How to Enable a Specific Attack Detector and Assign it an ACL **11-14**
 - How to Define the Action and Optionally the Thresholds for a Specific Attack Detector **11-14**
 - How to Define the Subscriber Notification Setting for a Specific Attack Detector **11-15**
 - How to Define the SNMP Trap Setting for a Specific Attack Detector **11-15**
 - How to Define the List of Destination Ports for TCP or UDP Protocols for a Specific Attack Detector **11-15**
 - How to Delete User-Defined Values **11-16**
 - How to Disable a Specific Attack Detector **11-16**
 - How to Disable All Non-default Attack Detectors **11-16**
 - How to Disable All Attack Detectors **11-16**
- Sample Attack Detector Configuration **11-17**
- How to Configure Subscriber Notifications **11-17**
 - About Subscriber Notifications **11-18**
 - How to Configure the Subscriber Notification Port **11-18**
 - Options **11-18**
 - How to Remove the Subscriber Notification Port **11-18**
- Preventing and Forcing Attack Detection **11-19**
 - Options **11-19**

| | |
|---|-------|
| Preventing Attack Filtering | 11-20 |
| How to Configure a don't-filter Setting for a Specified Situation | 11-20 |
| How to Remove a don't-filter Setting from a Specified Situation | 11-20 |
| How to Remove All don't-filter Settings | 11-20 |
| Forcing Attack Filtering | 11-20 |
| How to Configure a force-filter Setting for a Specified Situation | 11-21 |
| How to Remove a force-filter Setting from a Specified Situation | 11-21 |
| How to Remove All force-filter Settings | 11-21 |
| How to Monitor Attack Filtering | 11-21 |
| Monitoring Attack Filtering Using SNMP Traps | 11-22 |
| Monitoring Attack Filtering Using CLI Commands | 11-23 |
| How to display a specified attack detector configuration | 11-24 |
| How to display the default attack detector configuration | 11-26 |
| How to display all attack detector configurations | 11-26 |
| How to display filter state (enabled or disabled) | 11-27 |
| How to display configured threshold values and actions | 11-27 |
| How to display the current counters | 11-29 |
| How to display all currently handled attacks | 11-29 |
| How to display all existing force-filter settings | 11-29 |
| How to display all existing don't-filter settings | 11-29 |
| How to display the list of ports selected for subscriber notification | 11-29 |
| How to find out whether hardware attack filtering has been activated | 11-30 |
| Viewing the Attack Log | 11-30 |
| The Attack Log | 11-30 |
| How to View the Attack Log | 11-31 |
| How to Copy the Attack Log to a File | 11-31 |

CHAPTER 12**Value Added Services (VAS) Traffic Forwarding 12-1**

| | |
|--|------|
| Information About VAS Traffic Forwarding | 12-1 |
| VAS Service Goals | 12-2 |
| How VAS Traffic Forwarding Works | 12-2 |
| Requirements for VAS Servers | 12-4 |
| VAS Traffic Forwarding and SCA BB | 12-5 |
| VLAN Tags for VAS Traffic Forwarding | 12-5 |
| Service Flow | 12-6 |
| Data Flow | 12-6 |
| Non-VAS Data Flow | 12-7 |
| VAS Data Flow | 12-8 |
| Load Balancing | 12-8 |

- Load Balancing and Subscribers 12-9
 - Load Balancing and Subscriber Mode 12-9
 - VAS Redundancy 12-9
 - VAS Server Failure 12-10
 - VAS Server Group Failure 12-10
 - Ethernet Switch Failure 12-10
 - Disabling a VAS Server 12-11
 - VAS Status and VAS Health Check 12-11
 - VAS Server States 12-12
 - VAS Traffic Forwarding Topologies 12-12
 - Single SCE Platform, Multiple VAS Servers 12-12
 - Data Flow 12-13
 - Multiple SCE Platforms, Multiple VAS Servers 12-14
 - SNMP Support for VAS 12-15
 - Interactions Between VAS Traffic Forwarding and Other SCE Platform Features 12-15
 - Incompatible SCE Platform Features 12-15
 - VAS Traffic Forwarding and DDoS Processing 12-15
 - Specific IP DDoS Attack Detection 12-15
 - Specific IP Attack filter 12-16
 - VAS Traffic Forwarding and Bandwidth Management 12-16
 - Global Controllers and VAS flows 12-16
 - Configuring VAS Traffic Forwarding 12-16
 - Configuring VAS Traffic Forwarding from the SCA BB Console 12-17
 - Global Options 12-17
 - Enabling VAS Traffic Forwarding 12-17
 - Options 12-18
 - Disabling VAS Traffic Forwarding 12-18
 - How to Configure the VAS Traffic Link 12-19
 - Options 12-19
 - How to Select the Link for VAS Traffic 12-19
 - How to Revert to the Default Link for VAS Traffic 12-19
 - How to Configure a VAS Server 12-19
 - Options 12-20
 - How to Enable a VAS Server 12-20
 - How to Disable a VAS Server 12-20
 - How to Restore all VAS Server Properties to Default 12-20
 - How to Assign a VLAN ID to a VAS Server 12-21
 - Options 12-21
 - How to Configure the VLAN Tag Number for a Specified VAS Server 12-21

| | |
|--|-------|
| How to Remove the VLAN Tag Number from a Specified VAS Server | 12-21 |
| How to Configure the Health Check | 12-21 |
| How to Configure Pseudo IP Addresses for the Health Check Packets | 12-23 |
| How to Configure a VAS Server Group | 12-24 |
| About VAS Server Groups | 12-24 |
| How to Add and Remove Servers | 12-25 |
| How to Configure VAS Server Group Failure Parameters | 12-25 |
| How to Monitor VAS Traffic Forwarding | 12-27 |
| How to Display Global VAS Status and Configuration | 12-27 |
| Example | 12-27 |
| How to Display Operational and Configuration Information for a Specific VAS Server Group | 12-28 |
| Example | 12-28 |
| How to Display Operational and Configuration Information for All VAS Server Groups | 12-28 |
| How to Display Operational and Configuration Information for a Specific VAS Server | 12-28 |
| Example | 12-28 |
| How to Display Operational and Configuration Information for All VAS Servers | 12-29 |
| How to Display the VAS Servers Used by a Specified Subscriber | 12-29 |
| How to Display Health Check Counters for a Specified VAS Server | 12-29 |
| Example | 12-29 |
| How to Display Health Check Counters for All VAS Servers | 12-30 |
| How to Clear the Health Check Counters for a Specified VAS Server | 12-30 |
| How to Clear the Health Check Counters for All VAS Servers | 12-30 |
| How to Display Bandwidth per VAS Server and VAS Direction | 12-30 |
| Example | 12-31 |
| VAS over 10G | 12-31 |
| About VAS over 10G | 12-32 |
| Data Flow in VAS over 10G Topology | 12-33 |
| VAS Data Flow: To the VAS Server | 12-35 |
| VAS Data Flow: From the VAS Server | 12-36 |
| Failover Support | 12-37 |
| Health Check in VAS over 10G Topology | 12-38 |
| Configuring VAS over 10G: General Guidelines | 12-39 |
| Configuring the 7600/6500 for VAS over 10G | 12-39 |
| How to Configure VAS over 10G | 12-40 |
| How to Configure the VAS Traffic Link Auto-Select Parameters (VAS over 10G) | 12-40 |
| How to Configure the Minimum Time between Link Switches | 12-41 |
| How to Set the Active VAS Link | 12-42 |
| How to Configure Health Check for VAS over 10G | 12-42 |
| How to Configure the Health Check IP Address | 12-43 |
| How to Remove the IP Address Configuration | 12-44 |

How to Enable the Health Check for VAS over 10G Topology Options 12-44

Options 12-44

How to Enable Health Check Compatibility for VAS over 10G (MGSCP) 12-44

How to Remove the Health Check Compatibility Configuration 12-44

VAS Over 10G Sample Configuration 12-45

CHAPTER 13

MPLS/VPN Support 13-1

Service Control in the MPLS/VPN Environment 13-2

Definitions and Acronyms 13-3

What are the Challenges for Service Control for MPLS/VPN Support? 13-3

How MPLS/VPN Support Works 13-4

Flow Detection 13-4

VPN Detection 13-4

Subscriber Detection 13-5

What is an MPLS/VPN-based Subscriber? 13-5

Private IP Subscriber Support 13-5

How the Service Control MPLS/VPN Solution Works 13-5

How the Service Control MPLS/VPN Solution Works: A Summary 13-6

SCE Platform Tasks in the MPLS/VPN Solution 13-6

BGP LEG Tasks in the MPLS/VPN Solution 13-6

SM Tasks in the MPLS/VPN Solution 13-6

Service Control MPLS/VPN Concepts 13-7

Non-VPN-Based Subscribers 13-7

Bypassing Unknown VPNs 13-7

Additional MPLS Pattern Support 13-8

VPN Identifier (RD or RT) 13-8

Service Control MPLS/VPN Requirements 13-9

Topology 13-9

Capacity 13-10

Limitations 13-10

Backwards Compatibility 13-11

How to Configure MPLS/VPN Support 13-11

Configuring the MPLS Environment 13-12

How to Check the Running Configuration 13-12

How to Configure the MPLS Environment 13-12

How to Configure the SCE Platform for MPLS/VPN Support 13-13

About Configuring the SCE Platform for MPLS/VPN Support 13-13

How to Define the PE Routers 13-13

How to Configure the MAC Resolver 13-14

| | |
|--|-------|
| How to Monitor the MAC Resolver | 13-16 |
| How to Configure the SM for MPLS/VPN Support | 13-16 |
| Configuring the SM for MPLS/VPN Support | 13-16 |
| How to Edit the SM Configuration File | 13-16 |
| How to Configure the SM to Allow IP Ranges | 13-17 |
| How to Manage MPLS/VPN Support | 13-17 |
| How to Manage MPLS/VPN Support via SNMP | 13-18 |
| MPLS/VPN MIB Objects | 13-18 |
| MPLS/VPN Traps | 13-18 |
| How to Monitor MPLS/VPN Support via SCE Platform CLI | 13-18 |
| How to Display VPN-related Mappings | 13-19 |
| How to Clear Upstream VPN Mappings | 13-21 |
| How to Monitor Subscriber Counters | 13-22 |
| How to Monitor MPLS/VPN Counters | 13-23 |
| How to Monitor the PE Routers | 13-23 |
| How to Monitor Bypassed VPNs | 13-24 |
| How to Monitor Non-VPN Mappings | 13-24 |
| How to Manage MPLS/VPN Support via SM CLU | 13-24 |
| Managing VPNs | 13-25 |
| How to Add Mappings to VPN-based Subscribers | 13-26 |
| How to Remove VPN Mappings from Subscribers | 13-28 |
| How to Monitor Subscriber MPLS/VPN Mappings | 13-29 |

CHAPTER 14**Managing the SCMP 14-1**

| | |
|---|------|
| About SCMP | 14-1 |
| SCMP Terminology | 14-2 |
| Deployment Scenarios | 14-3 |
| Single ISG Router with a Single SCE Platform (1xISG – 1xSCE) | 14-3 |
| Single ISG Router with Two Cascaded SCE Platforms (1xISG – 2xSCE) | 14-4 |
| Multiple ISG Routers with Two Cascaded SCE Platforms (NxISG – 2xSCE) | 14-5 |
| Multiple ISG Routers with Multiple SCE Platforms via Load Balancing (NxISG – MxSCE) | 14-6 |
| SCMP Peer Devices | 14-7 |
| Connection Management | 14-7 |
| SCMP Subscriber Management | 14-8 |
| GUID and Subscriber ID | 14-8 |
| How to Configure the SCMP | 14-8 |
| How to Configure SCMP Parameters | 14-9 |
| How to Enable the SCMP | 14-9 |
| How to Disable the SCMP | 14-9 |

- How to Configure the SCMP Peer Device to Push Sessions 14-9
- How to Configure the SCMP Peer Device to Force Each Subscriber to Single SCE Platform 14-10
- How to Define the Keep-alive Interval Parameter 14-10
- How to Define the Reconnect Interval Parameter 14-11
- How to Define the Loss-of-Sync Timeout Parameter 14-11
- How to Add an SCMP Peer Device 14-11
 - How to Define an SCMP Peer Device 14-12
 - How to Assign the SCMP Peer Device to an Anonymous Group 14-12
- How to Delete Subscribers Managed by an SCMP Peer Device 14-13
 - Options 14-13
- How to Delete an SCMP Peer Device 14-13
- How to Define the Subscriber ID 14-13
 - Options 14-14
- How to Configure the RADIUS Client 14-14
 - Options 14-15
- How to Monitor the SCMP Environment 14-15
 - How to Monitor the SCMP 14-15
 - Options 14-15
 - How to display the general SCMP configuration 14-16
 - How to display the configuration all currently defined SCMP peer devices 14-16
 - How to display the configuration for a specified SCMP peer device 14-16
 - How to display the statistics for all SCMP peer devices 14-16
 - How to display the statistics for a specified SCMP peer device 14-17
 - How to Monitor the RADIUS Client 14-17
- SCE Platform Utilization Indicators A-2
 - CPU Utilization A-2
 - Flows Capacity A-2
 - Subscribers Capacity A-2
- Service Loss A-3
 - Monitoring Service Loss A-3
- pcube Enterprise MIB B-2
 - Application MIB Integration B-3
 - Application and Subscriber groups B-4
 - The Engage MIB (pcubeEngageMIB) B-5
 - Using this Reference B-6
- pcubeModules (1.3.6.1.4.1.5655.2) B-6
 - pcubeSeMIB (1.3.6.1.4.1.5655.2.3) B-6
 - pcubeSeMIB Object Groups (1.3.6.1.4.1.5655.2.3.1.1) B-7
 - pcubeCompliances (1.3.6.1.4.1.5655.2.3.1.2) B-15

| | |
|--|-------------|
| pcubeWorkgroup (1.3.6.1.4.1.5655.4) | B-16 |
| Notification Types | B-16 |
| operationalStatusOperationalTrap (1.3.6.1.4.1.5655.4.0.1) | B-18 |
| operationalStatusWarningTrap (1.3.6.1.4.1.5655.4.0.2) | B-18 |
| operationalStatusFailureTrap (1.3.6.1.4.1.5655.4.0.3) | B-18 |
| systemResetTrap (1.3.6.1.4.1.5655.4.0.4) | B-18 |
| chassisTempAlarmOnTrap (1.3.6.1.4.1.5655.4.0.5) | B-18 |
| chassisTempAlarmOffTrap (1.3.6.1.4.1.5655.4.0.6) | B-19 |
| chassisVoltageAlarmOnTrap (1.3.6.1.4.1.5655.4.0.7) | B-19 |
| chassisFansAlarmOnTrap (1.3.6.1.4.1.5655.4.0.8) | B-19 |
| chassisPowerSupplyAlarmOnTrap (1.3.6.1.4.1.5655.4.0.9) | B-19 |
| rdrActiveConnectionTrap (1.3.6.1.4.1.5655.4.0.10) | B-19 |
| rdrNoActiveConnectionTrap (1.3.6.1.4.1.5655.4.0.11) | B-19 |
| rdrConnectionUpTrap (1.3.6.1.4.1.5655.4.0.12) | B-19 |
| rdrConnectionDownTrap (1.3.6.1.4.1.5655.4.0.13) | B-19 |
| loggerUserLogIsFullTrap (1.3.6.1.4.1.5655.4.0.18) | B-19 |
| snmpClockDriftWarnTrap (1.3.6.1.4.1.5655.4.0.19) | B-19 |
| linkModeBypassTrap (1.3.6.1.4.1.5655.4.0.20) | B-20 |
| linkModeForwardingTrap (1.3.6.1.4.1.5655.4.0.21) | B-20 |
| linkModeCutoffTrap (1.3.6.1.4.1.5655.4.0.22) | B-20 |
| moduleAttackFilterActivatedTrap (1.3.6.1.4.1.5655.4.0.25) | B-20 |
| moduleAttackFilterDeactivatedTrap (1.3.6.1.4.1.5655.4.0.26) | B-21 |
| moduleEmAgentGenericTrap (1.3.6.1.4.1.5655.4.0.27) | B-21 |
| linkModeSniffingTrap (1.3.6.1.4.1.5655.4.0.28) | B-21 |
| moduleRedundancyReadyTrap (1.3.6.1.4.1.5655.4.0.29) | B-21 |
| moduleRedundantConfigurationMismatchTrap (1.3.6.1.4.1.5655.4.0.30) | B-21 |
| moduleLostRedundancyTrap (1.3.6.1.4.1.5655.4.0.31) | B-21 |
| moduleSmConnectionDownTrap (1.3.6.1.4.1.5655.4.0.32) | B-22 |
| moduleSmConnectionUpTrap (1.3.6.1.4.1.5655.4.0.33) | B-22 |
| moduleOperStatusChangeTrap (1.3.6.1.4.1.5655.4.0.34) | B-22 |
| portOperStatusChangeTrap (1.3.6.1.4.1.5655.4.0.35) | B-22 |
| chassisLineFeedAlarmOnTrap (1.3.6.1.4.1.5655.4.0.36) | B-22 |
| rdrFormatterCategoryDiscardingReportsTrap (1.3.6.1.4.1.5655.4.0.37) | B-22 |
| rdrFormatterCategoryStoppedDiscardingReportsTrap (1.3.6.1.4.1.5655.4.0.38) | B-22 |
| sessionStartedTrap (1.3.6.1.4.1.5655.4.0.39) | B-22 |
| sessionEndedTrap (1.3.6.1.4.1.5655.4.0.40) | B-22 |
| sessionDeniedAccessTrap (1.3.6.1.4.1.5655.4.0.41) | B-22 |
| sessionBadLoginTrap (1.3.6.1.4.1.5655.4.0.42) | B-23 |
| illegalSubscriberMappingTrap (1.3.6.1.4.1.5655.4.0.43) | B-23 |
| loggerLineAttackLogFullTrap (1.3.6.1.4.1.5655.4.0.44) | B-23 |

| | |
|--|-------------|
| vasServerOperationalStatusChangeTrap (1.3.6.1.4.1.5655.4.0.45) | B-23 |
| pullRequestNumber (1.3.6.1.4.1.5655.4.0.46) | B-23 |
| pullRequestRetryFailedTrap (1.3.6.1.4.1.5655.4.0.47) | B-23 |
| mplsVpnTotalHWMappingsThresholdExceededTrap (1.3.6.1.4.1.5655.4.0.48) | B-23 |
| pcubeSe Objects | B-24 |
| sysOperationalStatus (1.3.6.1.4.1.5655.4.1.1.1) | B-30 |
| sysFailureRecovery (1.3.6.1.4.1.5655.4.1.1.2) | B-30 |
| sysVersion (1.3.6.1.4.1.5655.4.1.1.3) | B-31 |
| pchassisSysType (1.3.6.1.4.1.5655.4.1.2.1) | B-31 |
| pchassisPowerSupplyAlarm (1.3.6.1.4.1.5655.4.1.2.2) | B-31 |
| pchassisFansAlarm (1.3.6.1.4.1.5655.4.1.2.3) | B-31 |
| pchassisTempAlarm (1.3.6.1.4.1.5655.4.1.2.4) | B-32 |
| pchassisVoltageAlarm (1.3.6.1.4.1.5655.4.1.2.5) | B-32 |
| pchassisNumSlots (1.3.6.1.4.1.5655.4.1.2.6) | B-32 |
| pchassisSlotConfig (1.3.6.1.4.1.5655.4.1.2.7) | B-33 |
| pchassisPsuType (1.3.6.1.4.1.5655.4.1.2.8) | B-33 |
| pchassisLineFeedAlarm (1.3.6.1.4.1.5655.4.1.2.9) | B-33 |
| pmoduleTable (1.3.6.1.4.1.5655.4.1.3.1) | B-34 |
| pmoduleEntry (1.3.6.1.4.1.5655.4.1.3.1.1) | B-34 |
| pmoduleIndex (1.3.6.1.4.1.5655.4.1.3.1.1.1) | B-34 |
| pmoduleType (1.3.6.1.4.1.5655.4.1.3.1.1.2) | B-35 |
| pmoduleNumTrafficProcessors (1.3.6.1.4.1.5655.4.1.3.1.1.3) | B-35 |
| pmoduleSlotNum (1.3.6.1.4.1.5655.4.1.3.1.1.4) | B-35 |
| pmoduleHwVersion (1.3.6.1.4.1.5655.4.1.3.1.1.5) | B-36 |
| pmoduleNumPorts (1.3.6.1.4.1.5655.4.1.3.1.1.6) | B-36 |
| pmoduleNumLinks (1.3.6.1.4.1.5655.4.1.3.1.1.7) | B-36 |
| pmoduleConnectionMode (1.3.6.1.4.1.5655.4.1.3.1.1.8) | B-36 |
| pmoduleSerialNumber (1.3.6.1.4.1.5655.4.1.3.1.1.9) | B-37 |
| pmoduleUpStreamAttackFilteringTime (1.3.6.1.4.1.5655.4.1.3.1.1.10) | B-37 |
| pmoduleUpStreamLastAttackFilteringTime (1.3.6.1.4.1.5655.4.1.3.1.1.11) | B-37 |
| pmoduleDownStreamAttackFilteringTime (1.3.6.1.4.1.5655.4.1.3.1.1.12) | B-37 |
| pmoduleDownStreamLastAttackFilteringTime (1.3.6.1.4.1.5655.4.1.3.1.1.13) | B-38 |
| pmoduleAttackObjectsClearTime (1.3.6.1.4.1.5655.4.1.3.1.1.14) | B-38 |
| pmoduleAdminStatus (1.3.6.1.4.1.5655.4.1.3.1.1.15) | B-38 |
| pmoduleOperStatus (1.3.6.1.4.1.5655.4.1.3.1.1.16) | B-38 |
| linkTable (1.3.6.1.4.1.5655.4.1.4.1) | B-39 |
| linkEntry (1.3.6.1.4.1.5655.4.1.4.1.1) | B-39 |
| linkModuleIndex (1.3.6.1.4.1.5655.4.1.4.1.1.1) | B-39 |
| linkIndex (1.3.6.1.4.1.5655.4.1.4.1.1.2) | B-40 |
| linkAdminModeOnActive (1.3.6.1.4.1.5655.4.1.4.1.1.3) | B-40 |

| | |
|---|-------------|
| linkAdminModeOnFailure (1.3.6.1.4.1.5655.4.1.4.1.1.4) | B-40 |
| linkOperMode (1.3.6.1.4.1.5655.4.1.4.1.1.5) | B-41 |
| linkStatusReflectionEnable (1.3.6.1.4.1.5655.4.1.4.1.1.6) | B-41 |
| linkSubscriberSidePortIndex (1.3.6.1.4.1.5655.4.1.4.1.1.7) | B-41 |
| linkNetworkSidePortIndex (1.3.6.1.4.1.5655.4.1.4.1.1.8) | B-42 |
| diskNumUsedBytes (1.3.6.1.4.1.5655.4.1.5.1) | B-42 |
| diskNumFreeBytes (1.3.6.1.4.1.5655.4.1.5.2) | B-42 |
| rdrFormatterEnable (1.3.6.1.4.1.5655.4.1.6.1) | B-42 |
| rdrFormatterDestTable (1.3.6.1.4.1.5655.4.1.6.2) | B-43 |
| rdrFormatterDestEntry (1.3.6.1.4.1.5655.4.1.6.2.1) | B-43 |
| rdrFormatterDestIPAddr (1.3.6.1.4.1.5655.4.1.6.2.1.1) | B-43 |
| rdrFormatterDestPort (1.3.6.1.4.1.5655.4.1.6.2.1.2) | B-44 |
| rdrFormatterDestPriority (1.3.6.1.4.1.5655.4.1.6.2.1.3) | B-44 |
| rdrFormatterDestStatus (1.3.6.1.4.1.5655.4.1.6.2.1.4) | B-44 |
| rdrFormatterDestConnectionStatus (1.3.6.1.4.1.5655.4.1.6.2.1.5) | B-45 |
| rdrFormatterDestNumReportsSent (1.3.6.1.4.1.5655.4.1.6.2.1.6) | B-45 |
| rdrFormatterDestNumReportsDiscarded (1.3.6.1.4.1.5655.4.1.6.2.1.7) | B-45 |
| rdrFormatterDestReportRate (1.3.6.1.4.1.5655.4.1.6.2.1.8) | B-45 |
| rdrFormatterDestReportRatePeak (1.3.6.1.4.1.5655.4.1.6.2.1.9) | B-46 |
| rdrFormatterDestReportRatePeakTime (1.3.6.1.4.1.5655.4.1.6.2.1.10) | B-46 |
| rdrFormatterNumReportsSent (1.3.6.1.4.1.5655.4.1.6.3) | B-46 |
| rdrFormatterNumReportsDiscarded (1.3.6.1.4.1.5655.4.1.6.4) | B-46 |
| rdrFormatterClearCountersTime (1.3.6.1.4.1.5655.4.1.6.5) | B-47 |
| rdrFormatterReportRate (1.3.6.1.4.1.5655.4.1.6.6) | B-47 |
| rdrFormatterReportRatePeak (1.3.6.1.4.1.5655.4.1.6.7) | B-47 |
| rdrFormatterReportRatePeakTime (1.3.6.1.4.1.5655.4.1.6.8) | B-47 |
| rdrFormatterProtocol (1.3.6.1.4.1.5655.4.1.6.9) | B-48 |
| rdrFormatterForwardingMode (1.3.6.1.4.1.5655.4.1.6.10) | B-48 |
| rdrFormatterCategoryTable (1.3.6.1.4.1.5655.4.1.6.11) | B-48 |
| rdrFormatterCategoryEntry (1.3.6.1.4.1.5655.4.1.6.11.1) | B-49 |
| rdrFormatterCategoryIndex (1.3.6.1.4.1.5655.4.1.6.11.1.1) | B-49 |
| rdrFormatterCategoryName (1.3.6.1.4.1.5655.4.1.6.11.1.2) | B-49 |
| rdrFormatterCategoryNumReportsSent (1.3.6.1.4.1.5655.4.1.6.11.1.3) | B-49 |
| rdrFormatterCategoryNumReportsDiscarded (1.3.6.1.4.1.5655.4.1.6.11.1.4) | B-50 |
| rdrFormatterCategoryReportRate (1.3.6.1.4.1.5655.4.1.6.11.1.5) | B-50 |
| rdrFormatterCategoryReportRatePeak (1.3.6.1.4.1.5655.4.1.6.11.1.6) | B-50 |
| rdrFormatterCategoryReportRatePeakTime (1.3.6.1.4.1.5655.4.1.6.11.1.7) | B-50 |
| rdrFormatterCategoryNumReportsQueued (1.3.6.1.4.1.5655.4.1.6.11.1.8) | B-51 |
| rdrFormatterCategoryDestTable (1.3.6.1.4.1.5655.4.1.6.12) | B-51 |
| rdrFormatterCategoryDestEntry (1.3.6.1.4.1.5655.4.1.6.12.1) | B-51 |

[rdrFormatterCategoryDestPriority \(1.3.6.1.4.1.5655.4.1.6.12.1.1\)](#) **B-51**
[rdrFormatterCategoryDestStatus \(1.3.6.1.4.1.5655.4.1.6.12.1.2\)](#) **B-52**
[loggerUserLogEnable \(1.3.6.1.4.1.5655.4.1.7.1\)](#) **B-52**
[loggerUserLogNumInfo \(1.3.6.1.4.1.5655.4.1.7.2\)](#) **B-52**
[loggerUserLogNumWarning \(1.3.6.1.4.1.5655.4.1.7.3\)](#) **B-53**
[loggerUserLogNumError \(1.3.6.1.4.1.5655.4.1.7.4\)](#) **B-53**
[loggerUserLogNumFatal \(1.3.6.1.4.1.5655.4.1.7.5\)](#) **B-53**
[loggerUserLogClearCountersTime \(1.3.6.1.4.1.5655.4.1.7.6\)](#) **B-53**
[subscribersInfoTable \(1.3.6.1.4.1.5655.4.1.8.1\)](#) **B-54**
[subscribersInfoEntry \(1.3.6.1.4.1.5655.4.1.8.1.1\)](#) **B-54**
[subscribersNumIntroduced \(1.3.6.1.4.1.5655.4.1.8.1.1.1\)](#) **B-54**
[subscribersNumFree \(1.3.6.1.4.1.5655.4.1.8.1.1.2\)](#) **B-55**
[subscribersNumIpAddrMappings \(1.3.6.1.4.1.5655.4.1.8.1.1.3\)](#) **B-55**
[subscribersNumIpAddrMappingsFree \(1.3.6.1.4.1.5655.4.1.8.1.1.4\)](#) **B-55**
[subscribersNumIpRangeMappings \(1.3.6.1.4.1.5655.4.1.8.1.1.5\)](#) **B-55**
[subscribersNumIpRangeMappingsFree \(1.3.6.1.4.1.5655.4.1.8.1.1.6\)](#) **B-56**
[subscribersNumVlanMappings \(1.3.6.1.4.1.5655.4.1.8.1.1.7\)](#) **B-56**
[subscribersNumVlanMappingsFree \(1.3.6.1.4.1.5655.4.1.8.1.1.8\)](#) **B-56**
[subscribersNumActive \(1.3.6.1.4.1.5655.4.1.8.1.1.9\)](#) **B-56**
[subscribersNumActivePeak \(1.3.6.1.4.1.5655.4.1.8.1.1.10\)](#) **B-57**
[subscribersNumActivePeakTime \(1.3.6.1.4.1.5655.4.1.8.1.1.11\)](#) **B-57**
[subscribersNumUpdates \(1.3.6.1.4.1.5655.4.1.8.1.1.12\)](#) **B-57**
[subscribersCountersClearTime \(1.3.6.1.4.1.5655.4.1.8.1.1.13\)](#) **B-57**
[subscribersNumTmplRangeMappings \(1.3.6.1.4.1.5655.4.1.8.1.1.14\)](#) **B-58**
[subscribersNumTmplRangeMappingsFree \(1.3.6.1.4.1.5655.4.1.8.1.1.15\)](#) **B-58**
[subscribersNumAnonymous \(1.3.6.1.4.1.5655.4.1.8.1.1.16\)](#) **B-58**
[subscribersNumWithSessions \(1.3.6.1.4.1.5655.4.1.8.1.1.17\)](#) **B-58**
[subscribersPropertiesTable \(1.3.6.1.4.1.5655.4.1.8.2\)](#) **B-59**
[subscribersPropertiesEntry \(1.3.6.1.4.1.5655.4.1.8.2.1\)](#) **B-59**
[spIndex \(1.3.6.1.4.1.5655.4.1.8.2.1.1\)](#) **B-59**
[spName \(1.3.6.1.4.1.5655.4.1.8.2.1.2\)](#) **B-59**
[spType \(1.3.6.1.4.1.5655.4.1.8.2.1.3\)](#) **B-60**
[subscriberPropertiesValuesTable \(1.3.6.1.4.1.5655.4.1.8.3\)](#) **B-60**
[subscriberPropertiesValueEntry \(1.3.6.1.4.1.5655.4.1.8.3.1\)](#) **B-60**
[spvIndex \(1.3.6.1.4.1.5655.4.1.8.3.1.1\)](#) **B-61**
[spvSubName \(1.3.6.1.4.1.5655.4.1.8.3.1.2\)](#) **B-61**
[spvPropertyName \(1.3.6.1.4.1.5655.4.1.8.3.1.3\)](#) **B-61**
[spvRowStatus \(1.3.6.1.4.1.5655.4.1.8.3.1.4\)](#) **B-61**
[spvPropertyStringValue \(1.3.6.1.4.1.5655.4.1.8.3.1.5\)](#) **B-62**
[spvPropertyUintValue \(1.3.6.1.4.1.5655.4.1.8.3.1.6\)](#) **B-62**

| | |
|--|-------------|
| spvPropertyCounter64Value (1.3.6.1.4.1.5655.4.1.8.3.1.7) | B-62 |
| tpInfoTable (1.3.6.1.4.1.5655.4.1.9.1) | B-62 |
| tpInfoEntry (1.3.6.1.4.1.5655.4.1.9.1.1) | B-63 |
| tpModuleIndex (1.3.6.1.4.1.5655.4.1.9.1.1.1) | B-63 |
| tpIndex (1.3.6.1.4.1.5655.4.1.9.1.1.2) | B-63 |
| tpTotalNumHandledPackets (1.3.6.1.4.1.5655.4.1.9.1.1.3) | B-64 |
| tpTotalNumHandledFlows (1.3.6.1.4.1.5655.4.1.9.1.1.4) | B-64 |
| tpNumActiveFlows (1.3.6.1.4.1.5655.4.1.9.1.1.5) | B-64 |
| tpNumActiveFlowsPeak (1.3.6.1.4.1.5655.4.1.9.1.1.6) | B-64 |
| tpNumActiveFlowsPeakTime (1.3.6.1.4.1.5655.4.1.9.1.1.7) | B-65 |
| tpNumTcpActiveFlows (1.3.6.1.4.1.5655.4.1.9.1.1.8) | B-65 |
| TpNumTcpActiveFlowsPeak (1.3.6.1.4.1.5655.4.1.9.1.1.9) | B-65 |
| tpNumTcpActiveFlowsPeakTime (1.3.6.1.4.1.5655.4.1.9.1.1.10) | B-65 |
| tpNumUdpActiveFlows (1.3.6.1.4.1.5655.4.1.9.1.1.11) | B-66 |
| tpNumUdpActiveFlowsPeak (1.3.6.1.4.1.5655.4.1.9.1.1.12) | B-66 |
| tpNumUdpActiveFlowsPeakTime (1.3.6.1.4.1.5655.4.1.9.1.1.13) | B-66 |
| tpNumNonTcpUdpActiveFlows (1.3.6.1.4.1.5655.4.1.9.1.1.14) | B-66 |
| tpNumNonTcpUdpActiveFlowsPeak (1.3.6.1.4.1.5655.4.1.9.1.1.15) | B-67 |
| tpNumNonTcpUdpActiveFlowsPeakTime (1.3.6.1.4.1.5655.4.1.9.1.1.16) | B-67 |
| tpTotalNumBlockedPackets (1.3.6.1.4.1.5655.4.1.9.1.1.17) | B-67 |
| tpTotalNumBlockedFlows (1.3.6.1.4.1.5655.4.1.9.1.1.18) | B-67 |
| tpTotalNumDiscardedPacketsDueToBwLimit (1.3.6.1.4.1.5655.4.1.9.1.1.19) | B-68 |
| tpTotalNumWredDiscardedPackets (1.3.6.1.4.1.5655.4.1.9.1.1.20) | B-68 |
| tpTotalNumFragments (1.3.6.1.4.1.5655.4.1.9.1.1.21) | B-68 |
| tpTotalNumNonIpPackets (1.3.6.1.4.1.5655.4.1.9.1.1.22) | B-68 |
| tpTotalNumIpCrcErrPackets (1.3.6.1.4.1.5655.4.1.9.1.1.23) | B-69 |
| tpTotalNumIpLengthErrPackets (1.3.6.1.4.1.5655.4.1.9.1.1.24) | B-69 |
| tpTotalNumIpBroadcastPackets (1.3.6.1.4.1.5655.4.1.9.1.1.25) | B-69 |
| tpTotalNumTtlErrPackets (1.3.6.1.4.1.5655.4.1.9.1.1.26) | B-69 |
| tpTotalNumTcpUdpCrcErrPackets (1.3.6.1.4.1.5655.4.1.9.1.1.27) | B-70 |
| tpClearCountersTime (1.3.6.1.4.1.5655.4.1.9.1.1.28) | B-70 |
| tpHandledPacketsRate (1.3.6.1.4.1.5655.4.1.9.1.1.29) | B-70 |
| tpHandledPacketsRatePeak (1.3.6.1.4.1.5655.4.1.9.1.1.30) | B-70 |
| tpHandledPacketsRatePeakTime (1.3.6.1.4.1.5655.4.1.9.1.1.31) | B-71 |
| tpHandledFlowsRate (1.3.6.1.4.1.5655.4.1.9.1.1.32) | B-71 |
| tpHandledFlowsRatePeak (1.3.6.1.4.1.5655.4.1.9.1.1.33) | B-71 |
| tpHandledFlowsRatePeakTime (1.3.6.1.4.1.5655.4.1.9.1.1.34) | B-71 |
| tpCpuUtilization (1.3.6.1.4.1.5655.4.1.9.1.1.35) | B-72 |
| tpCpuUtilizationPeak (1.3.6.1.4.1.5655.4.1.9.1.1.36) | B-72 |
| tpCpuUtilizationPeakTime (1.3.6.1.4.1.5655.4.1.9.1.1.37) | B-72 |

tpFlowsCapacityUtilization (1.3.6.1.4.1.5655.4.1.9.1.1.38) **B-72**

tpFlowsCapacityUtilizationPeak (1.3.6.1.4.1.5655.4.1.9.1.1.39) **B-73**

tpFlowsCapacityUtilizationPeakTime (1.3.6.1.4.1.5655.4.1.9.1.1.40) **B-73**

tpServiceLoss (1.3.6.1.4.1.5655.4.1.9.1.1.41) **B-73**

pportTable (1.3.6.1.4.1.5655.4.1.10.1) **B-73**

pportEntry (1.3.6.1.4.1.5655.4.1.10.1.1) **B-74**

pportModuleIndex (1.3.6.1.4.1.5655.4.1.10.1.1.1) **B-74**

pportIndex (1.3.6.1.4.1.5655.4.1.10.1.1.2) **B-74**

pportType (1.3.6.1.4.1.5655.4.1.10.1.1.3) **B-75**

pportNumTxQueues (1.3.6.1.4.1.5655.4.1.10.1.1.4) **B-75**

pportIfIndex (1.3.6.1.4.1.5655.4.1.10.1.1.5) **B-75**

pportAdminSpeed (1.3.6.1.4.1.5655.4.1.10.1.1.6) **B-75**

pportAdminDuplex (1.3.6.1.4.1.5655.4.1.10.1.1.7) **B-76**

pportOperDuplex (1.3.6.1.4.1.5655.4.1.10.1.1.8) **B-76**

pportLinkIndex (1.3.6.1.4.1.5655.4.1.10.1.1.9) **B-76**

pportOperStatus (1.3.6.1.4.1.5655.4.1.10.1.1.10) **B-76**

txQueuesTable (1.3.6.1.4.1.5655.4.1.11.1) **B-77**

txQueuesEntry (1.3.6.1.4.1.5655.4.1.11.1.1) **B-77**

txQueuesModuleIndex (1.3.6.1.4.1.5655.4.1.11.1.1.1) **B-77**

txQueuesPortIndex (1.3.6.1.4.1.5655.4.1.11.1.1.2) **B-77**

txQueuesQueueIndex (1.3.6.1.4.1.5655.4.1.11.1.1.3) **B-78**

txQueuesDescription (1.3.6.1.4.1.5655.4.1.11.1.1.4) **B-78**

txQueuesBandwidth (1.3.6.1.4.1.5655.4.1.11.1.1.5) **B-78**

txQueuesUtilization (1.3.6.1.4.1.5655.4.1.11.1.1.6) **B-78**

txQueuesUtilizationPeak (1.3.6.1.4.1.5655.4.1.11.1.1.7) **B-79**

txQueuesUtilizationPeakTime (1.3.6.1.4.1.5655.4.1.11.1.1.8) **B-79**

txQueuesClearCountersTime (1.3.6.1.4.1.5655.4.1.11.1.1.9) **B-79**

txQueuesDroppedBytes (1.3.6.1.4.1.5655.4.1.11.1.1.10) **B-79**

globalControllersTable (1.3.6.1.4.1.5655.4.1.12.1) **B-80**

globalControllersEntry (1.3.6.1.4.1.5655.4.1.12.1.1) **B-80**

globalControllersModuleIndex (1.3.6.1.4.1.5655.4.1.12.1.1.1) **B-80**

globalControllersPortIndex (1.3.6.1.4.1.5655.4.1.12.1.1.2) **B-81**

globalControllersIndex (1.3.6.1.4.1.5655.4.1.12.1.1.3) **B-81**

globalControllersDescription (1.3.6.1.4.1.5655.4.1.12.1.1.4) **B-81**

globalControllersBandwidth (1.3.6.1.4.1.5655.4.1.12.1.1.5) **B-81**

globalControllersUtilization (1.3.6.1.4.1.5655.4.1.12.1.1.6) **B-82**

globalControllersUtilizationPeak (1.3.6.1.4.1.5655.4.1.12.1.1.7) **B-82**

globalControllersUtilizationPeakTime (1.3.6.1.4.1.5655.4.1.12.1.1.8) **B-82**

globalControllersClearCountersTime (1.3.6.1.4.1.5655.4.1.12.1.1.9) **B-82**

globalControllersDroppedBytes (1.3.6.1.4.1.5655.4.1.12.1.1.10) **B-83**

| | |
|---|-------------|
| appInfoTable (1.3.6.1.4.1.5655.4.1.13.1) | B-83 |
| appInfoEntry (1.3.6.1.4.1.5655.4.1.13.1.1) | B-83 |
| appName (1.3.6.1.4.1.5655.4.1.13.1.1.1) | B-83 |
| appDescription (1.3.6.1.4.1.5655.4.1.13.1.1.2) | B-84 |
| appVersion (1.3.6.1.4.1.5655.4.1.13.1.1.3) | B-84 |
| appPropertiesTable (1.3.6.1.4.1.5655.4.1.13.2) | B-84 |
| appPropertiesEntry (1.3.6.1.4.1.5655.4.1.13.2.1) | B-84 |
| apIndex (1.3.6.1.4.1.5655.4.1.13.2.1.1) | B-85 |
| apName (1.3.6.1.4.1.5655.4.1.13.2.1.2) | B-85 |
| apType (1.3.6.1.4.1.5655.4.1.13.2.1.3) | B-85 |
| appPropertiesValuesTable (1.3.6.1.4.1.5655.4.1.13.3) | B-86 |
| appPropertiesValueEntry (1.3.6.1.4.1.5655.4.1.13.3.1) | B-86 |
| apvIndex (1.3.6.1.4.1.5655.4.1.13.3.1.1) | B-86 |
| apvPropertyName (1.3.6.1.4.1.5655.4.1.13.3.1.2) | B-87 |
| apvRowStatus (1.3.6.1.4.1.5655.4.1.13.3.1.3) | B-87 |
| apvPropertyStringValue (1.3.6.1.4.1.5655.4.1.13.3.1.4) | B-87 |
| apvPropertyUintValue (1.3.6.1.4.1.5655.4.1.13.3.1.5) | B-87 |
| apvPropertyCounter64Value (1.3.6.1.4.1.5655.4.1.13.3.1.6) | B-88 |
| trafficCountersTable (1.3.6.1.4.1.5655.4.1.14.1) | B-88 |
| trafficCountersEntry (1.3.6.1.4.1.5655.4.1.14.1.1) | B-88 |
| trafficCounterIndex (1.3.6.1.4.1.5655.4.1.14.1.1.1) | B-88 |
| trafficCounterValue (1.3.6.1.4.1.5655.4.1.14.1.1.2) | B-89 |
| trafficCounterName (1.3.6.1.4.1.5655.4.1.14.1.1.3) | B-89 |
| trafficCounterType (1.3.6.1.4.1.5655.4.1.14.1.1.4) | B-89 |
| attackTypeTable (1.3.6.1.4.1.5655.4.1.15.1) | B-89 |
| attackTypeEntry (1.3.6.1.4.1.5655.4.1.15.1.1) | B-90 |
| attackTypeIndex (1.3.6.1.4.1.5655.4.1.15.1.1.1) | B-90 |
| attackTypeName (1.3.6.1.4.1.5655.4.1.15.1.1.2) | B-90 |
| attackTypeCurrentNumAttacks (1.3.6.1.4.1.5655.4.1.15.1.1.3) | B-90 |
| attackTypeTotalNumAttacks (1.3.6.1.4.1.5655.4.1.15.1.1.4) | B-91 |
| attackTypeTotalNumFlows (1.3.6.1.4.1.5655.4.1.15.1.1.5) | B-91 |
| attackTypeTotalNumSeconds (1.3.6.1.4.1.5655.4.1.15.1.1.6) | B-91 |
| vasServersTable (1.3.6.1.4.1.5655.4.1.16.1) | B-91 |
| vasServerEntry (1.3.6.1.4.1.5655.4.1.16.1.1) | B-92 |
| vasServerIndex (1.3.6.1.4.1.5655.4.1.16.1.1.1) | B-92 |
| vasServerId (1.3.6.1.4.1.5655.4.1.16.1.1.2) | B-92 |
| vasServerAdminStatus (1.3.6.1.4.1.5655.4.1.16.1.1.3) | B-92 |
| vasServerOperStatus (1.3.6.1.4.1.5655.4.1.16.1.1.4) | B-93 |
| mplsVpnSoftwareCountersTable (1.3.6.1.4.1.5655.4.1.17.1) | B-93 |
| mplsVpnSoftwareCountersEntry (1.3.6.1.4.1.5655.4.1.17.1.1) | B-93 |

[mplsVpnMaxHWMappings \(1.3.6.1.4.1.5655.4.1.17.1.1.1\)](#) **B-93**
[mplsVpnCurrentHWMappings \(1.3.6.1.4.1.5655.4.1.17.1.1.2\)](#) **B-94**
[Supported Standards](#) **B-95**



About this Guide

This preface describes who should read the *Cisco Service Control Engine (SCE) 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 SCE platform.

Document Revision History

| Cisco Service Control Release | Part Number | Publication Date |
|-------------------------------|-------------|------------------|
| Release 3.1.5 LA | OL-7827-07 | August, 2007 |

Description of Changes

The following chapters were updated to describe the updated VPN functionality:

- Configuring the Line Interface (How to Configure Tunneling Protocols)
- Managing Subscribers
- MPLS/VPN Support

The following chapter was updated to describe the updated TOS marking functionality:

- Configuring the Line Interface

Minor changes were made in the following chapters to clarify certain topics and issues:

- Configuring the Line Interface: Hardware packet drops
- Value Added Services (VAS) Traffic Forwarding: Requirements for VAS Servers
- Redundancy and Fail-Over: Emphasize L1 connection of cascade ports.
- Proprietary MIB Reference

| Cisco Service Control Release | Part Number | Publication Date |
|-------------------------------|-------------|------------------|
| Release 3.1.0 | OL-7827-06 | May, 2007 |

Description of Changes

Added the following new feature:

- Asymmetric Routing Topology

The following chapter was updated to include the NetflowV9 protocol option:

- Raw Data Formatting: The RDR Formatter and NetFlow Exporting

| Cisco Service Control Release | Part Number | Publication Date |
|-------------------------------|-------------|------------------|
| Release 3.0.5 | OL-7827-05 | November, 2006 |

Description of Changes

Added the following new feature:

- Managing the SCMP

The following sections were added or updated to explain various CLI commands that had not previously appeared in this guide:

- Monitoring the Operational Status of the SCE Platform
- Monitoring the Connection Mode
- Link Failure Reflection in Linecard-Aware Mode (SCE 2000 only)
- Removing Subscribers with Tunnel Mappings
- Traffic Rules

| Cisco Service Control Release | Part Number | Publication Date |
|-------------------------------|-------------|------------------|
| Release 3.0.3 | OL-7827-04 | May, 2006 |

Description of Changes

Added the following new features:

- MPLS/VPN Support (including MPLS/VPN-related changes in *Managing Subscribers and Configuring Tunneling Protocols*).
- Configuring VLAN Translation
- VAS over 10G

The Proprietary MIB Reference was reorganized to reflect reorganization of the pcube Enterprise MIB.

| Cisco Service Control Release | Part Number | Publication Date |
|--------------------------------------|--------------------|-------------------------|
| Release 3.0 | OL-7827-03 | December, 2005 |

Description of Changes

Added the following new features:

- Value Added Services (VAS)
- Traffic Forwarding
- Monitoring SCE Platform Utilization
- Configuring the Management Ports for Redundancy
- Management Interface Security
- <SKIP>+ Authentication, Authorization and Accounting
- Dynamic Mapping of RDRs to Categories

| Cisco Service Control Release | Part Number | Publication Date |
|--------------------------------------|--------------------|-------------------------|
| Release 2.5.7 | OL-7827-02 | May, 2005 |

Description of Changes

Complete reorganization and revision of product documentation.

Organization

The major sections of this guide are as follows:

Table 1 **Documentation Organization**

| Chapter | Title | Description |
|------------|--|---|
| Chapter 1 | General Overview | Overview of SCE platform management.. |
| Chapter 2 | Command Line Interface | Detailed explanation of how to use the Cisco SCE Command-line Interface. |
| Chapter 3 | 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 | Configuring the Line Interface | Explanation of how to configure tunneling, TOS marking, and traffic rules.. |
| Chapter 7 | Configuring the Connection | Explanation of how to configure the connection mode, link mode, and failure behaviors. |
| Chapter 8 | Raw Data Formatting: The RDR Formatter and NetFlow Exporting | Explanation of how to configure the RDR Formatter so that RDRs are sent to the proper destinations. |
| Chapter 9 | Managing Subscribers | Explanation of how to import and export subscriber information and how to monitor subscribers. |
| Chapter 10 | Redundancy and Fail-Over | Explanation of how to configure and manage a redundant system. This chapter applies only to the SCE 2000 platform. |
| Chapter 11 | Identifying and Preventing Distributed-Denial-Of-Service Attacks | Explanation of how to configure attack filtering. |

Table 1 **Documentation Organization (continued)**

| Chapter | Title | Description |
|----------------|---|---|
| Chapter 12 | Value Added Services (VAS) Traffic Forwarding | Explanation of Value Added Services (VAS) and how to configure VAS traffic forwarding. |
| Chapter 13 | MPLS/VPN Support | Explanation of MPLS/VPN support, and how to configure and monitor MPLS/VPN subscribers and support. |
| Chapter 14 | Managing the SCMP | Explanation of Service Control Management Protocol (SCMP), which is a protocol that integrates the 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. |
| Appendix A | Monitoring SCE Platform Utilization | Explanation of how to monitor SCE platforms that are installed in real traffic. |
| Appendix B | Proprietary MIB Reference | Definition of the proprietary Service Control Enterprise MIB. |

Related Publications

Your 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 Service Control Engine (SCE) CLI Command Reference*
- For complete installation information, including initial configuration, refer to the relevant installation guide:
 - *Cisco SCE 2000 4xGBE Installation and Configuration Guide*
 - *Cisco SCE 2000 4/8xFE Installation and Configuration Guide*
 - *Cisco SCE 1000 2xGBE Installation and Configuration Guide*



Note

You can access Cisco software configuration and hardware installation and maintenance documentation on the World Wide Web at Cisco Website URL. Translated documentation is available at the following URL: International Cisco Website

- For initial installation and startup information, refer to the relevant quick start guide:
 - *Cisco SCE 2000 4xGBE Quick Start Guide*
 - *Cisco SCE 2000 4/8xFE Quick Start Guide*
 - *Cisco SCE 1000 2xGBE Quick Start Guide*
- For international agency compliance, safety, and statutory information for wide-area network (WAN) interfaces for the SCE 2000 platform, refer to the regulatory and safety information document:
 - *Regulatory Compliance and Safety Information for Cisco Service Control Engine (SCE)* .
- 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 Service Control Application Reporter 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 xxxviii
 - The Cisco Information Packet that shipped with your SCE 2000 platform.

Conventions

This document uses the following conventions:

Table 2 **Documentation Conventions**

| Convention | Description |
|-----------------------------|--|
| boldface font | Commands and keywords are in boldface . |
| <i>italic</i> font | Arguments for which you supply values are in <i>italics</i> . |
| [] | Elements in square brackets are optional. |
| {x y z} | 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. |
| screen font | Terminal sessions and information that the system displays are in <code>screen</code> font. |
| boldface screen font | Information you must enter is in boldface screen font. |
| <i>italic screen</i> font | Arguments for which you supply values are in <i>italic screenfont</i> . |
| <> | 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* . Notes contain helpful suggestions or references to materials not covered in this manual.

Means *reader be careful* . In this situation, you might do something that could result in equipment damage or loss of data.

Means *reader be warned* . In this situation, you might do something 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 a Really Simple Syndication (RSS) feed and set content to be delivered directly to your desktop using a reader application. The RSS feeds are a free service and Cisco currently supports RSS Version 2.0.



CHAPTER 1

General Overview

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

It also briefly describes the hardware capabilities of the Service Control Engine (SCE) platform and the Cisco specific applications that together compose the total Cisco Service Control solution.

- [Information About the Cisco Service Control Concept, page 1-1](#)
- [Cisco Service Control Capabilities, page 1-2](#)
- [The SCE Platform, page 1-3](#)
- [Information About Management and Collection, page 1-4](#)

Information About the Cisco Service Control Concept

- [The Cisco Service Control Solution, page 1-1](#)
- [Service Control for Broadband Service Providers, page 1-2](#)

The Cisco Service Control Solution

The Cisco Service Control solution is delivered through a combination of purpose-built hardware and specific software solutions that address various service control challenges faced by service providers. The SCE platform is designed to support classification, analysis, and control of Internet/IP traffic.

Service Control enables service providers to create profitable new revenue streams while capitalizing on their existing infrastructure. With the power of Service Control, service providers have the ability to analyze, charge for, and control IP network traffic at multigigabit wire line speeds. The Cisco Service Control solution also gives service providers the tools they need to identify and target high-margin content-based services and to enable their delivery.

As the downturn in the telecommunications industry has shown, IP service providers' business models need to be reworked to make them profitable. Having spent billions of dollars to build ever larger data links, providers have incurred massive debts and faced rising costs. At the same time, 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. However, capturing real profits from IP services requires more than simply running those services over data links; it requires detailed monitoring and precise, real-time control and awareness of services as they are delivered. Cisco provides Service Control solutions that allow the service provider to bridge this gap.

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 new 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 SLAs
- Implement different service levels for different types of customers, content, or applications
- Identify network abusers who are violating the Acceptable Use Policy
- Identify and manage peer-to-peer, NNTP (news) traffic, and spam abusers
- Enforce the Acceptable Use Policy (AUP)
- Integrate Service Control solutions easily with existing network elements and BSS/OSS systems

Cisco Service Control Capabilities

The core of the Cisco Service Control solution is the purpose-built network hardware device: the Service Control Engine (SCE). The core capabilities of the 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 in order to maintain the state of each subscriber transmitting traffic through the 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 via 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 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 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 easily adapt to new services and applications in the ever-changing 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 SCE provides a set of open and well-documented APIs that allows a quick and robust integration process.
- Scalable high-performance service engines—The ability to perform all these operations at wire speed.

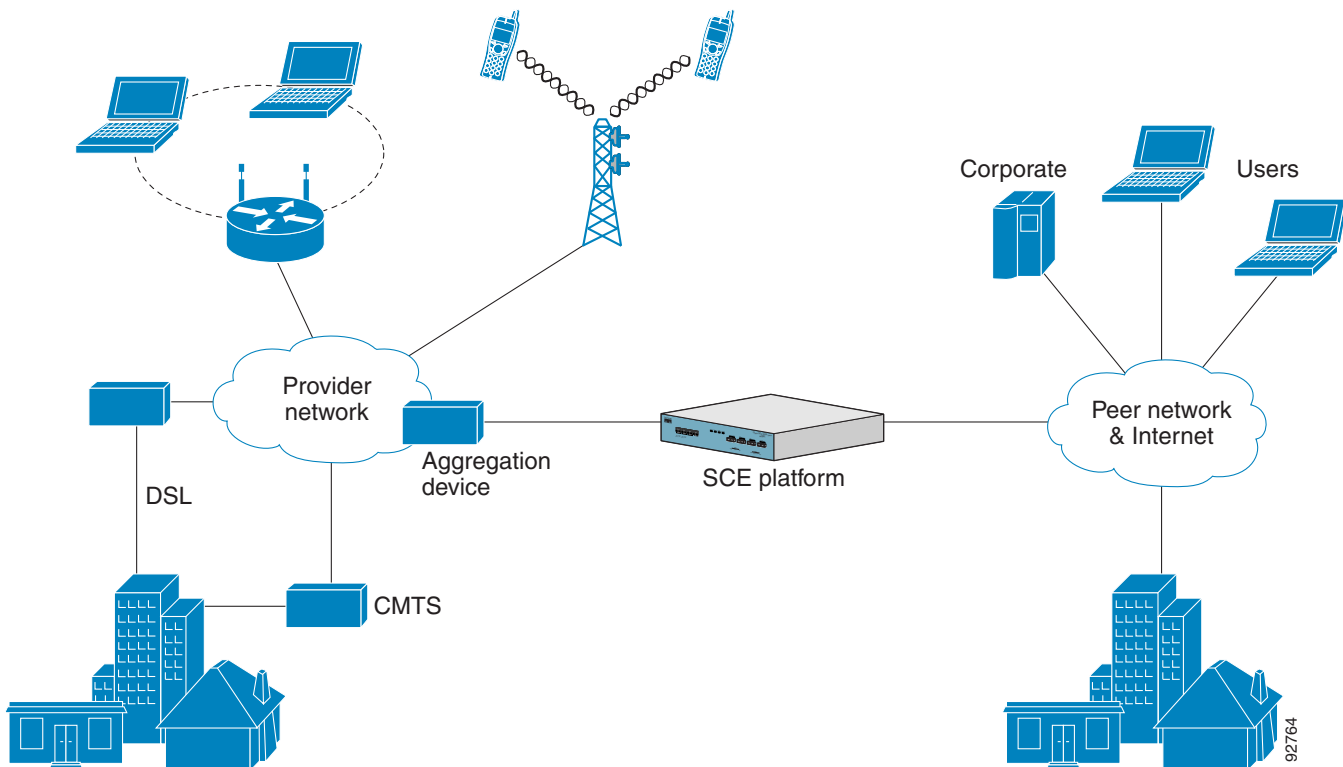
The SCE Platform

The SCE family of programmable network devices is capable of performing application-layer stateful-flow inspection of IP traffic, and controlling that traffic based on configurable rules. The SCE platform is a purpose-built network device that uses ASIC components and RISC processors to go beyond packet counting and delve deeper into the contents of network traffic. Providing programmable, stateful inspection of bidirectional traffic flows and mapping these flows with user ownership, the SCE platforms provide real-time classification of network usage. This information provides the basis of the SCE platform advanced traffic-control and bandwidth-shaping functionality. Where most bandwidth shaper functionality ends, the SCE platform provides more control and shaping options, including:

- Layer 7 stateful wire-speed packet inspection and classification
- Robust support for over 600 protocols and applications, including:
 - General—HTTP, HTTPS, FTP, TELNET, NNTP, SMTP, POP3, IMAP, WAP, and others
 - P2P file sharing—FastTrack-KazaA, Gnutella, BitTorrent, Winny, Hotline, eDonkey, DirectConnect, Piolet, and others
 - P2P VoIP—Skype, Skinny, DingoTel, and others
 - Streaming and Multimedia—RTSP, SIP, HTTP streaming, RTP/RTCP, and others
- Programmable system core for flexible reporting and bandwidth control
- Transparent network and BSS/OSS integration into existing networks
- Subscriber awareness that relates traffic and usage to specific customers

The following diagram illustrates a common deployment of an SCE platform in a network.

Figure 1-1 SCE Platform in the Network




Information About 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 Control management

These management interfaces are designed to comply with common management standards and to integrate easily with existing OSS infrastructure.

Figure 1-2 Service Control Management Infrastructure

Network Management

Cisco provides complete network FCAPS (Fault, Configuration, Accounting, Performance, Security) Management.

Two interfaces are provided for 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 (via SNMP traps) and performance monitoring functionality.

Subscriber Management

Where the Cisco Service Control Application for Broadband (SCA BB) enforces different 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 the OSS and the 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 SCE platform automatically upon logon of a subscriber.
- Pull Mode—The SM sends subscriber information to the SCE platform in response to a query from the 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 an SCE platform. The SCA BB application provides tools to automate the distribution of these configuration files to SCE platforms. This simple, standards-based approach makes it easy to manage multiple devices in a large network.

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

Data Collection

The Cisco Service Control solution generates usage data and statistics from the SCE platform and forwards them as Raw Data Records (RDRs), using a simple TCP-based protocol (RDR-Protocol). The Cisco Service Control Management Suite (SCMS) Collection Manager (CM) software implements the collection system, listening in on RDRs from one or more SCE platforms and processing them on the local machine. The data is then stored for analysis and reporting functions, and for the collection and presentation of data to additional OSS systems such as billing.



CHAPTER 2

Command Line Interface

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

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

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

- [Information About Getting Help, page 2-2](#)
- [Information About Authorization and Command Levels \(Hierarchy\), page 2-3](#)
- [Information About CLI Help Features, page 2-7](#)
- [Information About Navigational and Shortcut Features, page 2-8](#)
- [Information About Managing Command Output, page 2-11](#)
- [CLI Authorization Levels, page 2-13](#)
- [Exiting Modes, page 2-15](#)
- [Information About Navigating Between Configuration Modes, page 2-16](#)
- [How to create a CLI script, page 2-21](#)

Information About Getting Help

To obtain a list of commands that are available for each command mode, enter a question mark (?) at the system prompt. You also can obtain a list of keywords and arguments associated with any command using the context-sensitive help feature.

The following table lists commands you can enter to get help that is specific to a command mode, a command, a keyword, or an argument.

Table 2-1 **Getting Help**

| Command | Purpose |
|--------------------------------|---|
| abbreviated-command-entry? | 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> | Complete a partial command name. |
| ? | List all commands available for a particular command mode |
| command? | List the keywords associated with the specified command. |
| command keyword? | List the arguments associated with the specified keyword. Leave a space between the keyword and question mark |

Information About Authorization and Command 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 at all authorization levels. See [CLI Authorization Levels, page 2-13](#)

- **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. [CLI Command Hierarchy, 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 host name. See [CLI Authorization Levels, page 2-13](#)

**Note**

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

CLI Command 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 SCE platform, file system operations, and other operations that cannot alter the configuration of the 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 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 SCE platform restores the saved configuration.

The following table 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 SCE platform. | Admin/ Root | SCE # |
| Global Configuration | Configuration of general system parameters, such as DNS, host name, and time zone. | Admin/ Root | 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)# |
| Interface Configuration | Configuration of specific system interface parameters, such as the Line Card, and the Ethernet interfaces. | Admin/ Root | SCE (config if)# |
| Line Configuration | Configuration of Telnet lines, such as an access-list. | Admin/ Root | SCE (config-line)# |

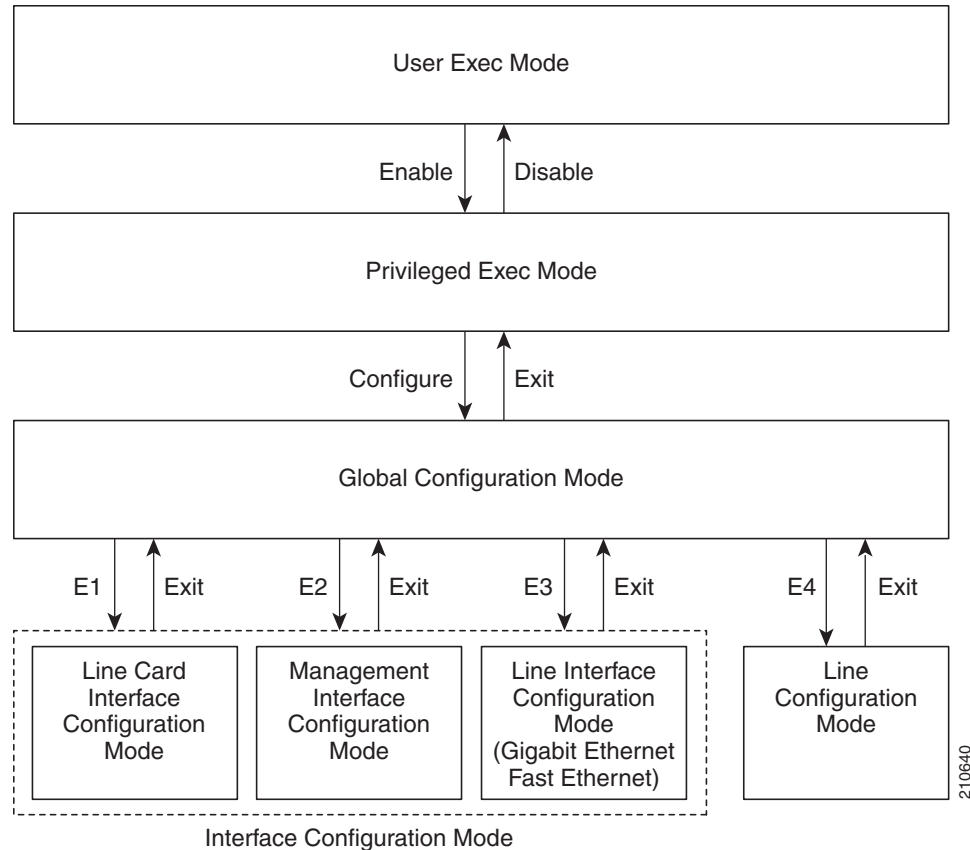
When you login to the system, you have the User authorization level and enter User Exec mode. Changing the authorization level to Viewer does not change the mode. Changing the authorization level to Admin automatically moves you to Privileged Exec mode. To move to any of the configuration modes, you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

A telnet session begins with a request for password, and will not continue until the proper user password is supplied. This enhances the security of the system by not revealing its identity to unauthorized people.

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

The figure below, 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 different configure interface modes and the Line Configuration Mode:

- **E1 interface Linecard 0**
- E2 interface Mng 0/1 or 0/2**(management port, all platforms)
- E3 interface GigabitEthernet 0/1 or 0/2**(line ports, SCE 1000 platform)
- E3 interface GigabitEthernet 0/1, 0/2, 0/3, or 0/4**(line ports, SCE 2000 4xGBE platform)
- E3 interface FastEthernet 0/1, 0/2, 0/3, or 0/4**(line ports, SCE 2000 4/8xFE platform)
- E4 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 the above figure).

Example:

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

- Enter global configuration mode
- Configure the SCE platform time zone
- Enter Mng Interface configuration mode for Mng port 1
- Configure the speed of the management interface
- Exit the Mng Interface configuration mode to the global configuration mode
- Enter the Linecard Interface configuration
- Define the link mode
- Exit Linecard Interface configuration mode to the global configuration mode
- Exit global configuration 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 all-links forwarding
SCE(config if)#exit
SCE(config)#exit
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 follows:

| This prompt... | Indicates this... |
|----------------|----------------------------------|
| > | indicates User and Viewer levels |
| number | indicates Admin level |
| #> | indicates Root level |

Command hierarchy levels are indicated as follows:

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

Example:

The prompt `SCE1(config if)#` indicates:

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

Information About CLI Help Features

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

- Partial help
- Argument help

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 command's associated keywords or parameters, 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 ?  
CommunityDefine community string  
ContactSet system contact  
Enable Enable the SNMP agent  
Host Set traps destination  
LocationSet system location  
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-configCopy running configuration file
      STRING      Source file name
SCE#copy
```

The [no] Prefix

Many CLI commands offer the option of adding the word **no** before the command to disable the feature controlled by the command or revert it to its default configuration. This notation is shown in the CLI Command Reference (on page 1) as **[no]** to denote it is optional.

For example, **no service telnetd** disables the telnet server. Enabling the telnet server is done by typing **service telnetd** .

Information About Navigational and Shortcut Features

- [Command History, page 2-9](#)
- [Keyboard Shortcuts, page 2-9](#)
- [Tab Completion, page 2-10](#)
- [FTP User Name and Password, page 2-11](#)

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 the following table.

Table 2-3 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 SCE platform has several keyboard shortcuts that make it easier to navigate and use the system. The following table shows the keyboard shortcuts available.

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

Table 2-4 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 |
| 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 |

Table 2-4 Keyboard Shortcuts (continued)

| Description | Shortcut key |
|--|--------------|
| 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 |
| Deletes all characters from the cursor to the beginning of the line. (Same functionality as CTRL-U.) | CTRL-X |
| 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 |

Tab Completion

The CLI interface features tab completion. When you type in the first letters of a command and type <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 type <Enter> instead of <Tab>, and there is no ambiguity, the system actually carries out the command which would be filled in by the rest of the word.

Example:

The following example displays how the system completes a partial (unique) command for the **enable** command. Because enable does not require any parameters, the system simply carries out the **enable** command when the user presses Enter.

```
SCE>en<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 vk
SCE# ip FTP username vk
SCE# copy ftp://@10.1.1.253/h:/config.tmp myconf.txt
connecting 10.1.1.253 (user name vk password vk) to retrieve config.tmp
SCE#
```

Information About Managing Command Output

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 (on page) 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 <file-name>`
- `<command>| append <file-name>`

Following is an example of how to do the following:

- Filter the **more** command to display from a `csv subscriber` file only the gold package subscribers.
- 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_10.10.2004 include gold append current_gold_subscribers
```

CLI Authorization Levels

The SCE platform has four authorization levels, which represent the user access permissions. When you initially connect to the 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 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 the following table.

Table 2-5 **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 2-5 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 SCE platform. | 10 | number |
| 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 | #> |

How to change from User to Viewer level authorization

A telnet session begins with a request for password, and will not continue until the proper user password is supplied. This enhances the security of the system by not revealing its identity to unauthorized people.

Step 1 From the `sce>` prompt, type **enable 5** and press Enter.

The system prompts for a password by showing the prompt `Password:`

Step 2 Type in the password for the Viewer level and press Enter.

Note that the password is an access-level authorization setting, not an individual user password.

The system prompt `sce>` does not change when you move from User to Viewer level.

How to log in with Root level authorization

- Step 1** Initiate a telnet connection.
- Step 2** A `password:` prompt appears. Type in the user level password and press Enter.
The `SCE>` prompt appears.
You now have user level authorization.
- Step 3** From the `SCE>` prompt, type **enable 15** and press Enter.
The system prompts for a password by showing the prompt `Password:`
- Step 4** Type in the password for the Root level and press Enter.
Note that the password is an access-level authorization setting, not an individual user password.
The system prompt changes to `SCE#>` to show you are now in Root level.
-

This example illustrates how to change the authorization level from User to Root, and then revert back to Admin. No password is required for moving to a lower authorization level.

```
SCE>enable 15
Password: <Cisco>
SCE#>disable
SCE#
```

Exiting Modes

This section describes how to revert to a previous mode.

- To exit from one authorization level to the previous one, use the **disable** command.
- To exit from one mode to another with the Admin authorization level (these are the various configuration modes), use the **exit** command.

How to exit from the Privileged Exec mode and revert to the Viewer mode

- Step 1** At the `SCE#` prompt, type **disable**, and press **Enter**.
The `SCE>` prompt for the Viewer and User Exec mode appears.
-

How to exit from the Global Configuration Mode

-
- Step 1** At the SCE (config)# prompt, type **exit** , and press **Enter**.
The appropriate prompt for the previous level appears.
-

Example:

This example illustrates how to change the authorization level from User to Root, and then revert back to Admin. No password is required for moving to a lower authorization level.

```
SCE(config if)#exitSCE(config)#
```

Information About Navigating Between Configuration Modes

- [Entering and Exiting Global Configuration Mode, page 2-16](#)
- [Information About Interface Configuration Modes, page 2-17](#)

Entering and Exiting Global Configuration Mode

- [How to enter the Global Configuration Mode, page 2-16](#)
- [How to exit the Global Configuration Mode, page 2-16](#)

How to enter the Global Configuration Mode

-
- Step 1** At the SCE# prompt, type **configure** , and press Enter.
The SCE (config)# prompt appears.
-

How to exit the Global Configuration Mode

-
- Step 1** At the SCE# prompt, type **exit** , and press Enter.
The SCE# prompt appears.
-

Information About Interface Configuration Modes

The components that are configured by the Interface Configuration Modes are:

- Card
 - Linecard — **Interface Linecard 0**
- Ports
 - [Configuring the Physical Ports, page 2-17](#)
- Telnet
 - • Line Configuration Mode — **Line vty 0**

The Line Configuration Mode enables you to configure Telnet parameters.

Configuring the Physical Ports

The SCE platform contains the following physical port interfaces:

- Management:

Interface Mng 0/1 or 0/2

The Management Interface mode configures the settings for the interface to a remote management console. The two management ports support management interface redundancy.

The following commands are used to configure the management port:

- ip address
- duplex
- speed
- active-port (SCE 2000 platform only)
- auto-fail-over

- Fast Ethernet (SCE 2000 4/8xFE):

Interface FastEthernet 0/1, 0/2, 0/3, or 0/4

The FastEthernet Interface mode configures the settings for the FastEthernet interface to the Internet traffic on the wire. Each of the four ports can be set individually.

The following commands are used to configure the Fast Ethernet line ports:

- bandwidth
- duplex
- queue
- speed

- Gigabit Ethernet (SCE 1000 platform):

Interface GigabitEthernet 0/1, or 0/2

The GigabitEthernet Interface mode configures the settings for the GigabitEthernet interface to the Internet traffic on the wire. Each of the two ports can be set individually.

- Gigabit Ethernet (SCE 2000 4xGBE platform):

- The GigabitEthernet Interface mode configures the settings for the GigabitEthernet interface to the Internet traffic on the wire. Each of the four ports can be set individually.

The following commands are used to configure the Gigabit Ethernet line ports:

- auto-negotiate (GigabitEthernet only)
- bandwidth
- queue

**Note**

You must specify the slot number/interface number when referencing any interface. The slot number is always 0, and the interfaces are numbered as follows:

**Note**

Management Interface: **1,2**

**Note**

Ethernet Line Interfaces:

SCE 1000 platform: **1,2**

SCE 2000 platform: **1,2,3,4**

Entering Management Interface Configuration Mode

Before you can configure the parameters for the management interface, you must be in the Mng Interface Configuration Mode.

How to enter Mng Interface Configuration Mode

Step 1 To enter Global Configuration Mode, type **configure** and press Enter.

The `SCE(config)#` prompt appears.

Step 2 Type **interface Mng [0/1/0/2]** and press Enter.

The `SCE(config-if)#` prompt appears.

The system prompt changes to reflect the higher level mode.

How to return to the Global Configuration mode

Step 1 Type **Exit** .

How to enter Linecard Interface Configuration mode

The following procedure is for entering Line Card Interface Configuration mode. The procedures for entering the other interfaces are the same except for the interface command as described above and in [How to navigate from one Interface Configuration Mode to another, page 2-20](#).

-
- Step 1** To enter Global Configuration Mode, at the SCE# prompt, type **configure** , and press Enter.
The SCE(config)# prompt appears.
- Step 2** Type **interface Linecard 0**, and press Enter.
The SCE(config if)# prompt appears.
- Step 3** To return to Global Configuration Mode, type **exit** and press Enter.
The SCE(config)# prompt appears.
- Step 4** To exit Global Configuration Mode, type **exit** and press Enter.
-

Entering Ethernet Line Interface Configuration Mode

How to enter the FastEthernet Interface Configuration Mode

-
- Step 1** To enter Global Configuration Mode, at the SCE# prompt, type **configure** , and press Enter.
The SCE(config)# prompt appears.
- Step 2** For the SCE 2000, type interface **FastEthernet [0/1/0/2/0/3/0/4]**and press Enter.
The SCE(config if)# prompt appears.
-

The following example shows how to enter Configuration Mode for the FastEthernet Interface number 3.

```
SCE(config)#interface FastEthernet 0/3SCE(config if)#
```

How to enter the GigabitEthernet Interface Configuration Mode:

-
- Step 1** To enter Global Configuration Mode, at the SCE# prompt, type **configure** , and press Enter.
The SCE(config)# prompt appears.
- Step 2** For the SCE 1000, type interface **GigabitEthernet [0/1/0/2]** and press Enter.
- Step 3** For the SCE 2000, type **GigabitEthernet [0/1/0/2/0/3/0/4]**and press Enter.
The SCE(config if)# prompt appears.
-

The following example shows how to enter Configuration Mode for the GigabitEthernet Interface number 2.

```
SCE(config)#interface GigabitEthernet 0/2  
SCE(config if)#
```

How to navigate from one Interface Configuration Mode to another

-
- Step 1** Type **Exit** .
You are returned to the Global Configuration Mode.
- Step 2** Type the appropriate command to enter a different Interface Configuration Mode.
-

The "do" Command: Executing Commands Without Exiting

There are four configuration command modes:

- Global configuration mode
- Management interface configuration mode
- Interface configuration mode
- Line configuration mode

When you are in one of these 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)#`) prompt, type **do** *<command>*.
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
```

How to create 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 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 SCE platforms. The available script commands are:

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

Step 1 At the `SCE#` prompt, type **script capture** *sample1.scr* where *sample1.scr* is the name of the script.


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 '/tffs0/images/new.pkg'...
Verifying package file...
Package file verified OK.
Device '/tffs0/' 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#
```

CHAPTER 3

Operations

This module describes basic operations necessary for managing the SCE platform.

- [How to Manage Configurations, page 3-1](#)
- [How to Upgrade the SCE Platform Firmware, page 3-8](#)
- [How to Manage Application Files, page 3-10](#)
- [How to Monitor the Operational Status of the SCE Platform, page 3-13](#)
- [How to Display the SCE Platform Version Information, page 3-15](#)
- [How to Display the SCE Platform Inventory, page 3-16](#)
- [How to Display the System Uptime, page 3-17](#)
- [How to Reboot and Shut Down the SCE Platform, page 3-18](#)

How to Manage Configurations

This section explains how to view, save, and recover configuration files, as well as how to create a backup configuration file.

- [Managing Configurations, page 3-2](#)
- [How to View Configurations, page 3-2](#)
- [How to Remove a Configuration, page 3-3](#)
- [How to Save the Configuration Settings, page 3-4](#)
- [How to Restore a Previous Configuration, page 3-5](#)
- [How to Backup Configuration Files, page 3-7](#)

Managing Configurations

The SCE platform uses two configuration files:

- Startup configuration — This file contains the non-default configuration as saved by the user. The **startup-config** file is loaded each time the SCE platform reboots.
- Running configuration — This file contains results of configuration commands entered by the user. The **running-config** file is saved in the SCE platform volatile memory and is effective only as long as the SCE platform is up and running.

Use the following commands to view and save the configuration files.

You can also recover a previous configuration from a saved configuration file, as well as completely remove all current user configuration.

How to View Configurations

```
show running-config
```

Viewing Configurations

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

The SCE platform provides commands for:

- Viewing the running configuration
- Viewing the startup configuration

After configuring the SCE platform, you may query for the running configuration using the command **show running-config**. This command displays the non-default running configuration. To view all SCE platform running configuration, whether it is the default or not, you may use the option **all-data** in the **show running-config** command.

Step 1 At the SCE# prompt, type `show running-config` and press Enter.

The specified configuration file is displayed.

Viewing Configurations: Example

This example shows how to view the running configuration.

```
SCE#show running-config
#This is a general configuration file (running-config).
#Created on 15:50:56 CET MON December 11 2006
#cli-type 1
#version 1
clock timezone CET 1
snmp-server community "public" ro
snmp-server host 10.1.1.253 traps version 1 "public"
interface LineCard 0
connection-mode active
no silent
no shutdown
flow-aging default-timeout UDP 60
interface FastEthernet 0/0
ip address 10.1.5.109 255.255.0.0
interface FastEthernet 0/1
interface FastEthernet 0/2
exit
line vty 0 4
no timeout
exit
SCE#
```

How to Remove a Configuration

Removing the Configuration

You can completely remove all current configuration by removing all configuration files. The following data is deleted by this command::

- General configuration files
- Application configuration files
- Static party DB files
- Management agent installed MBeans

**Note**

After using this command, the SCE platform should be reloaded immediately to ensure that it returns to the 'factory default' state.

Step 1

At the SCE(config)# prompt, type `erase startup-config-all` and press Enter.

All configuration files are removed, including configuration files not explicitly managed by the user, as listed above.

How to Save the Configuration Settings

- [Saving the Configuration Settings, page 3-4](#)
- [Saving the Configuration Settings: Example, page 3-5](#)

Saving the Configuration Settings

When you make changes to the current running configuration and you want those changes to continue to be valid when the system restarts, you must save the changes before leaving the management session, that is, you must save the running configuration to the startup configuration file.

For backup purposes, the old startup-config file is saved under the directory: **tffs0:system/prevconf**. Refer to [How to Restore a Previous Configuration, page 3-5](#) for an explanation of how to recover a previous configuration.

The SCE platform provides multiple interfaces for the purpose of configuration and management. All interfaces supply an API to the same database of the SCE platform, so that any configuration change 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.

-
- Step 1** At the SCE# prompt, type `show running-config` and press Enter.
Displays the running configuration.
- Step 2** Check the displayed configuration to make sure that all parameters are set to the desired values. If not, make the changes you want before saving.
Refer to the relevant sections of this guide for more information regarding specific configuration parameters.
- Step 3** At the SCE# prompt, type `copy running-config startup-config` and press Enter.
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.txt** located in the directory: **tffs0:system**.
-

Saving the Configuration Settings: Example

The following example shows how to review and then save the running configuration file.

```
SCE#show running-config
#This is a general configuration file (running-config).
#Created on 15:50:56 CET MON February 11 2006
#cli-type 1
#version 1
clock timezone CET 1
snmp-server community "public" ro
snmp-server host 10.1.1.253 traps version 1 "public"
interface LineCard 0
connection-mode active
no silent
no shutdown
flow-aging default-timeout UDP 60
interface FastEthernet 0/0
ip address 10.1.5.109 255.255.0.0
interface FastEthernet 0/1
interface FastEthernet 0/2
exit
line vty 0 4
no timeout
exit
SCE#
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#
```

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
```

How to Restore a Previous Configuration

- [Restoring a Previous Configuration, page 3-6](#)
- [Restoring a Previous Configuration: Example, page 3-6](#)

Restoring a Previous Configuration

When you save a new configuration, the system automatically backs up the old configuration in the directory `tffs0: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.

Use the CLI command `more` to view the old startup configuration files.

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

Since the restore operation overwrites the current configuration file, you cannot undo the configuration restore operation. It is recommended to always backup the current configuration file first.

-
- Step 1** At the SCE# prompt, type `more tffs0:system/prevconf/ filename` and press Enter.
- Displays the contents of the specified backup configuration file. Backup configuration filenames are `config.tx1-config.tx9`.
- Step 2** Read the configuration information to make sure it is the configuration you want to restore.
- Be sure that you are restoring the proper configuration file, since the restore operation is not reversible. It is recommended to always backup the current configuration file first.
- Step 3** At the SCE# prompt, type `copy tffs0:system/prevconf/ filename tffs0:system/ filename` and press Enter.
- Overwrites the current startup configuration file with the contents of the specified backup configuration file.
- Backup configuration filenames are `config.tx1-config.tx9`.
-

Restoring a Previous Configuration: Example

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

```
SCE#more tffs0:system/prevconf/config.tx1
#This is a general configuration file (running-config).
#Created on 19:36:07 UTC THU February 14 2006
#cli-type 1
#version 1
interface LineCard 0
no silent
no shutdown
interface FastEthernet 0/0
ip address 10.1.5.109 255.255.0.0
interface FastEthernet 0/1
interface FastEthernet 0/2
exit
line vty 0 4
exit
SCE#copy tffs0:system/prevconf/config.tx1 tffs0:system/config.txt
SCE#
```

How to Backup Configuration Files

- [Backing Up Configuration Files, page 3-7](#)
- [Options, page 3-7](#)
- [How to Create a Backup Configuration File, page 3-7](#)
- [How to Upload a Backup Configuration File, page 3-8](#)

Backing Up Configuration Files

Although a backup of the configuration file is created automatically under certain circumstances, it is useful to be able to explicitly create a backup configuration file.

For example, it can be used in a cascaded solution to copy the configuration from one SCE platform to the other, as follows:

1. To create a backup configuration file, execute this command on the first SCE platform, specifying an FTP backup file:
copy startup-config *backup-file*
2. To upload the backup configuration file to the cascaded SCE platform, execute this command on that SCE platform, specifying the previously created backup file:
copy *backup-file* **startup-config**

Options

The following option is available:

- *backup-file*— The name of the backup configuration file to be created. The file name should be in 8.3 format, that is, there are a maximum of 8 characters before the period and three characters following it.

The backup file may be created via FTP or it may be a local file, as shown in the following examples:

- via FTP: **ftp://user:pass@host/drive:/dir/bckupcfg.txt**
- local: **/tffs0/bckupcfg.txt**

How to Create a Backup Configuration File

-
- Step 1** At the SCE# prompt, type **copy startup-config** *backup-file* and press Enter.
Backs up the startup-config file to the specified file.
-

How to Upload a Backup Configuration File

Step 1 At the SCE# prompt, type `copy backup-filestartup-config` and press Enter.

Copying a Backup Configuration File to a Different SCE Platform: Example

This example shows how to copy the configuration from one SCE platform to another.

On the first SCE platform, enter the following command:

```
SCE1#copy startup-config ftp://adminuser:mypassword@10.10.10.10/c:/config/bckupcfg.txt  
SCE1number
```

On the second SCE platform, enter the following command:

```
SCE2#copy ftp://adminuser:mypassword@10.10.10.10/c:/config/bckupcfg.txt startup-config  
SCE2number
```

How to Upgrade the SCE Platform Firmware

- [Upgrading SCE Platform Firmware, page 3-9](#)
- [Upgrading SCE Platform Firmware: Example, page 3-10](#)

Upgrading SCE Platform Firmware

Cisco distributes upgrades to the software and firmware on the SCE platform. Cisco distributes upgrade software as a file with the extension `.pkg` that is installed directly from the ftp site without being copied to the disk. This procedure walks you through installation and rebooting of the SCE platform with the new firmware.

-
- Step 1** At the command prompt, type **configure** and press Enter.
Enters Global Configuration mode.
- Step 2** **boot** Type **boot system ftp://user:password@host/drive:dir/seNum.pkg**, where *seNum.pkg* is the file name on the ftp site.
The **boot** command verifies that the package is a legal, appropriate update for the SCE platform and that the file was not corrupted. It does not perform an upgrade, but does keep in the system memory that a `pkg` file is available.
- Step 3** Type **exit** .
Exits Global Configuration mode.
- Step 4** Type **copy running-config startup-config** .
This command re-verifies that the package is valid, and extracts the upgrade to the Flash file system.
The system notifies you that it is performing the extraction as follows:
Backing-up configuration file..
Writing configuration file..
Extracting new system image..
Extracted OK.
- Step 5** Type **reload** .
Reboots the system.
The system first prompts you for confirmation:
Are you sure?
- Step 6** Type **Y** and press Enter.
The system sends the following message and reboots.
The system is about to reboot, this will end your CLI session
-

Upgrading SCE Platform Firmware: Example

The following example shows the full procedure for performing a firmware update.

```
SCE#configure
SCE(config)# boot system ftp://vk:vk@10.1.1.230/downloads/SE2000.pkg
SCE(config)#exit
SCE#copy running-config startup-config
Backing-up configuration file..
Writing configuration file..
Extracting new system image..
Extracted OK.
SCE#>reload
Are you sure? y
the system is about to reboot, this will end your CLI session
```

How to Manage Application Files

This module explains how to install and upgrade application files.

- [Configuring Applications, page 3-10](#)
- [Managing Application Files, page 3-11](#)
- [How to Display Information about an Application File, page 3-11](#)
- [How to Install an Application, page 3-11](#)
- [How to Uninstall an Application, page 3-12](#)
- [How to Upgrade an Application, page 3-12](#)
- [How to Undo an Upgrade of an Application, page 3-12](#)
- [How to Display the Last pqi File that was Installed, page 3-12](#)

Configuring Applications

The SCE platform can be configured to run with different Service Control applications by installing the appropriate file. All SCE platform application files are **pqi** files, that is, the filename must end with the pqi extension.

Once a specific Service Control application is installed it can be configured by applying a configuration file. The configuration file is application-specific, and is produced by application-specific means, not covered in this documentation. Configuration files have no specific extension.



Note

These configuration changes are automatically saved to the start-up configuration after execution, and therefore do not appear when the running configuration is displayed (**more running-config** command).



Note

These configurations cannot be manipulated by changing the **system/config.txt** file

Managing Application Files

Use the following commands to install, uninstall, and upgrade an application. You can use the `show pqi file info` command before installing or upgrading an application to display the options that are available when installing the `pqi` file. These options can then be specified in the `install` or `upgrade` command as needed.

The documentation of the application will tell the user whether the application is stand-alone (in which case `install` should be used), or an upgrade to an existing application that is assumed to be installed already (in this case `upgrade` should be used).

You should always run the **`pqi uninstall`** command before installing a new `pqi` file. This prevents old files from accumulating on the disk.

The following commands are relevant for installing and uninstalling an application:

- **`pqi install file`** (interface linecard configuration mode)
- **`pqi uninstall file`** (interface linecard configuration mode)
- **`pqi upgrade file`** (interface linecard configuration mode)
- **`pqi rollback file`** (interface linecard configuration mode)
- **`show pqi file info`** (viewer mode)
- **`show pqi last-installed`** (viewer mode)

How to Display Information about an Application File

-
- Step 1** From the `SCE>` prompt, type `show pqi file filename info` and press Enter.
Displays installation options available for this application file, if any.
-

How to Install an Application

`pqi install file`

-
- Step 1** From the `SCE(config if)#` prompt, type `pqi install file filename[options]` and press Enter.
Installs the specified `pqi` file, using the installation options specified (if any). Use the **`show pqi file filename info`** command to display installation options available for the application file. (See [How to Display Information about an Application File](#), page 3-11.)
Note that this may take up to five minutes.



-
- Note** Always run the **`pqi uninstall`** command before installing a new `pqi` file. (See [How to Uninstall an Application](#), page 3-12.)
-

How to Uninstall an Application

pqi uninstall file

- Step 1** From the SCE(config if)# prompt, type `pqi uninstall file filename` and press Enter.
Uninstalls the specified pqi file.
You must specify the last pqi file that was installed.
Note that this may take up to five minutes.
-

How to Upgrade an Application

pqi upgrade file

- Step 1** From the SCE(config if)# prompt, type `pqi upgrade file filename[options]` and press Enter.
Upgrades the current application with the specified pqi file using the installation options specified (if any). Use the **show pqi file filename info** command to display installation options available for the application file. (See [How to Display Information about an Application File, page 3-11.](#))
Note that this may take up to five minutes.
-

How to Undo an Upgrade of an Application

pqi rollback file

- Step 1** From the SCE(config if)# prompt, type `pqi rollback file filename` and press Enter.
Undoes the upgrade of the specified pqi file.
Note that this may take up to five minutes.
-

How to Display the Last pqi File that was Installed

pqi last-installed

- Step 1** From the SCE>prompt, type `pqi last-installed` and press Enter
-

How to Monitor the Operational Status of the SCE Platform

- [Monitoring the Operational Status of the SCE Platform, page 3-13](#)
- [How to Display the Current Operational Status of the SCE Platform, page 3-15](#)
- [Displaying the Current Operational Status of the SCE Platform: Example, page 3-15](#)

Monitoring the Operational Status of the SCE Platform

The following table lists the operational states of the SCE platform. You can monitor the operational status of the SCE platform via:

- The Status LED on the SCE platform front panel
- The **show system operation-status** CLI command

Table 3-1 SCE Platform Operational States

| SCE platform Operational Status | Description | Status LED State |
|---------------------------------|---|------------------|
| Booting | Initial state after reset | Orange |
| Operational | SCE platform becomes operational after completing the following process: <ul style="list-style-type: none"> • Boot is completed • Power self-tests are completed without failure • Platform configuration is applied | Flashing green |

Table 3-1 SCE Platform Operational States (continued)

| SCE platform Operational Status | Description | Status LED State |
|---------------------------------|--|------------------|
| Warning | <p>SCE platform is fully operational (as above) but one of the following occurred:</p> <ul style="list-style-type: none"> • Link on one of the line ports is down • Management port link is down • Temperature raised above threshold • Voltage not in required range • Fans problem • Power supply problem • Insufficient space on the disk <p>Note: If the condition that caused the SCE platform to be in Warning state is resolved (for example, link is up) the SCE platform reverts to Operational state.</p> | Flashing orange |
| Failure | <p>System is in Failure state after Boot due to one of the following conditions:</p> <ul style="list-style-type: none"> • Power on test failure • Three abnormal reboots in less than 20 minutes • Platform configured to enter Failure mode consequent to failure-induced reboot (this is configurable using CLI command) <p>Note: Depending on the cause of failure, the management interface and the platform configuration may or may not be active/available.</p> | Red |

How to Display the Current Operational Status of the SCE Platform

```
show system operation-status
```

Step 1 From the SCE>prompt, type `show system operation-status` and press Enter

Displaying the Current Operational Status of the SCE Platform: Example

This example shows how to display the current operational status of the SCE platform.

```
SCE>show system operation-status
System Operation status is Operational
Port status is:
Link on port #1 is down
Link on port #2 is down
```

How to Display the SCE Platform Version Information

```
show version
```

Displaying the SCE Platform Version Information

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

Step 1 From the SCE>prompt, type `show version` and press Enter

Displaying the SCE Platform Version Information: Example

This example shows how to display the SCE platform version information.

```
SCE>show version
System version: Version 3.0.0 Build 240
Build time: Jan 11 2006, 07:34:47
Software version is: Version 2.5.2 Build 240
Hardware information is:
rx          : 0x0075
dp          : 0x1808
tx          : 0x1708
ff          : 0x0077
cls        : 0x1721
cpld       : 0x0025
Lic        : 0x0176
rev        : G001
Bootrom    : 2.1.0
L2 cache   : Samsung 0.5
lic type   : MFE
optic mode : MM
Product S/N : CAT093604K3
Product ID  : SCE2020-4XGBE-MM
Version ID  : V01
Deviation   :
Part number : 800-26601-01
Revision    : B0
Software revision : G001
LineCard S/N : CAT09370L1Q
Power Supply type : AC
SML Application information is:
Application file: /tffs0/temp.sli
Application name:
Application help:
Original source file: H:\work\Emb\jrt\V2.5\sml\actions\drop\drop_basic_anyflow.san
Compilation date: Wed, November 12 2006 at 21:25:21
Compiler version: SANc v2.50 Build 32 gcc_codelets=true built on: Tue September 23 2006
09:51:57 AM.;SME plugin v1.1
Default capacity option used.
Logger status: Enabled
Platform: SCE 2000 - 4xGBE
Management agent interface version: SCE Agent 3.0.5 Build 18
Software package file: ftp://vk:vk@10.1.8.22/P:/EMB/LatestVersion/3.0.5/se1000.pkg
SCE 2000 uptime is 21 minutes, 37 seconds
SCE>
```

How to Display the SCE Platform Inventory

```
show inventory
```

Displaying the 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, page 5-40](#))
- CLI **show inventory** command

The **show inventory** CLI command displays the following information:

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

Step 1 From the SCE>prompt, type `show inventory` and press Enter

Displaying the SCE Platform Inventory: Example

This example shows how to display the inventory (UDI) of the SCE platform.

```
SCE>show inventory
NAME: "Chassis",
DESCR: "Cisco SCE 2020 Service Control Engine, Multi Mode, 4-port GE"
PID: SCE2020-4XGBE-MM , VID: V01, SN: CAT093604K3
SCE>
```

How to Display the System Uptime

`show system-uptime`

Displaying the System Uptime

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

Step 1 From the SCE>prompt, type `show system-uptime` and press Enter

Displaying the System Uptime: Example

```
SCE>show system-uptime
SCE uptime is 21 minutes, 37 seconds
SCE>
```

How to Reboot and Shut Down the SCE Platform

- [How to Reboot the SCE Platform, page 3-18](#)
- [How to Shut Down the SCE Platform, page 3-18](#)

How to Reboot the SCE Platform

`reload`

Rebooting the SCE Platform

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



Note

When the 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 the Configuration Settings, page 3-4](#).

-
- Step 1** From the SCE# prompt, type `reload` and press Enter
A confirmation message appears.
- Step 2** Type `Y` to confirm the reboot request and press Enter.
-

Rebooting the SCE Platform: Example

```
SCE# reload
Are you sure? y
the system is about to reboot, this will end your CLI session
```

How to Shut Down the SCE Platform

`reload shutdown`

Shutting Down the SCE Platform

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

**Note**

When the 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 the Configuration Settings, page 3-4](#)

-
- Step 1** Connect to the serial console port (The CON connector on the SCE platform front panel, 9600 baud). Provides connection to a local terminal for restarting the SCE platform.
- Step 2** From the SCE# prompt, type `reload shutdown` and press Enter. A confirmation message appears.
- Step 3** Type **Y** to confirm the shutdown request and press Enter. Performs the shutdown operation.
-

Shutting Down the SCE Platform: Examples

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 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 a user issues this command from a Telnet session, and then realizes that he or she has no physical access to the SCE platform.



CHAPTER 4

Utilities

This module describes the following utilities:

- Setup utility command
- File operation utility commands
- User log utility commands

How to Use the Setup Command

- [Setup Command Parameters, page 4-1](#)
- [How to Enter the Setup Command, page 4-4](#)
- [How to Define Lists in the Setup Utility, page 4-5](#)

Setup Command Parameters

The setup utility is an interactive wizard that guides the user through the basic configuration process. This utility runs automatically upon initial connection to the local terminal. It may also be invoked explicitly via Telnet or via the local terminal to make changes to the system configuration.

The following table lists all the command parameters for the setup utility.

Table 4-1 **Setup Command Parameters**

| Parameter | Definition |
|-----------------|---|
| IP address | IP address of the SCE platform. |
| subnet mask | Subnet mask of the SCE platform. |
| default gateway | Default gateway. |
| hostname | Character string used to identify the SCE platform. Maximum length is 20 characters. |
| admin password | Admin level password Character string from 4-100 characters beginning with an alpha character. |

Table 4-1 Setup Command Parameters (continued)

| Parameter | Definition |
|--|---|
| .root password | Root level password. Character string from 4-100 characters beginning with an alpha character. |
| password encryption status | Enable or disable password encryption? |
| Time Settings | |
| time zone name and offset | Standard time zone abbreviation and minutes offset from UTC. |
| local time and date | Current local time and date. Use the format: 00:00:00 1 January 2007 |
| SNTP Configuration | |
| broadcast client status | Sets the status of the SNTP broadcast client. If enabled, the SCE will synchronize its local time with updates received from SNTP broadcast servers. |
| unicast query interval | Interval in seconds between unicast requests for update (64 – 1024) |
| unicast server IP address | IP address of the SNTP unicast server. |
| DNS Configuration | |
| DNS lookup status | Enable or disable IP DNS-based hostname translation. |
| default domain name | Default domain name to be used for completing unqualified host names |
| IP address | IP address of domain name server. (maximum of three servers) |
| RDR Formatter Destination Configuration | |
| IP address | IP address of the RDR-formatter destination |
| TCP port number | TCP port number of the RDR-formatter destination |
| Access Control Lists | |
| Access Control List number | How many ACLs will be necessary? What IP addresses will be permitted/denied access for each management interface? You may want ACLs for the following: <ul style="list-style-type: none"> • Any IP access • Telnet access • SNMP GET access • SNMP SET access |
| list entries (maximum 20 per list) | IP address, and whether permitted or denied access. |

Table 4-1 Setup Command Parameters (continued)

| Parameter | Definition |
|--|--|
| IP access ACL | ID number of the ACL controlling IP access. |
| telnet ACL | ID number of the ACL controlling telnet access. |
| SNMP Configuration | |
| SNMP agent status | SNMP agent status Enable or disable SNMP management. |
| GET community names | Community strings to allow GET access and associated ACLs (maximum 20). |
| SET community names | Community strings to allow SET access and associated ACLs (maximum 20). |
| trap managers (maximum 20) | Trap manager IP address, community string, and SNMP version. |
| Authentication Failure trap status | Sets the status of the Authentication Failure traps. |
| enterprise traps status | Sets the status of the enterprise traps. |
| system administrator | Name of the system administrator. |
| Topology Configuration (All Platforms) | |
| connection mode | Is the SCE platform installed in bump-in-the-wire topology (inline) or out of line using an optical splitter or external switch (receive-only)? |
| Admin status of the SCE platform after abnormal boot | After a reboot due to a failure, should the SCE platform remain in a Failure status or move to operational status provided no other problem was detected? |
| Topology Configuration (SCE 1000) | |
| link bypass mode on operational status | When the SCE 1000 is operational, should it bypass traffic or not? |
| redundant SCE 1000 platform? | Is there a redundant SCE 1000 installed as a backup? |
| link bypass mode on non-operational status | When the SCE 1000 is not operational, should it bypass traffic or cut it off? |
| Topology Configuration (SCE 2000) | |
| type of deployment | Is this a cascade topology, with two SCE platforms connected via the cascade ports? Or is this a single platform topology? |
| physically connected link (cascade topology only) | In a cascade deployment this parameter sets the index for the link that this SCE 2000 is deployed on. The options for the SCE 2000 are link-0 or link-1. In a single-SCE 2000 Platform deployment this parameter is not relevant since one SCE 2000 is deployed on both links. In this case the link connected to port1-port2 is by default link-0 and the link connected to port3-port4 is by default link-1. |

Table 4-1 Setup Command Parameters (continued)

| Parameter | Definition |
|---|---|
| priority (cascade topology only) | If this is a cascaded topology, is this SCE 2000 the primary or secondary SCE 2000? |
| on-failure behavior (inline connection mode only) | If this SCE 2000 is deployed inline, should the failure behavior be bypass or cutoff of the link? |

Information regarding these parameters can be found in the appropriate sections throughout this guide.

For more information regarding SCE platform topology, and for a step-by-step description of the setup utility, see the Cisco SCE 2000/SCE 1000 Installation and Configuration Guides.

How to Enter the Setup Command

setup

Step 1 From the SCE# prompt, type `setup` and press Enter.

The following dialog appears:

```

--- System Configuration Dialog ---
At any point you may enter a question mark '?' followed by 'Enter' for help.
Use ctrl-C to abort configuration dialog at any prompt.
Use ctrl-Z to jump to the end of the configuration dialog at any prompt.
Default settings are in square brackets '[]'.
Would you like to continue with the System Configuration Dialog? [yes/no]: y
system configuration dialog begins.

```

How to Define Lists in the Setup Utility

Multiple entry parameters (Lists)

When explicitly invoked, the setup utility offers the option of multiple entries (lists) for certain parameters.

Several parameters, such as the Access Control Lists, are actually lists containing several entries. If these lists are empty (initial configuration) or contain only one entry, they act the same as any scalar parameter, except that you are given the option of adding additional entries to the list.

If these lists already contain more than one entry, the entire list is displayed, and you are then presented with several options. Following is an excerpt from the SNMP trap manager menu, illustrating how to configure list entries.

Step 1 The entries in the list are displayed.

```
There are 2 SNMP trap managers in the current configuration as follows:
IP address: 10.10.10.10 Community: privateVersion: 1
IP address: 10.11.10.1 Community: pcubeVersion: 2c
```

Step 2 Three options are presented.



Note

If only one entry exists in the table, it is displayed as the default [] to be either accepted or changed. The three list options are not displayed.

```
Please choose one of the following options:
1. Leave the running configuration unchanged.
2. Clear the existing lists and configure new ones.
3. Add new entries.
Enter your choice:
```

Step 3 You are prompted to continue the setup, depending on the choice you entered:

- 1. Leave the running configuration unchanged:
The dialog proceeds to the next question. The list remains unchanged.
 - 2. Clear the existing entries and configure new ones:
The dialog prompts you for a new entry in the list.
After completing the first entry, you are asked whether you would like to add another new entry.
 - Would you like to add another SNMP trap manager? [no]:**y**
Since the list was empty, you may enter the maximum number of entries.
 - 3. Add new entries:
The dialog prompts you for a new entry in the list.
After the completing one entry, you are asked whether you would like add another new entry.
 - Would you like to add another SNMP trap manager? [no]:**y**
You may enter only enough additional entries to reach the maximum number
-

Working with 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 SCE disk do not cause the disk to exceed 70%.

- [Working with Directories, page 4-6](#)
- [Working with Files, page 4-8](#)

Working with Directories

- [How to Create a Directory, page 4-6](#)
- [How to Delete a Directory, page 4-6](#)
- [How to Change Directories, page 4-7](#)
- [How to Display your Working Directory, page 4-7](#)
- [How to List the Files in a Directory, page 4-7](#)

How to Create a Directory

mkdir

-
- Step 1** From the SCE# prompt, type `mkdir directory-name` and press Enter.
-

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

delete

-
- Step 1** From the SCE# prompt, type `delete directory-name /recursive` and press Enter.
- The recursive flag deletes all files and sub-directories contained in the specified directory.
-

How to Delete an Empty Directory

`rmdir`

-
- Step 1** From the SCE# prompt, type `rmdir directory-name` and press Enter.
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`

-
- Step 1** From the SCE# prompt, type `cd new path` and press Enter.
-

How to Display your Working Directory

`pwd`

-
- Step 1** From the SCE# prompt, type `pwd` and press Enter.
-

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

`dir`

-
- Step 1** From the SCE# prompt, type `dir` and press Enter.
-

How to List the Applications in the Current Directory

-
- Step 1** From the SCE# prompt, type `dir applications` and press Enter.
-

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

-
- Step 1** From the SCE# prompt, type `dir -r` and press Enter.
-

Working with Files

- [How to Rename a File, page 4-8](#)
- [How to Delete a File, page 4-8](#)
- [Copying Files, page 4-8](#)
- [How to Display File Contents, page 4-9](#)
- [How to Unzip a File, page 4-9](#)

How to Rename a File

rename

- Step 1** From the SCE# prompt, type `rename current-file-name new-file-name` and press Enter.
-

How to Delete a File

delete

- Step 1** From the SCE# prompt, type `delete file-name` and press Enter.
-

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

copy

- [Copying a File: Example, page 4-8](#)

- Step 1** From the SCE# prompt, type `copy source-file-name destination-file-name` and press Enter.
-

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 an FTP site. In this case, either the source or destination filename must begin with *ftp://*.

-
- Step 1** From the SCE# prompt, type `copy ftp://source destination-file-name` and press Enter.
To upload a file to an FTP site, specify the FTP site as the destination (*ftp://destination*)
-

How to Upload a File to a Passive FTP Site

copy-passive

-
- Step 1** From the SCE# prompt, type `copy-passive source-file-name ftp://destination` and press Enter.
To download a file from a passive FTP site, specify the FTP site as the source (*ftp://source*)
-

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 10.1.1.105, specifying Passive FTP.

```
SCE#copy-passive /appli/analysis.sli ftp://myname:mypw@10.1.1.105/p:/appli/analysis.sli
SCE#
```

How to Display File Contents


more

-
- Step 1** From the SCE# prompt, type `more file-name` and press Enter.
-

How to Unzip a File

unzip

-
- Step 1** From the SCE# prompt, type `unzip file-name` and press Enter.
-



CHAPTER 5

Configuring the Management Interface and Security

This module 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.

- [About Management Interface and Security, page 5-2](#)
- [How to Configure the Management Ports, page 5-2](#)
- [How to Enter Management Interface Configuration Mode, page 5-3](#)
- [How to Configure the Management Port Physical Parameters, page 5-4](#)
- [How to Configure Management Interface Redundancy, page 5-8](#)
- [How to Configure Management Interface Security, page 5-10](#)
- [How to Configure the Available Interfaces, page 5-13](#)
- [How to Configure and Manage the SNMP Interface, page 5-34](#)
- [How to Manage Passwords, page 5-48](#)
- [IP Configuration, page 5-54](#)
- [How to Configure Time Clocks and Time Zone, page 5-60](#)
- [How to Configure SNTP, page 5-66](#)
- [How to Configure Domain Name Server \(DNS\) Settings, page 5-69](#)
- [How to Configure the Management Port Physical Parameters, page 5-73](#)

About Management Interface and Security

The SCE platform is equipped with two RJ-45 management (MNG) ports. These ports provide access from a remote management console to the SCE platform via a LAN.

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

In addition to the Layer 1 security of a backup management link, the Service Control platform provides a further management interface security feature; an IP filter that monitors for various types of TCP/IP attacks. This filter can be configured with thresholds rates both for defining an attack and defining the end of an attack.



Note

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

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
 - Enable IP fragment filtering
 - Configure the permitted and not-permitted IP address monitor

How to Configure the Management Ports

Configuring the Management Ports

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
- Configure redundant management interface behavior (optional):
 - Fail-over mode
- If fail-over mode is disabled, specify the active port (optional).

To configure the system with management interface redundancy, see [How to Configure Management Interface Redundancy, page 5-8](#) Configuring the Management Ports for Redundancy.

-
- Step 1** Cable the desired management port, connecting it to the remote management console via the LAN.
- Step 2** Disable the automatic fail-over mode. (See [How to Disable Automatic Fail-Over Mode, page 5-10.](#))
- Step 3** Configure the management port physical parameters. (See [How to Configure the Management Port Physical Parameters, page 5-4.](#))
-

How to Enter Management Interface Configuration Mode

`interface Mng`

About Management Interface Configuration Mode

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

- 0/1 — Mng port 1
- 0/2 — Mng port 2

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-fail-over**

-
- Step 1** Type `configure` and press Enter.
Enables Global Configuration mode.
The command prompt changes to `SCE(config)#`.
- Step 2** Type `interface Mng {0/1|0/2}` and press Enter.
Enables Management Interface Configuration mode.
The command prompt changes to `SCE(config if)#`
-

How to Configure the Management Port Physical Parameters

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

- [How to Set the IP Address and Subnet Mask of the Management Interface, page 5-4](#)
- [How to Configure the Management Interface Speed and Duplex Parameters, page 5-5](#)
- [How to Specify the Active Management Port, page 5-7](#)

How to Set the IP Address and Subnet Mask of the Management Interface

- [Setting the IP Address and Subnet Mask of the Management Interface, page 5-4](#)
- [Options, page 5-4](#)
- [Setting the IP Address and Subnet Mask of the Management Interface: Example, page 5-5](#)

Setting the IP Address and Subnet Mask of the Management Interface

The user 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.
- **subnet mask** — subnet mask of the management interface.

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.

Step 1 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.



Note 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 SCE platform so that the change will take effect properly in all internal and external components of the SCE platform. (See [How to Reboot and Shut Down the SCE Platform, page 3-18.](#))

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

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

```
SCE(config if)#ip address 10.1.1.1 255.255.0.0
```

How to Configure 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

The following table summarizes the relationship between the interface state and speed and duplex.

Table 5-1 *Interface State Relationship to Speed and Duplex*

| Speed | Duplex | Actual FE Interface State |
|-------|--------|--------------------------------|
| Auto | Auto | Auto negotiation |
| Auto | Full | Auto negotiation |
| Auto | Half | Auto negotiation |
| 10 | Auto | Auto-negotiation (duplex only) |
| 10 | Full | 10 Mbps and full duplex |
| 10 | Half | 10 Mbps and half duplex |
| 100 | Auto | Auto-negotiation (speed only) |
| 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 (see).

-
- Step 1** 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 if)#speed 100
```


How to Configure the Duplex Operation of the Management Interface

Options

The following options are available:

- **duplex** — duplex operation of the currently selected 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 (see).

-
- Step 1** 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 a management port to half duplex mode.

```
SCE(config if)#duplex half
```

How to Specify the Active Management Port

- [Specifying the Active Management Port, page 5-7](#)
- [Options, page 5-8](#)
- [Specifying the Active Management Port: Example, page 5-8](#)

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 fail-over enabled) or not (auto fail-over disabled).

- auto fail-over enabled (automatic mode) — the specified port becomes the currently active port, in effect forcing a fail-over action even if a failure has not occurred.
- auto fail-over 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.



Note

This command is a Privileged Exec command, unlike the other commands in this section, which are Mng Interface Configuration commands. If in Mng interface configuration mode, you must exit to the privileged exec mode and see the SCE# prompt displayed.

Options

The following options are available:

- **slot-number/interface-number** — The interface number (0/1 or 0/2) of the management port that is specified as the active port.

Step 1 From the SCE# prompt, type `Interface Mng {0/1 | 0/2} active-port` and press Enter.
Specify the desired MNG interface.

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# Interface Mng 0/2 active-port
```

How to Configure Management Interface Redundancy

- [About Management Port Redundancy, page 5-9](#)
- [How to Configure the Management Ports for Redundancy, page 5-9](#)
- [How to Configure the Fail-Over Mode, page 5-10](#)

About Management Port Redundancy

The 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 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.

How to Configure 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 fail-over mode.
- See [How to Configure the Fail-Over Mode, page 5-10](#).
- 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 [How to Set 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 [How to Configure the Management Interface Speed and Duplex Parameters, page 5-5](#).
-

How to Configure the Fail-Over Mode

- [Configuring the Fail-Over Mode, page 5-10](#)
- [Options, page 5-10](#)
- [How to Enable Automatic Fail-Over Mode, page 5-10](#)
- [How to Disable Automatic Fail-Over Mode, page 5-10](#)

Configuring the Fail-Over Mode

Use the following command to enable automatic fail-over. 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 fail-over switching mode
 - Default — auto (automatic mode)

How to Enable Automatic Fail-Over Mode

Step 1 From the SCE(config if)# prompt, type `auto-fail-over` and press Enter.

How to Disable Automatic Fail-Over Mode

Step 1 From the SCE(config if)# prompt, type `no auto-fail-over` and press Enter.

How to Configure Management Interface Security

- [About Management Interface Security, page 5-11](#)
- [How to Configure the IP Fragment Filter, page 5-11](#)
- [How to Configure the Permitted and Not-permitted IP Address Monitor, page 5-12](#)
- [How to Monitor Management Interface IP Filtering, page 5-12](#)

About Management Interface Security

Management security is defined as the capability of the SCE platform to cope with malicious management conditions that might lead to global service failure. Resiliency to attacks on the management port includes the following features:

- The SCE platform remains stable during flooding attack.
- The number of TCP/IP stack control protocol vulnerabilities is minimized.
- The availability of reporting capabilities on attacks on the management port.

There are two parallel security mechanisms:

- Automatic security mechanism — monitors the TCP/IP stack rate at 200 msec intervals and throttles the rate from the device if necessary.

This mechanism always functions and is not user-configurable.

- User-configurable security mechanism — accomplished via two IP filters at user-configurable intervals:
 - IP fragment filter — Drops all IP fragment packets
 - IP filter monitor — Measures the rate of accepted and dropped packets for both permitted and not-permitted IP addresses.

How to Configure the IP Fragment Filter

- [Options, page 5-11](#)
- [How to Enable the IP Fragment Filter, page 5-11](#)
- [How to Disable the IP Fragment Filter, page 5-11](#)

Options

The following options are available:

- **enable/disable** — Enable or disable IP fragment filtering
 - Default — disable

How to Enable the IP Fragment Filter

Step 1 From the SCE(config)# prompt, type `ip filter fragment enable` and press Enter.

How to Disable the IP Fragment Filter

Step 1 From the SCE(config)# prompt, type `ip filter fragment disable` and press Enter.

How to Configure the Permitted and Not-permitted IP Address Monitor

Options

The following options are available:

- **ip permitted/ip not-permitted** — Specifies whether the configured limits apply to permitted or not-permitted IP addresses.
If neither keyword is used, it is assumed that the configured limits apply to both permitted and not-permitted IP addresses.
- **low rate** — lower threshold; the rate in Mbps that indicates the attack is no longer present.
 - Default — 20
- **high rate** — upper threshold; the rate in Mbps that indicates the presence of an attack.
 - Default — 20
- **burst size** — duration of the interval in seconds that the high and low rates must be detected in order for the threshold rate to be considered to have been reached
 - Default — 10

Step 1 From the SCE(config)# prompt, type `ip filter monitor {ip_permitted|ip_not_permitted} low_rate low_rate high_rate high_rate burst burst size` and press Enter.

How to Monitor Management Interface IP Filtering

Monitoring Management Interface IP Filtering

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

- IP fragment filter enabled or disabled
- configured attack threshold (permitted and not-permitted IP addresses)
- configured end of attack threshold (permitted and not-permitted IP addresses)
- burst size in seconds (permitted and not-permitted IP addresses)

Step 1 From the SCE>prompt, type `show ip filter` and press Enter.

How to Configure the Available Interfaces

The system allows you to configure the Telnet and SNMP interfaces according to the manner in which you are planning to manage the SCE platform and the external components of the system.

- [How to Configure TACACS+ Authentication, Authorization, and Accounting, page 5-13](#)
- [How to Configure Access Control Lists \(ACLs\), page 5-27](#)
- [How to Configure the Telnet Interface, page 5-30](#)
- [How to Configure the SSH Server, page 5-31](#)
- [Enabling the SNMP Interface, page 5-34](#)

How to Configure TACACS+ Authentication, Authorization, and Accounting

- [Information About TACACS+ Authentication, Authorization, and Accounting, page 5-13](#)
- [How to Configure the SCE Platform TACACS+ Client, page 5-17](#)
- [How to Manage the User Database, page 5-20](#)
- [How to Configure AAA Login Authentication, page 5-24](#)
- [How to Configure AAA Privilege Level Authorization Methods, page 5-25](#)
- [How to Configure AAA Accounting, page 5-26](#)
- [How to Monitor TACACS+ Servers, page 5-27](#)
- [How to Monitor TACACS+ Users, page 5-27](#)

Information About TACACS+ Authentication, Authorization, and Accounting

- [TACACS+ Authentication, Authorization, and Accounting, page 5-14](#)
- [Login Authentication, page 5-14](#)
- [Accounting, page 5-15](#)
- [Privilege Level Authorization, page 5-15](#)
- [General AAA Fallback and Recovery Mechanism, page 5-16](#)
- [About Configuring TACACS+, page 5-16](#)

TACACS+ Authentication, Authorization, and Accounting

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 SCE platform, providing a secure means of managing the 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 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 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 SCE platform user log and the telnet session is terminated (unless the user is connected to the console port.)

The 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 SCE platform. If an ERROR response is received, the 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, page 5-16](#).

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 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 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 SCE platform sends an authentication request to the TACACS+ server specifying the requested privilege level. The 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, page 5-16](#).

General AAA Fallback and Recovery Mechanism

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

The 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**



Note

Important: If the server goes to AAA fault, the 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 SCE platform becomes un-accessible, the shell function "AAA_MethodsReset" will allow the user to delete the current AAA method settings and set the AAA method used to "Enable".

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.

2. For complete details on server configuration, refer to the appropriate configuration guide for the particular TACACS+ server that you will be using.

3. Configure the 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.)

4. (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 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
5. Configure the authentication methods on the SCE platform.
 - login authentication methods
 - privilege level authorization methods
 6. Review the configuration.

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

How to Configure the SCE Platform TACACS+ Client

- [Configuring the SCE Platform TACACS+ Client, page 5-17](#)
- [How to Add a new TACACS+ Server Host, page 5-18](#)
- [How to Remove a TACACS+ Server Host, page 5-18](#)
- [How to Configure the Global Default Key, page 5-19](#)
- [How to Configure the Global Default Timeout, page 5-19](#)

Configuring the SCE Platform TACACS+ Client

The user must configure the remote servers for the TACACS+ protocol. Then the 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 SCE platform TACACS+ client are explained in the following sections:

- [How to Monitor Management Interface IP Filtering, page 5-12](#)
- [How to Remove a TACACS+ Server Host, page 5-18](#)
- [How to Configure the Global Default Key, page 5-19](#)
- [How to Configure the Global Default Timeout, page 5-19](#)

How to Add a new TACACS+ Server Host

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

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

- [Options, page 5-18](#)

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 [How to Define the Global Default Timeout, page 5-20.](#))
- **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 [How to Define a Global Default Key, page 5-19.](#))

Step 1 From the SCE(config)# prompt, type `TACACS-server host host-name [port portnumber] [timeout timeout-interval] [key key-string]` and press Enter.

How to Remove a TACACS+ Server Host

- [Options, page 5-19](#)

Options

The following options are available:

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

Step 1 From the SCE(config)# prompt, type `no TACACS-server host host-name` and press Enter.

How to Configure 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, page 5-19](#)
- [How to Define a Global Default Key, page 5-19](#)
- [How to Clear the Global Default Key, page 5-19](#)

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

How to Define a Global Default Key

Step 1 From the SCE(config)# prompt, type `TACACS-server key key-string` and press Enter.

How to Clear the Global Default Key

Step 1 From the SCE(config)# prompt, type `no TACACS-server key` and press Enter.

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.

How to Configure 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, page 5-20](#)
- [How to Define the Global Default Timeout, page 5-20](#)
- [How to Clear the Global Default Timeout, page 5-20](#)

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

How to Define the Global Default Timeout

-
- Step 1** From the SCE(config)# prompt, type `TACACS-server timeout timeout-interval` and press Enter.
-

How to Clear the Global Default Timeout

-
- Step 1** From the SCE(config)# prompt, type `no TACACS-server timeout` and press Enter.

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.

How to Manage 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 configured as encrypted or unencrypted
- Privilege level

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

- [How to Add a New User to the Local Database, page 5-20](#)
- [How to Define the User Privilege Level, page 5-22](#)
- [How to Add a New User with Privilege Level and Password, page 5-22](#)
- [How to Delete a User, page 5-24](#)

How to Add 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.

- [Options, page 5-21](#)
- [How to Add a User with a Clear Text Password, page 5-21](#)
- [How to Add a User with No Password, page 5-21](#)
- [How to Add a User with an MD5 Encrypted Password Entered in Clear Text, page 5-21](#)
- [How to Add a User with an MD5 Encrypted Password Entered as an MD5 Encrypted String, page 5-22](#)

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

Step 1 From the SCE(config)# prompt, type `username name password password` and press Enter.

How to Add a User with No Password

Step 1 From the SCE(config)# prompt, type `username name nopassword` and press Enter.

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

Step 1 From the SCE(config)# prompt, type `username name secret 0 password` and press Enter.

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

Step 1 From the SCE(config)# prompt, type `username name secret 5 encrypted-secret` and press Enter.

How to Define the User Privilege Level

- [About the User Privilege Level, page 5-22](#)
- [Options, page 5-22](#)

About the User Privilege Level

Privilege level authorization in the 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 SCE platform sends an authentication request to the TACACS+ server specifying the requested privilege level. The 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 to 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

Step 1 From the SCE(config)# prompt, type `username name privilege level` and press Enter.

How to Add 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.

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

-
- Step 1** From the SCE(config)# prompt, type `username name privilege level password password` and press Enter.
-

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

-
- Step 1** From the SCE(config)# prompt, type `username name privilege level secret 0 password` and press Enter.
-

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

-
- Step 1** From the SCE(config)# prompt, type `username name privilege level secret 5 encrypted-secret` and press Enter.
-

How to Delete a User

Options

The following options are available:

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

Step 1 From the SCE(config)# prompt, type `no username name` and press Enter.

How to Configure 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](#).)

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

- [Configuring Maximum Login Attempts, page 5-24](#)
- [How to Configure the Login Authentication Methods, page 5-24](#)

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

Step 1 From the SCE(config)# prompt, type `aaa authentication attempts login number-of-attempts` and press Enter.

How to Configure the 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, page 5-16](#)).

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

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

Step 1 From the SCE(config)# prompt, type `aaa authentication login default method1 [method2...]` and press Enter.

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

Step 1 From the SCE(config)# prompt, type `no aaa authentication login default` and press Enter.

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.

How to Configure AAA Privilege Level Authorization Methods

- [Options, page 5-25](#)
- [How to Specify AAA Privilege Level Authorization Methods, page 5-26](#)
- [How to Delete the AAA Privilege Level Authorization Methods List, page 5-26](#)

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

Step 1 From the SCE(config)# prompt, type `aaa authentication enable default method1 [method2...]` and press Enter.

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

Step 1 From the SCE(config)# prompt, type `no aaa authentication enable default` and press Enter.

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.

How to Configure AAA Accounting

Use this command to enable or disable TACACS+ accounting.

- [About AAA Accounting, page 5-26](#)
- [Options, page 5-26](#)
- [How to Enable AAA Accounting, page 5-26](#)
- [How to Disable AAA Accounting, page 5-27](#)

About AAA Accounting

If TACACS+ accounting is enabled, the 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

Step 1 From the SCE(config)# prompt, type `aaa authentication accounting commands level default stop-start group TACACS+` and press Enter.

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

-
- Step 1** From the SCE(config)# prompt, type `aaa authentication accounting commands leveldefault` and press Enter.
-

How to Monitor TACACS+ Servers

Use these commands to display statistics for the TACACS+ servers.

- [How to Display Statistics for TACACS+ Servers, page 5-27](#)
- [How to Display Statistics, Keys and Timeouts for TACACS+ Servers, page 5-27](#)

How to Display Statistics for TACACS+ Servers

-
- Step 1** From the SCE# prompt, type `show TACACS` and press Enter.
-

How to Display Statistics, Keys and Timeouts for TACACS+ Servers

-
- Step 1** From the SCE# prompt, type `show TACACSall` and press Enter.

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.

How to Monitor TACACS+ Users

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

-
- Step 1** From the SCE# prompt, type `show users` and press Enter.

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.

How to Configure Access Control Lists (ACLs)

- [About Access Control Lists, page 5-28](#)
- [Options, page 5-29](#)
- [How to Add Entries to an ACL, page 5-29](#)
- [How to Remove an ACL, page 5-29](#)
- [How to Define a Global ACL, page 5-30](#)

About Access Control Lists

The SCE platform can be configured with Access Control Lists (ACLs), which are used to permit or deny incoming connections on any of the management interfaces. 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. ([How to Add Entries to an ACL, page 5-29](#)).
2. Associating the access list with a management interface. (See [How to Define a Global ACL, page 5-30](#) and [How to Assign an ACL to the Telnet Interface, page 5-31.](#))

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 access list 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 access list, access is denied.

You can create up to 99 access lists. Access lists 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 SCE platform first checks if there is permission for access from that IP address. If not, the SCE does not respond to the request. Configuring the 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 and SNMP. The basic IP interface is low-level, blocking the IP packets before they reach the interfaces.
- Interface level: Access to each management interface (Telnet, SNMP, etc.) can be restricted to an access list. 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 access list is associated with a specific management interface, that interface checks the access list to find out if there is permission for a specific external IP address trying to access the management interface.

It is possible to configure several management interfaces to the same access list, if this is the desired behavior of the SCE platform.

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



Note

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

How to Add 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
```

How to Remove an ACL

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

Step 1 From the SCE(config)# prompt, type `no access-list number`, and press Enter.

Removes the specified ACL with all its entries.

How to Define a Global ACL

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

-
- Step 1** From the SCE(config)# prompt, type `ip access-class number`, and press Enter.
- Applies the specified ACL to all traffic attempting to access the SCE platform, rather than to a specific type of traffic, such as Telnet traffic.
-

How to Configure the Telnet Interface

- [About the Telnet Interface, page 5-30](#)
- [How to Prevent Telnet Access, page 5-30](#)
- [How to Assign an ACL to the Telnet Interface, page 5-31](#)
- [How to Configure the Telnet Timeout, page 5-31](#)

About the Telnet Interface

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

How to Prevent Telnet Access

Use this command to disable access by Telnet altogether.

-
- Step 1** From the SCE(config)# prompt, type `no service telnetd`, and press Enter.
- Current Telnet sessions are not disconnected, but no new Telnet sessions are allowed.
-

How to Assign an ACL to the Telnet Interface

-
- Step 1** From the SCE(config)# prompt, type `line vty 0` , and press Enter.
Enables Line Configuration mode.
- Step 2** Type `access-class acl-number in`.
`acl-number` is the ID number of an existing access list.
-

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
```

How to Configure the Telnet Timeout

The 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

-
- Step 1** From the SCE(config-line)# prompt, type `timeout timeout`, and press Enter.
-

How to Configure the SSH Server

- [Information About the SSH Server, page 5-32](#)
- [How to Manage the SSH Server, page 5-32](#)
- [How to Monitor the Status of the SSH Server, page 5-34](#)

Information About the SSH Server

- [The SSH Server, page 5-32](#)
- [Key Management, page 5-32](#)

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 SSH-1 and SSH-2 protocols.

An Access Control List (ACL) can be configured for SSH as for any other management protocol, limiting SSH access to a specific set of IP addresses (see [How to Configure Access Control Lists \(ACLs\), page 5-27](#)).

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 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.

How to Manage 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
- Assign/remove an ACL to the SSH server
- Delete existing SSH keys

How to Generate a Set of SSH Keys

Remember that you must generate a set of SSH keys before you enable the SSH server.

Step 1 From the SCE(config)# prompt, type `ip ssh key generate` and press Enter.

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.

How to Enable the SSH Server

Step 1 From the SCE(config)# prompt, type `ip ssh` and press Enter.

How to Disable the SSH Server

Step 1 From the SCE(config)# prompt, type `no ip ssh` and press Enter.

How to Assign an ACL to the SSH Server

Step 1 From the SCE(config)# prompt, type `ip ssh access-class acl-number` and press Enter.
Assigns the specified ACL to the SSH server, so that access the SSH server is limited to the IP addresses defined in the ACL.

How to Remove the ACL Assignment from the SSH Server

Step 1 From the SCE(config)# prompt, type `no ip ssh access-class` and press Enter.
Removes the ACL assignment from the SSH server, so that any IP address may now access the SSH server.

How to Delete the Existing SSH Keys

Step 1 From the SCE(config)# prompt, type `ip ssh key remove` and press Enter.

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 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 SCE platform is restarted (using **reload**):

- Generate a new set of keys.
 - Disable the SSH server and save the configuration.
-

How to Monitor the Status of the SSH Server

Use this command to monitor the status of the SSH sever, including current SSH sessions.

Step 1 From the SCE>prompt, type `show ip ssh` and press Enter.

This is a User Exec command. Make sure that you are in User Exec command mode by exiting any other modes.

Enabling the SNMP Interface

Use this command to enable the SNMP interface. For more information on configuring and managing the SNMP parameters, including hosts, communities, contact, location, and trap destination hosts, see [How to Configure and Manage the SNMP Interface, page 5-34](#).

- [How to Enable the SNMP Interface, page 5-34](#)
- [How to Disable the SNMP Interface, page 5-34](#)

How to Enable the SNMP Interface

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

Step 1 From the SCE(config)# prompt, type `snmp-server community community-string` and press Enter.

You must define at least one community string to allow SNMP access. For complete information on community strings see [How to Configure SNMP Community Strings, page 5-43](#).

How to Disable the SNMP Interface

Step 1 From the SCE(config)# prompt, type `no snmp-server` and press Enter.

How to Configure and Manage the SNMP Interface

- [Information About the SNMP Interface, page 5-35](#)
- [How to Configure SNMP Community Strings, page 5-43](#)
- [How to Configure SNMP Notifications, page 5-45](#)

Information About the SNMP Interface

This section explains how to configure the SNMP agent parameters. It also provides a brief overview of SNMP notifications and the supported MIBs, and explains the order in which the MIB must be loaded.

- [The SNMP Interface, page 5-35](#)
- [SNMP Protocol, page 5-35](#)
- [Security Considerations, page 5-36](#)
- [Information About CLI, page 5-36](#)
- [Information About MIBs, page 5-37](#)
- [Configuration via SNMP, page 5-43](#)

The SNMP Interface

The SCE platform operating system includes a Simple Network Management Protocol (SNMP) agent that supports the following:

- RFC 1213 standard (MIB-II)
- RFC 2737 standard (ENTITY-MIB version 2)
- *pcube* enterprise MIBs

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.

SCE platform supports the original SNMP protocol (also known as SNMPv1), and a newer version called Community-based SNMPv2 (also known as SNMPv2C).

- **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** — is the revised protocol, which 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.

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 and SNMPv2C specifications define the following basic operations that are supported by SCE platform:

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 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 | 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 newly defined 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 SCE platform supports the option to configure community of managers for read-write accessibility or for read-only accessibility. Furthermore, an ACL (Access List) may be associated with a community to allow SNMP management to a restricted set of managers IP addresses.

Information About CLI

- [CLI, page 5-36](#)
- [CLI Commands for Configuring SNMP, page 5-37](#)
- [CLI Commands for Monitoring SNMP, page 5-37](#)

CLI

The 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 SNMP. 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

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 (also available when SNMP agent is disabled)
- show snmp community
- show snmp contact
- show snmp enabled
- show snmp host
- show snmp location
- show snmp mib
- show snmp traps

Information About MIBs

- [MIBs, page 5-38](#)
- [MIB Data Objects, page 5-38](#)
- [Information About MIB-II, page 5-38](#)
- [Information About ENTITY-MIB, page 5-40](#)
- [Information About pcube Enterprise MIB, page 5-41](#)

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.

The SCE platform supports the following MIBs:

- Standard MIBs:
 - MIB-II (as defined in RFC 1213, Management Information Base for Network Management of TCP/IP-based Internets) and some of its extensions.
 - ENTITY-MIB version 2 (as defined in RFC 2737)
- Proprietary MIBs – *Pcube enterprise MIBs* defined by Cisco, for the Cisco Service Control products (see [Proprietary MIB Reference](#)).

Pcube enterprise MIB (pcube) can be divided into different kinds of MIBs:

- Proprietary SCOS MIBs – These MIBs contain platform specific information. They also contain the generic definitions of the pcube subtree.

The *SE* MIB and the *Dispatcher* MIB are two examples of OS MIBs.

- Proprietary Application MIB(s) – These MIBs contain application specific information.

Currently, there is one application MIB – *Engage* MIB.

- Proprietary Common MIB(s) – These MIBs contain functionality that is common across more than a single Cisco platform.

Currently there is one common MIB – *configuration copy* MIB.

Since the acquisition of P-cube, Inc by **Cisco Systems, Inc**, the existing proprietary MIBs have undergone a process of updating to make them conform to Cisco standards. Note that all Pcube MIBs since SCOS version 3.0.3 are compiled using SMICNG and are in conformation with Cisco standards and styling.

**Note**

While the designations "Pcube" and "SC" have been retained in the MIB for the sake of consistency, they refer to the corresponding Cisco SCE products.

MIB Data Objects

The data objects that make up the MIB may be identified in two ways:

- OID (Object Identifier) — The unique string that describes a specific data object in the agent database.
- OIDs are written in dotted format such as: 1.3.6.1.4.1.5655.4.1.10.1
- MIB descriptor — A name defined in the MIB file for the OID. It is often used instead of the explicit OID.

For instance: "ifTable" stands for the OID of the MIB-II interface table.

Information About MIB-II

- [MIB-II, page 5-39](#)
- [IF-MIB, page 5-39](#)

MIB-II

SCE platform fully supports MIB-II (RFC1213), including the following groups

- System Interface (for both the management and line ports) AT (management port) IP (management port) ICMP (management port)
- TCP (management port)
- UDP (management port)
- SNMP (management port)

IF-MIB

The MIB-II standard has been extended by several different MIBs. The SCOS supports the IF-MIB, defined in RFC-2233.

The IF-MIB defines the following four tables:

| | |
|-------------------|---|
| iftable | An update to the MIB-II ifTable |
| ifxtable | An addition to the ifTable, intended for high capacity interfaces |
| ifStackTable | A table containing information about sublayers of interfaces |
| ifRcvAddressTable | A table meant for interfaces that support more than one receive address |

These are the details of specific objects in this MIB:

| | |
|------------------------|--|
| ifindex | The numbering of the interfaces is such that the port(s) come first. |
| ifPhysAddress | For Management interfaces, this is the MAC address. For traffic interfaces, this is an all zeros address. |
| IfAdminStatus | Write operation to this object is not supported. This is OK according to Ethernet MIB RFC2665 section 3.2.7 |
| IfOutQLen | Always returns 0. |
| Under ifXTable: ifname | The same as ifDescr. |
| ifpromiscuousmode | Management interface – “false”. Traffic interfaces – “true”. |
| ifRcvAddressTable | Not implemented |
| iftesttable | Was deprecated by RFC-2233, and is therefore not implemented |

Information About ENTITY-MIB

- [ENTITY-MIB, page 5-40](#)
- [entityPhysical group, page 5-40](#)
- [entityGeneral group, page 5-41](#)

ENTITY-MIB

The Entity-MIB contains five groups of MIB objects:

- entityPhysical group
- entityLogical group
- entityMapping group
- entityGeneral group
- entityNotifications group

The SCOS implements only the physical and the general groups of the Entity-MIB, since the other groups are not relevant to the SCE platform.

entityPhysical group

The entityPhysical group describes the physical entities managed by a single agent. It contains a single table, the *entPhysicalTable*, that identifies physical system components.

These are the details of specific objects in the *entPhysicalTable*, as implemented in SCOS:

| | |
|-----------------------------|--|
| entPhysicalIndex (1) | 1 (SCE main board) |
| entPhysicalDescr (2) | The description corresponding to the Product ID, as it appears in the product catalog. |
| entPhysicalVendorType (3) | cevChassisSCE2000 = {cevChassis 511} (1.3.6.1.4.1.9.12.3.1.3.511) cevChassisSCE1000 = {cevChassis 512} (1.3.6.1.4.1.9.12.3.1.3.512) |
| entPhysicalContainedIn (4) | 0 (not contained) |
| entPhysicalClass (5) | 3 (chassis) |
| entPhysicalParentRelPos (6) | 1 |
| entPhysicalName (7) | "Chassis" |
| entPhysicalHardwareRev (8) | Version ID, as identified in EPROM |
| entPhysicalFirmwareRev (9) | empty string |
| entPhysicalSoftwareRev (10) | Software version, as seen in "show version" |
| entPhysicalSerialNum (11) | Serial number, as identified in EPROM |
| entPhysicalMfgName (12) | "Cisco Systems, inc." |
| entPhysicalModelName (13) | Product ID, as identified in EPROM |
| entPhysicalAlias (14) | empty string |
| entPhysicalAssetID (15) | empty string |
| entPhysicalIsFRU (16) | 2 (false) |

entityGeneral group

The entityGeneral group contains general information relating to the other object groups. The entGeneral group contains a single scalar object.

These are the details of specific object in the *entityGeneral* group, as implemented in SCOS:

| | |
|-------------------|---|
| entLastChangeTime | sysUpTime. This reflects the fact that the entries in the Entity-MIB do not change in the SCE platform after their creation at boot time. |
|-------------------|---|

Information About pcube Enterprise MIB

- [pcube Enterprise MIB, page 5-41](#)
- [Loading the MIB Files, page 5-42](#)

pcube Enterprise MIB

The SCE proprietary *pcube* MIB enables external management systems to retrieve general information regarding the SCE platform operating status and resources utilization, extract real time measurements of bandwidth utilization and network statistics, and receive notifications of critical events and alarms.

**Note**

The following object identifier represents the pcube Enterprise MIB:

**Note**

1.3.6.1.4.1.5655, or iso.org.dod.internet.private.enterprise.pcube

The *pcube* Enterprise MIB splits into four main groups:

- Products
- Modules
- Management
- Workgroup

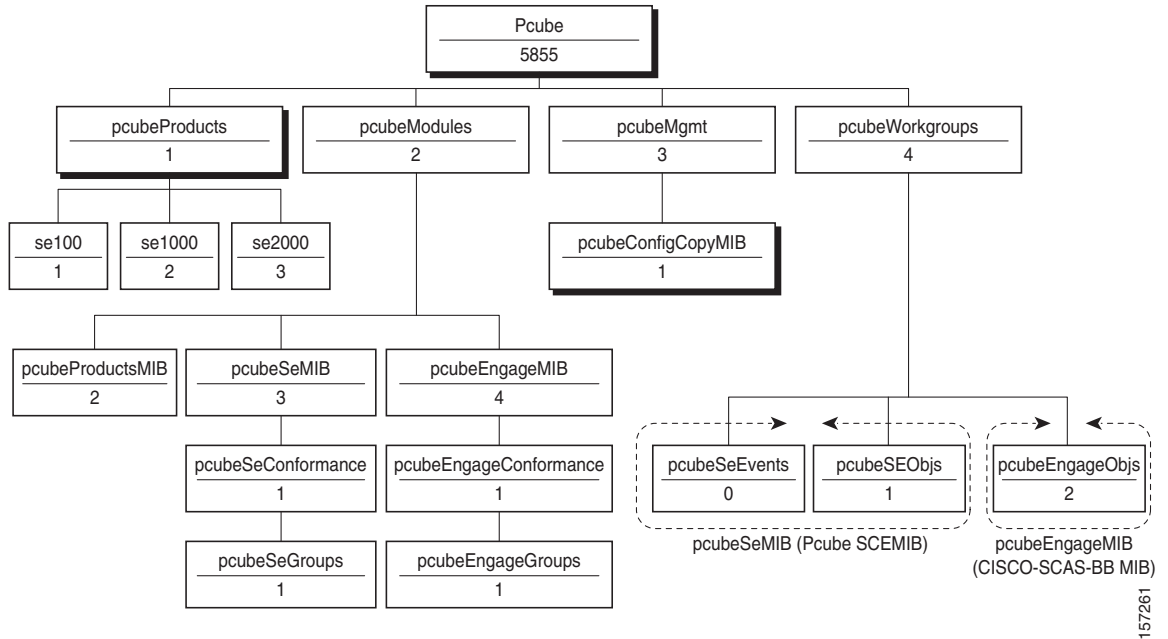
The *pcube* enterprise tree structure is defined in a MIB file named *pcube.mib*.

Refer to the [Proprietary MIB Reference](#) for a complete description of the *pcube* enterprise MIB.

The figure below illustrates the *pcube* Enterprise MIB structure. Conventions used in the diagram:

- Dotted arrows surrounding a unit or units indicate that the component is described in the MIB file specified below the line.
- A shadowed box indicates that the component is described in its own MIB file.

Figure 5-1 Pcube Enterprise MIB Structure



- *pcubeProducts* subtree — contains the OIDs of Cisco Service Control products.
- *pcubeModules* subtree — provides a root object identifier under which MIB modules can be defined.
- *pcubeMgmt* subtree — the root for *pcube* MIBs that are relevant to multiple products.
 - *pcubeConfigCopyMIB* — a subset of the Cisco Config-Copy-MIB that supports local copying of running config to startup config.
- *pcubeWorkgroups* subtree — contains the actual MIBs for Cisco Service Control devices and sub-devices.
- *pcubeSeMIB* — comprises two branches:
 - *pcubeSeEvents* — Contains the OIDs used for sending enterprise-specific notifications.
 - *pcubeSEObjs* — Contains the OIDs that belong to the SCE platform, divided into groups according to functionality.

Loading the MIB Files

The Service Control proprietary MIB uses definitions that are defined in other MIBs, such as *pcube* MIB (*pcube.mib*), and the *SNMPv2.mib*. Therefore, the order in which the MIBs are loaded is important. To avoid errors, the MIBs must be loaded in the proper order.

1. Load *SNMPv2.my*.
2. Load *SNMP-FRAMEWORK-MIB.my*.
3. Load *PCUBE-SMI.my*.
4. Load *PCUBE-SE-MIB.my*.



Note

Information and proprietary MIB files supported by the SCOS can be downloaded from: <http://www.cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml> under the Cisco Service Routing Products section.

Configuration via SNMP

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 (see [Figure 5-1Pcube Enterprise MIB Structure, page 5-42](#)).

It should be noted also that the 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.

How to Configure SNMP Community Strings

- [Configuring SNMP Community Strings, page 5-43](#)
- [How to Define a Community String, page 5-43](#)
- [How to Remove a Community String, page 5-44](#)
- [How to Display the Configured Community Strings, page 5-45](#)

Configuring SNMP Community Strings

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.

How to Define a Community String

- [Options, page 5-44](#)
- [Defining a Community String: Example, page 5-44](#)

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** — ID number of the Access Control List to be assigned to the SNMP interface. It should list the IP addresses of the SNMP managers permitted to use the community string to gain access to the agent.

If no ACL is specified, all IP addresses can access the agent using the defined community string. For more information about ACLs, see [How to Configure Access Control Lists \(ACLs\)](#), page 5-27

The following keywords are available:

- **ro** — read only (default accessibility)
- **rw** — read and write

Step 1 From the SCE(config)# prompt, type `snmp-server community community-stringrolrw acl-number` and press Enter.

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 and access list number “1”.

Since read-only is the default, it does not need to be defined explicitly.

```
SCE(config)#snmp-server community mycommunity 1
```

How to Remove a Community String

Step 1 From the SCE(config)# prompt, type `no snmp-server community community-string` and press Enter.

Removing a Community String: Example

The following example shows how to remove a community string called “mycommunity”.

```
SCE(config)#no snmp-server community mycommunity
```

How to Display the Configured Community Strings

Step 1 From the SCE>prompt, type `show snmp-server community community-string` and press Enter.

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>
```

How to Configure SNMP Notifications

Use these commands to configure:

- The destinations that will receive SNMP notifications (hosts)
- Which types of notifications will be sent (traps)

About SNMP Notifications

Notifications are unsolicited messages that are generated by the SNMP agent that resides inside the 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 SCE platform is not configured to send any SNMP notifications. You must define the Network Management System to which the 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 SCE platform, an SNMP notification is sent from the SCE platform to the list of IP addresses that you define.

SCE platform supports two general categories of notifications:

- Standard SNMP notifications — As defined in RFC1157 and using the conventions defined in RFC1215.
- Proprietary SCE enterprise notifications — As defined in the SCE proprietary MIB (see [\\$\\$\\$LINK](#) in MIB Ref\$\$\$Notification Types).

After a host or hosts are configured to receive notifications, by default, the SCE platform sends to the host or hosts all the notifications supported by the SCE platform except for the AuthenticationFailure notification. The SCE platform provides the option to enable or disable the sending of this notification, as well as some of the SCE enterprise notifications, explicitly.

SCE platform can be configured to generate either SNMPv1 style or SNMPv2c style notifications. By default, the SCE platforms sends SNMPv1 notifications.

Following are some sample procedures illustrating how to do the following:

- 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

How to Define SNMP Hosts

Use this command to define the hosts that will receive notifications from the SCE platform.

- [Options, page 5-46](#)
- [How to Configure the SCE Platform to Send Notifications to a Host \(NMS\), page 5-46](#)
- [How to Configure the SCE Platform to Stop Sending Notifications to a Host, page 5-47](#)
- [How to Configure SNMP Traps, page 5-47](#)

Options

The following options are available:

- **ip-address** — the IP address of the SNMP server host
- **community-string** — a 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 or 2c.
 - Default — 1 (SNMPv1)

How to Configure the SCE Platform to Send Notifications to a Host (NMS)

Step 1 At the SCE(config)# prompt, type `snmp-server host ip-address community-string` and press Enter.

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 SCE Platform to Send Notifications to Multiple Hosts: Example

The following example shows how to configure the SCE platform to send SNMPv1 notifications to several hosts.

```
SCE(config)#snmp-server host 10.10.10.10 mycommunity
SCE(config)#snmp-server host 20.20.20.20 mycommunity
SCE(config)#snmp-server host 30.30.30.30 mycommunity
SCE(config)#snmp-server host 40.40.40.40 mycommunity
```


How to Configure the SCE Platform to Stop Sending Notifications to a Host

Step 1 At the SCE(config)# prompt, type `no snmp-server host ip-address` and press Enter.

Configuring the 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
```

How to Configure SNMP Traps

Use this command to configure the notifications that will be sent to the defined host.

- [Options, page 5-47](#)
- [How to Enable the SNMP Server to Send Authentication Failure Notifications, page 5-48](#)
- [How to Enable the SNMP Server to Send All Enterprise Notifications, page 5-48](#)
- [How to Enable the SNMP Server to Send a specific Enterprise Notification, page 5-48](#)
- [How to Restore All Notifications to the Default Status, page 5-48](#)

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, SNTTP, 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

Step 1 At the SCE(config)# prompt, type `snmp-server enable traps snmp authentication` and press Enter.

How to Enable the SNMP Server to Send All Enterprise Notifications

Step 1 At the SCE(config)# prompt, type `snmp-server enable traps enterprise` and press Enter.

How to Enable the SNMP Server to Send a specific Enterprise Notification

Step 1 At the SCE(config)# prompt, type `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]` and press Enter.
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

Step 1 At the SCE(config)# prompt, type `default snmp-server enable traps` and press Enter.
Resets all notifications supported by the SCE platform to their default status.

How to Manage Passwords

- [About Passwords, page 5-49](#)
- [Changing Your Password, page 5-49](#)
- [Password Encryption, page 5-50](#)
- [Password Recovery, page 5-51](#)

About Passwords

Cisco CLI passwords are an access-level authorization setting, not individual user passwords. All Admin users, for example, log in with the same password. This means that the system does not identify you as an individual, but as a user with certain privileges.

Passwords are needed for all authorization levels to prevent unauthorized users from accessing the SCE platform. It is highly recommended that you change the default password upon initial installation, and that you change the passwords periodically to secure the system.

**Note**

The default password for all levels is “Cisco”.

When a telnet user logs on, he sees only a Password: prompt, no logo is displayed. This provides extra security by not revealing the system identity to users that do not know the password.

Password guidelines:

- Password length must be between 4 and 100 characters long.
- Passwords can contain any visible keyboard character.
- Passwords must begin with a letter.
- Passwords cannot contain spaces.
- Passwords are case-sensitive.

Users with Admin or higher authorization level can view the configured passwords using the **show running-config** or the **show startup-config** commands. Therefore, if you want passwords to remain completely confidential, you must activate the encryption feature, described in [How to Enable Password Encryption, page 5-51](#).

Changing Your Password

- [How to Change Your Password, page 5-49](#)
- [Verifying that the Password has been Successfully Changed, page 5-50](#)

How to Change Your Password

Use the **enable password** command to change the password. Note that if the password has been changed, the default password will no longer be accepted.

Options

The following options are available:

- level — the number of the authorization level for which the password is defined.
The authorization levels are:
 - 0: user
 - 5: viewer
 - 10: admin
 - 15: root
- password — the new password for the specified authorization level.

-
- Step 1** At the SCE(config)# prompt type `enable password level level password` and press Enter.
The Network Administrator should record passwords in a secure location.
-

Verifying that the Password has been Successfully Changed

-
- Step 1** Initiate a new telnet connection, while maintaining the one you used to set the password.
This is important so that if the verification fails, you would still have admin level authorization to re-enter the password.
- Step 2** At the SCE# prompt type `enable level` and press Enter.
Specify the level for which you have just changed the password.
- Step 3** At the prompt, type your new password and press Enter.
If your new password has been entered successfully, then the appropriate prompt appears.
If you enter an incorrect password, the password prompt will appear again.
Repeat these steps as necessary to check additional passwords.
The encryption feature will encrypt the passwords in the platform configuration files.
-

Password Encryption

- [About Password Encryption, page 5-51](#)
- [How to Enable Password Encryption, page 5-51](#)
- [How to Disable Password Encryption, page 5-51](#)

About Password Encryption

Once the encryption feature is activated, passwords entered into the system are encrypted to the startup configuration file the next time the configuration is saved. When encryption feature is turned off, passwords previously encrypted to the startup configuration file are not deciphered.

By default, the password encryption feature is disabled.

How to Enable Password Encryption

Step 1 From the `sce(config)#` prompt, type `service password encryption` and press Enter.

How to Disable Password Encryption

Step 1 From the `sce(config)#` prompt, type `no service password encryption` and press Enter.

This does not remove the encryption from the configuration file. You must save to the startup configuration file if you want the password to be stored un-encrypted on the startup configuration file.

Password Recovery

Use the following procedures if it becomes necessary to recover the **enable** passwords for the SCE platform. Be sure to use the appropriate procedure depending on the version of SCOS you are using:

- [How to Recover the Passwords: SCOS versions before 2.5.5, page 5-51](#)
- [How to Recover the Passwords: SCOS versions 2.5.5 or later, page 5-54](#)

How to Recover the Passwords: SCOS versions before 2.5.5

The procedure used for recovering the passwords depends on whether or not there are user-configured parameters that need to be saved.

- [How to Recover the Passwords: Reverting to Default the Configuration, page 5-52](#)
- [How to Recover the Passwords: Saving the Current Configuration, page 5-52](#)

How to Recover the Passwords: Reverting to Default the Configuration

For SCE platforms running a SCOS version before 2.5.5, you can recover the passwords by simply removing the *config.txt* file and then rebooting.

**Note**

This procedure resets the configuration of the SCE platform to factory defaults. Therefore, this procedure should be performed only after making sure that no traffic can be affected by the behavior of the SCE platform.

**Note**

This procedure resets the configuration of the SCE platform to factory defaults. Therefore, all current user configuration is lost. To recover the passwords without losing the user configuration, use the procedure described in [How to Recover the Passwords: Saving the Current Configuration, page 5-52](#)

-
- Step 1** Connect a serial terminal to the ‘AUX’ port at 9600 baud.
- Step 2** Press Enter so that the prompt appears.
Connects to the SCE platform.
- Step 3** Type `rm "/tffs0/system/config.txt"` and press Enter.
Deletes the configuration file, which contains the unknown passwords.
- Step 4** Type `reboot` and press Enter.
Reboot the system to restore the default configuration, including default passwords.

**Note**

To block unauthorized users from connecting to the SCE platform using the default password, a new password should be configured immediately for all levels for which such a password is required. The configuration should be saved (use the CLI command **copy running-config startup-config**) to make the new passwords permanent.

How to Recover the Passwords: Saving the Current Configuration

Simply deleting the *config.txt* file, while quick and easy to do, resets the configuration of the SCE platform to factory defaults, replacing all current user configuration. Use this procedure, which saves the configuration file, to recover the passwords without losing the current configuration.

**Note**

Although this procedure does save the current configuration, the process involves a temporary reset of the configuration of the SCE platform to factory defaults. Therefore, this procedure should be performed only after making sure that no traffic can be affected by the behavior of the SCE platform.

-
- Step 1** Connect a serial terminal to the ‘AUX’ port at 9600 baud.
- Step 2** Press Enter so that the prompt appears.
Connects to the SCE platform.
- Step 3** Type `cd system` and press Enter.
Navigates to the *system* directory.

- Step 4** Type `rename "config.txt" "config2.txt"` and press Enter.
Renames the configuration file so that it will not be lost when the system is rebooted.
- Step 5** Type `reboot` and press Enter.
Reboots the system to reset passwords. The default configuration is restored, in which all passwords are *Cisco*. This will provide access to the system.
(The IP address configuration remains as it was last configured.)
- Step 6** Establish a Telnet session to the SCE platform and enable the **admin** authorization level.
Use the default password (Cisco).
- Step 7** At the SCE# prompt type `copy /system/config2.txt ftp://<user>:<ftp_password>@ip_address/<path>/config2.txt` and press Enter.
Copies the file `/system/config2.txt` to the workstation using the SCE platform FTP client.
- Step 8** On the workstation, open the file. Look for the lines that start with **enable password**.
If passwords are not encrypted, you will be able to see the passwords, and can simply take note of them.
If passwords are encrypted, do the following
- a. Edit the file by removing the lines that begin with **enable password**.
 - a. Save the file.
 - a. At the SCE# prompt, type `copy ftp://<user>:<ftp_password>@ip_address/<path>/config2.txt /system/config2.txt` and press Enter.
Copies the file from the workstation back to the SCE platform disk space using the SCE platform FTP client.
- Step 9** At the SCE# prompt, type `rename /system/config2.txt /system/config.txt` and press Enter.
Renames the configuration file on the SCE platform back to the original name `config.txt`.
- Step 10** At the SCE# prompt, type `reload` and press Enter.
Reboots the SCE platform to restore the saved user configuration
- If passwords were not encrypted—the user-configured passwords that you viewed in the copied file are restored, since the configuration file was not changed.
 - If passwords were encrypted—the default password **Cisco** remains, since the encrypted lines were removed from the configuration file before it was copied back to the SCE platform.
- To block unauthorized users from connecting to the SCE platform using the default password, a new password should be configured immediately for all levels for which such a password is required. The configuration should be saved (use the CLI command **copy running-config startup-config**) to make the new passwords permanent.
-

How to Recover the Passwords: SCOS versions 2.5.5 or later

In SCOS versions 2.5.5 or later, a specific command is available to restore the default passwords. However, it is important to note that this default password configuration is only temporary. New passwords should be configured and saved immediately both for security and also so that the unknown passwords will not be restored in case of system reboot.



Note

This procedure does not affect configuration parameters other than login passwords. It is, therefore, safe to execute during traffic control.

Step 1 Connect a serial terminal to the 'AUX' port at 9600 baud.

Step 2 Press Enter so that the prompt appears.

Connects to the SCE platform.

Step 3 At the prompt, type `PSWD_ResetAll` and press Enter.

Resets the **enable** passwords.

The following message will appear:

```
All 'enable' passwords have been reset.
```

The SCOS is now using the default passwords for all levels. Note that this is a temporary state that is not preserved after a reboot. Rebooting the SCE platform without changing and saving the passwords will restore the unknown passwords.

In order to block unauthorized users from connecting to the SCE platform using the default password, a new password should be configured immediately for all levels for which such a password is required. The configuration should be saved (use the CLI command **copy running-config startup-config**) to make the new passwords permanent.

IP Configuration

- [How to Configure the IP Routing Table, page 5-54](#)
- [IP Advertising, page 5-57](#)
- [How to Configure the IP Address of the Management Interface, page 5-59](#)

How to Configure the IP Routing Table

- [About the IP Routing Table, page 5-55](#)
- [How to Configure the Default Gateway, page 5-55](#)
- [How to Add an Entry to the IP Routing Table, page 5-55](#)
- [How to Display the IP Routing Table, page 5-56](#)

About the IP Routing Table

For handling IP packets on the out-of-band MNG port, the 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 SCE platform cannot determine where to route a packet, it sends the packet to the default gateway.

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 10 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, page 5-55](#)
- [Configuring the Default Gateway: Example, page 5-55](#)

Options

The following option is available:

- `ip-address`—the IP address of the default gateway.

Step 1 From the SCE(config)# prompt, type `ip default-gateway ip-address`, and press Enter.

Enables privileged EXEC mode.

- Enter your password if prompted.
-

Configuring the Default Gateway: Example

The following example shows how to set the default gateway IP of the 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, page 5-56](#)
- [How to Add an Entry to the IP Routing Table: Example, page 5-56](#)

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.

-
- Step 1** From the SCE(config)# prompt, type `ip route prefixmasknext-hop`, and press Enter.
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 5-56](#)
- [How to Display the IP Routing Table for a Specified Subnet, page 5-56](#)

How to Display the Entire IP Routing Table

- [Displaying the Entire IP Routing Table: Example, page 5-56](#)

-
- Step 1** From the SCE# prompt, type `show ip route` and press Enter.
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, page 5-57](#)
- [Displaying the IP Routing Table for a Specified Subnet: Example, page 5-57](#)

Options

The following options are available:

- **prefix** — IP address of the routing entry, in dotted notation.
- **mask** — The relevant subnet mask, in dotted notation

-
- Step 1** From the SCE# prompt, type `show ip route prefixmask` and press Enter.
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

- [About IP Advertising, page 5-57](#)
- [How to Configure IP Advertising, page 5-58](#)
- [How to Display the Current IP Advertising Configuration, page 5-59](#)

About IP Advertising

IP advertising is the act of periodically sending ping requests to a configured address at configured intervals. This maintains the 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

How to Configure 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, page 5-58](#)
- [How to Enable IP Advertising, page 5-58](#)
- [How to Configure the IP Advertising Destination, page 5-58](#)
- [How to Configure the IP Advertising Interval, page 5-58](#)
- [Configuring IP Advertising: Example, page 5-58](#)

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

-
- Step 1** From the SCE(config)# prompt, type `ip advertising` , and press Enter.
Enables IP advertising.
-

How to Configure the IP Advertising Destination

-
- Step 1** From the SCE(config)# prompt, type `ip advertising destination destination`, and press Enter.
Configures the destination for the IP advertising pings.
-

How to Configure the IP Advertising Interval

-
- Step 1** From the SCE(config)# prompt, type `ip advertising interval interval` , and press Enter.
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 interval 240
```

How to Display the Current IP Advertising Configuration

-
- Step 1** From the SCE# prompt, type `show ip advertising` and press Enter.
- Displays the status of IP advertising (enabled or disabled), the configured destination, and the configured interval
-

How to Configure the IP Address of the Management Interface

- [About the Management Interface IP Address, page 5-59](#)
- [Options, page 5-59](#)
- [Configuring the IP Address of the Management Interface: Example, page 5-60](#)

About the Management Interface IP Address

The user 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.

**Note**

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 SCE platform (see [How to Reboot and Shut Down the SCE Platform, page 3-18](#)) so that the change will take effect properly in all internal and external components of the SCE platform.

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

-
- Step 1** Connect the SCE platform directly to a local console.
- Establishes contact with the SCE platform that is not dependent on the configured IP address.

- Step 2** From the SCE(config if)# prompt, type `ip address ip-addresssubnet-mask` and press Enter.
- Configures a new IP address for the management interface
- 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..
-

Configuring the IP Address of the Management Interface: Example

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

```
SCE(config if)#ip address 10.1.1.1 255.255.0.0
```

How to Configure Time Clocks and Time Zone

- [About Time Clocks and Time Zone, page 5-60](#)
- [How to Display the System Time, page 5-61](#)
- [How to Display the Calendar Time, page 5-61](#)
- [How to Set the System Clock, page 5-61](#)
- [How to Set the Calendar, page 5-62](#)
- [How to Set the Time Zone, page 5-62](#)
- [How to Remove the Current Time Zone Setting, page 5-63](#)
- [How to Configure Daylight Saving Time, page 5-63](#)

About Time Clocks and Time Zone

The 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 SCE platform does not track Daylight Saving Time automatically, so you must update the time zone when the time changes bi-annually.

The 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 SCE platform is not powered up. When the 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.

How to Display the System Time

Step 1 From the SCE(config)# prompt, type `show clock` and press Enter.

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#
```

How to Display the Calendar Time

Step 1 From the SCE(config)# prompt, type `show calendar` and press Enter.

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#
```

How to Set the System Clock

- [Options, page 5-61](#)
- [Setting the System Clock: Example, page 5-62](#)

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 `clock set time-date`, and press Enter.
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
```

How to Set the Calendar

The calendar is a system clock that continues functioning even when the system shuts down.

- [Options, page 5-62](#)
- [Setting the Calendar: Example, page 5-62](#)

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 2007
SCE#clock read-calendar
SCE#show calendar
10:21:06 UTC THU May 13 2007
```

How to Set the Time Zone

- [Options, page 5-63](#)
- [Setting the Time Zone: Example, page 5-63](#)

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

-
- Step 1** From the SCE(config)# prompt, type `clock timezone zonehoursminutes`, and press Enter.
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)#
```

How to Remove the Current Time Zone Setting

-
- Step 1** From the SCE(config)# prompt, type `no clock timezone` and press Enter.
Removes the timezone configuration and resets the timezone to the default value (UTC).
-

How to Configure Daylight Saving Time

The 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 5-64](#)
- [Guidelines, page 5-65](#)
- [How to Define Recurring Daylight Saving Time Transitions, page 5-65](#)
- [How to Define Non-Recurring Daylight Saving Time Transitions, page 5-66](#)
- [How to Cancel the Daylight Saving Time Configuration, page 5-66](#)
- [How to Display the Current Daylight Saving Time Configuration, page 5-66](#)

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

Step 1 From the SCE(config)# prompt, type `clock summer-time zonerecurring [week1 day1 month1 time1 week2 day2 month2 time2 [offset]]` and press Enter.

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

-
- Step 1** From the SCE(config)# prompt, type `clock summer-time zone [date1 month1 year1 time1 date2 month2 year2 time2[offset]]` and press Enter.
- Enables privileged EXEC mode.
-

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

-
- Step 1** From the SCE(config)# prompt, type `no clock summer-time` and press Enter.
- Removes all daylight saving configuration.
-

How to Display the Current Daylight Saving Time Configuration

-
- Step 1** From the SCE# prompt, type `show timezone` and press Enter.
- Displays the current time zone and daylight saving time configuration.
-

How to Configure SNTP

- [About SNTP, page 5-67](#)
- [How to Enable the SNTP Multicast Client, page 5-67](#)
- [How to Disable the SNTP Multicast Client, page 5-67](#)
- [How to Enable the SNTP Unicast Client, page 5-67](#)
- [How to Disable the SNTP Unicast Client, page 5-68](#)
- [How to Define the SNTP Unicast Update Interval, page 5-68](#)
- [How to Display SNTP Information, page 5-69](#)

About SNTP

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 [How to Configure Access Control Lists \(ACLs\)](#), page 5-27)

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

- Step 1** From the SCE(config)# prompt, type `sntp broadcast client` , and press Enter.
Enables the SNTP multicast client. It will accept time updates from any broadcast server.

How to Disable the SNTP Multicast Client

- Step 1** From the SCE(config)# prompt, type `no sntp broadcast client` , and press Enter.
Disables the SNTP multicast client. It will not accept any broadcast time updates.

How to Enable the SNTP Unicast Client

- [Options](#), page 5-68
- [Enabling SNTP Unicast Client: Example](#), page 5-68

Options

The following option is available:

- **ip-address** — the IP address of the SNMP unicast server.

-
- Step 1** From the SCE(config)# prompt, type `snmp server ip-address`, and press Enter
Defines the SNMP unicast server so that SNMP client is able to query that server.
-

Enabling SNMP Unicast Client: Example

The following example shows how to enable an SNMP server at IP address 128.182.58.100.

```
SCE(config)# snmp server 128.182.58.100
```

How to Disable the SNMP Unicast Client

- [How to Disable the SNMP Unicast Client and Remove All Servers, page 5-68](#)
- [How to Remove One SNMP Server, page 5-68](#)

How to Disable the SNMP Unicast Client and Remove All Servers

-
- Step 1** From the SCE(config)# prompt, type `no snmp server all` , and press Enter.
Removes all SNMP unicast servers, preventing unicast SNMP query.
-

How to Remove One SNMP Server

- [Options, page 5-68](#)

Options

The following option is available:

- **ip-address** — the IP address of the SNMP unicast server.

-
- Step 1** From the SCE(config)# prompt, type `no snmp server ip-address`, and press Enter
Removes the specified SNMP unicast server .
-

How to Define the SNMP Unicast Update Interval

- [Options, page 5-69](#)
- [Defining the SNMP Unicast Update Interval: Example, page 5-69](#)

Options

The following option is available:

- **interval** — the time in seconds between updates (64 through 1024)
default interval = 900 seconds

-
- Step 1** From the SCE(config)# prompt, type `sntp update-interval interval`, and press Enter.
Enables privileged EXEC mode.
- Configures the SNTP unicast client to query the server at the defined intervals.
-

Defining the SNTP Unicast Update Interval: 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

-
- Step 1** From the SCE# prompt, type `show sntp`, and press Enter.
Displays the configuration of both the SNTP unicast client and the SNTP multicast client.
-

Displaying SNTP Information: 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
```

How to Configure Domain Name Server (DNS) Settings

- [About Domain Name Servers, page 5-70](#)
- [How to Configure DNS Lookup, page 5-70](#)
- [How to Configure Name Servers, page 5-71](#)
- [How to Add a Host to the Host Table, page 5-72](#)
- [How to Display Current DNS Settings, page 5-72](#)

About Domain Name Servers

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

How to Configure DNS Lookup

- [How to Enable DNS Lookup, page 5-70](#)
- [How to Disable DNS Lookup, page 5-70](#)

How to Enable DNS Lookup

-
- Step 1** From the SCE(config)# prompt, type `ip domain-lookup` and press Enter.
Enables DNS lookup.
-

How to Disable DNS Lookup

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

How to Configure Name Servers

- [Options, page 5-71](#)
- [How to Define Domain Name Servers, page 5-71](#)
- [How to Remove a Domain Name Server, page 5-71](#)
- [How to Remove All Domain Name Servers, page 5-72](#)

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.

- [Defining Domain Name Servers: Example, page 5-71](#)

Step 1 From the SCE(config)# prompt, type `ip name-server server-address1 [server-address2 [server-address3]]`, and press Enter.

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

Step 1 From the SCE(config)# prompt, type `no ip name-server server-address1 [server-address2 [server-address3]]`, and press Enter.

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

-
- Step 1** From the SCE(config)# prompt, type `no ip name-server` , and press Enter.
Removes all configured DNS servers.
-

How to Add a Host to the Host Table

- [Options, page 5-72](#)
- [Adding Hosts to Removing them from the Host Table: Examples, page 5-72](#)

Options

The following options are available:

- **hostname** — The name of the host.
- **ip-address** — The IP address of the host

-
- Step 1** From the SCE(config)# prompt, type `ip host hostname ip-address`, and press Enter.
Adds the specified host to the host table.
-

Adding Hosts to Removing them from the Host Table: Examples

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

-
- Step 1** From the SCE# prompt, type `show hosts`, and press Enter.
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#
```

How to Configure the Management Port Physical Parameters

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

The procedures for configuring this interface are explained in the following sections.

- [How to Configure the Management Interface Speed and Duplex Parameters, page 5-73](#)
- [How to Monitor the Management Interface, page 5-75](#)

How to Configure 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.

- [How to Configure the Duplex Operation of the Management Interface, page 5-73](#)
- [How to Configure the Speed of the Management Interface, page 5-74](#)

How to Configure the Duplex Operation of the Management Interface

- [Options, page 5-73](#)
- [Configuring the Duplex Operation of the Management Interface: Example, page 5-74](#)

Options

The following options are available:

- **duplex** — duplex operation of the currently selected 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 (see).

-
- Step 1** From the SCE(config if)# prompt, type `duplex auto|full|half` and press Enter.
Configures the duplex operation of the currently selected management interface.
-

Configuring the Duplex Operation of the Management Interface: Example

The following example shows how to use this command to configure both management ports to half duplex mode.

```
SCE#config
SCE(config)#interface mng 0/1
SCE(config if)#duplex half
SCE(config if)#exit
SCE(config)#interface mng 0/2
SCE(config if)#duplex half
```

How to Configure the Speed of the Management Interface

- [Options, page 5-74](#)
- [Configuring the Speed of the Management Interface: Example, page 5-74](#)

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 (see).

-
- Step 1** From the SCE(config if)# prompt, type `speed 10|100|auto` and press Enter.
Configures the speed of the currently selected management interface.
-

Configuring the Speed of the Management Interface: Example

The following example shows how to use this command to configure both management ports to The following example shows how to use this command to configure both management ports to half duplex mode..

```
SCE#config
SCE(config)#interface mng 0/1
SCE(config if)#speed 100
SCE(config if)#exit
SCE(config)#interface mng 0/2
SCE(config if)#speed 100
```

How to Monitor the Management Interface

Use this command to display the following information for the specified management interface. Speed and duplex parameters are specific to the selected interface (port), while other parameters apply to both ports and are displayed by a command to either interface.

- speed
- duplex
- IP address
- active port

Step 1 From the SCE# prompt, type `show interface Mng {01 | 0/2} [auto-fail-over|duplex|ip address|speed]` and press Enter.

Displays the specified management interface configuration.

- If no option is specified, all management interface information is displayed.
 - The displayed speed and duplex mode are for the management interface specified in the command.
-



CHAPTER 6

Configuring the Line Interface

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

- [How to Configure the Line Interfaces, page 6-1](#)
- [How to Configure Tunneling Protocols, page 6-3](#)
- [How to Configure VLAN Translation, page 6-9](#)
- [How to Configure Traffic Rules and Counters, page 6-12](#)
- [TOS Marking, page 6-20](#)
- [How to Count Dropped Packets, page 6-20](#)

How to Configure the Line Interfaces

- [Line Interfaces, page 6-1](#)
- [How to Configure the Gigabit Ethernet Line Interfaces, page 6-2](#)
- [How to Configure the Fast Ethernet Line Interfaces, page 6-2](#)

Line Interfaces

The Line Interfaces (Subscriber and Network) are used to connect the SCE platform to the network. See the description of network topologies in the *Topology* section of the *Cisco SCE 2000/SCE 1000 Installation and Configuration Guides*.

The SCE 1000 2xGBE and the SCE 2000 4xGBE have Gigabit Ethernet line interfaces. You should configure **auto-negotiate** for these interfaces.

The SCE 2000 4/8x FE has Fast Ethernet line interfaces. You should configure the **speed** and **duplex** for these interfaces.

How to Configure the Gigabit Ethernet Line Interfaces



Note The maximum packet size supported by the SCE platform is 1600 bytes

- Step 1** At the SCE# prompt, type `configure` , and press Enter.
Enters Global Configuration mode.
- Step 2** At the SCE(config)# prompt, type `interface GigabitEthernet 0/ portnumber`, and press Enter, where *portnumber* is the number of the selected port (1-4).
Enters Interface Configuration mode for the selected GBE interface.
- Step 3** At the SCE(config if)# prompt, type `auto-negotiate`, and press Enter.
Enables auto-negotiation.
- Step 4** At the SCE(config if)# prompt, type `exit` and press Enter.
Exits to global configuration mode, from which you can access a different Gigabit Ethernet interface.
- Step 5** Repeat steps 2 and 3 for the remaining Gigabit Ethernet interfaces.
Auto-negotiation must enabled separately and explicitly for each Gigabit Ethernet interface.
-

How to Configure the Fast Ethernet Line Interfaces

Note that both sides of the FE link (both the SCE 2000 4/8xFE and the remote device) should have the same configuration. Use either of the following two configuration options:

- Autonegotiation = ON
- Autonegotiation = ON, speed = 100

-
- Step 1** At the SCE# prompt, type `configure` , and press Enter.
Enters Global Configuration mode.
- Step 2** At the SCE(config)# prompt, type `interface FastEthernet 0/ portnumber`, and press Enter, where *portnumber* is the number of the selected port (1-4).
Enters Interface Configuration mode for the selected FE interface.
- Step 3** At the SCE(config if)# prompt, type `duplex auto/full/half`, and press Enter.
Configures the duplex mode for the selected FE interface.
- Step 4** At the SCE(config if)# prompt, type `speed auto/100`, and press Enter.
Configures the speed for the selected FE interface.
- Step 5** At the SCE(config if)# prompt, type `exit` and press Enter.
Exits to global configuration mode, from which you can access a different Fast Ethernet interface.
- Step 6** Repeat steps 2 through 5 for the remaining Fast Ethernet interfaces.
Speed and duplex mode must enabled separately and explicitly for each Fast Ethernet interface.
-

How to Configure Tunneling Protocols

- [How to Select the Tunneling Mode, page 6-5](#)
- [How to Display the Tunneling Configuration, page 6-8](#)
- [How to Display the Logged-in VPNs, page 6-9](#)
- [About Tunneling Protocols, page 6-3](#)
- [VPNs, page 6-4](#)

About Tunneling Protocols

Tunneling technology is used across various telecommunications segments to solve a wide variety of networking problems. The SCE platform is designed to recognize and process various tunneling protocols in several ways. The SCE platform is able to either ignore the tunneling protocols ("skip" the header) or treat the tunneling information as subscriber information ("classify"). A special case of classification by tunneling information is MPLS/VPN with private IP support

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

Table 6-1 Tunneling Protocol Summary

| Protocol | Supported handling | Mode name |
|----------|---|--------------------------------------|
| L2TP | Ignore tunnel | IP-tunnel L2TP skip |
| | Don't ignore tunnel – classify by external IP | No IP-Tunnel |
| VLAN | Ignore tunnel | VLAN symmetric skip |
| | Ignore tunnel – asymmetric | VLAN a-symmetric skip |
| | VLAN tag used for VPN classification | VLAN symmetric classify |
| MPLS | Ignore tunnel (inject unlabeled) | MPLS Traffic-engineering skip |
| | Ignore tunnel (inject labeled) | MPLS VPN skip |
| | MPLS labels used for VPN classification | MPLS VPN auto-learn |

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

L2TP

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

VLAN

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

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

MPLS

MPLS/VPN-based subscribers with private IP support are a special case of classification by tunneling information.

MPLS labels are supported up to a maximum of 15 labels per packet.

For more information regarding MPLS/VPN-based subscribers, see the [MPLS/VPN Support, page 13-1](#)

VPNs

- [Private IP Addresses, page 6-4](#)
- [Capacity, page 6-4](#)
- [Limitations for VPN mode, page 6-4](#)

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

A VPN may contain several MPLS mappings or 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.

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

Private IP Addresses

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

- MPLS VPN auto-learn
- VLAN symmetric classify

Capacity

The system supports:

- 2015 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:

- TCP Bypass-establishment
- DDoS
- Value Added Services (VAS) mode

Number of MPLS labels

- The choice of the unique VPN site must be based on the BGP label only. The BGP label must be the innermost label.
- The MPLS/VPN solution supports various combinations of labels.
- The system does not support VPNs for which other MPLS-related features, such as MPLS-TE or MPLS-FRR, are enabled

Subscriber-related limitations

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

Topology-related limitations

- An asymmetrical routing topology in which the traffic may be unidirectional, is not supported, since the identification of the VPN relies on the bidirectional nature of the traffic for various mechanisms.

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

- Two VPN sites must be aggregated into one VPN if the following conditions are both true:
 - They are both connected to the same SCE platform
 - They both communicate with a common remote site using the same upstream labels and P router.
- In MPLS-based VPNs (MPLS auto-learn mode), a subscriber MAY NOT have IP mappings over more than one VPN
- 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.)

How to Select the Tunneling Mode

Use these commands to configure tunneling:

- ip tunnel
- vlan
- mpls
- L2TP identify-by

How to Configure IP Tunnels

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

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

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

Enables IP tunnel mode

- To disable IP tunnels, use the following command:
no ip tunnel
-

How to Configure the VLAN Environment

Options

There are three options:

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

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

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

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

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



Note

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

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

Specify the desired VLAN mode.

Configuring the VLAN Environment: Example

The following example selects VLAN-based classification.

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

How to Configure the MPLS Environment

Use this command to set the MPLS environment.

- [Options, page 6-7](#)
- [Configuring the MPLS Environment: Example, page 6-7](#)

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.
- **VPN auto-learn** —Use in an MPLS/VPN environment where auto-learning is required due to the existence of private IP addresses and/or VPN based subscribers.

When this option is configured, both **ip-tunnel** and **VLAN** must be set to their default values.

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.

The *MPLS VPN auto-learn* option is required in an MPLS/VPN environment.

Step 1 From the SCE(config if)# prompt, type **mpls {traffic-engineering skip|vpn skip|vpn auto-learn}** and press Enter.

Specify the desired MPLS mode.

Configuring the MPLS Environment: Example

The following example selects the MPLS/VPN tunnel environment.

```
SCE(config if)#mpls vpn auto-learn
```

About Changing VPN Modes

VPNs can only exist in either VLAN symmetric classify or MPLS VPN auto-learn, but these two modes cannot be enabled simultaneously. When changing from one of these VPN-related modes to another, keep the following guidelines in mind:

- All VPN-based subscribers must be cleared to change the tunneling mode. If the connection with the SM is down, use the `no subscriber all with-vpn-mappings` CLI command (see [How to Remove VPN-based Subscribers, page 9-11](#))
- All VPN mappings must also be removed. This can only be done via the SM CLU (which means that the connection with the SM must be up). (See [How to Manage VPN Mappings, page 13-26](#)

How to Restore the Default VLAN or MPLS Environment

Use this command to restore the default VLAN or MPLS configuration.

It is not usually necessary to explicitly restore the default environment, as this is done automatically when executing a VLAN or MPLS command. When such an automatic reset to default occurs, a warning message appears similar to the following:

Warning - Disabled previously configured IP-tunnel support or tunneling classification mode

-
- Step 1** From the SCE(config if)# prompt, type `default {mpls | vlan}` and press Enter.
-

How to Configure the L2TP Environment

Use this command to set the port number that the LNS and LAC use for L2TP tunnels.

Options

The following option is available:

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

Default port# = 1701

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

**Note**

If external fragmentation exists in the L2TP environment, it is required to configure a Traffic Rule (see [How to Configure Traffic Rules and Counters, page 6-12](#)) 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.

How to Display the Tunneling Configuration

-
- Step 1** From the SCE# prompt, type `show interface linecard 0 MPLS|VLAN|L2TP|IP-tunnel` and press Enter.

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.

Step 1 From the SCE>prompt, type `show interface linecard 0 VPN {name vpn-name | all-names}` and press Enter.

How to Configure VLAN Translation

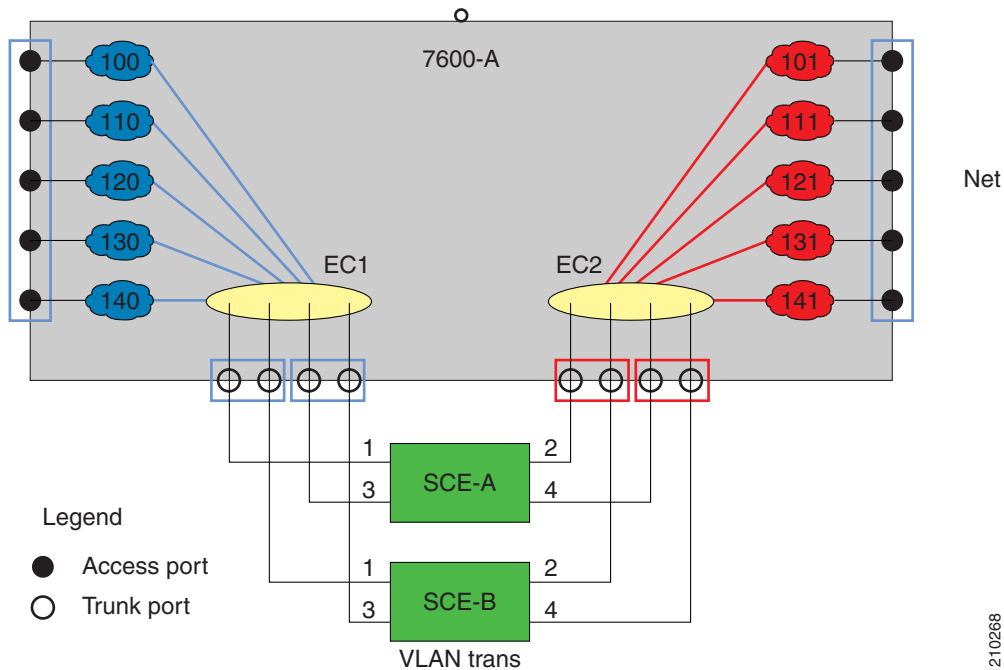
- [About VLAN Translation, page 6-10](#)
- [VLAN Translation Features and Limitations, page 6-10](#)
- [How to Set the VLAN Translation Constant, page 6-11](#)
- [How to Disable VLAN Translation, page 6-12](#)
- [How to Monitor VLAN Translation, page 6-12](#)

About VLAN Translation

Some topologies require the SCE platform to be able to translate between different VLAN tags.

The following drawing illustrates an example of such a system, in which one router acts as a dispatcher, forwarding traffic and performing load balancing between two SCE 2000 platforms.

Figure 6-1 VLAN Translation



In this example, traffic enters the router via the access ports; it is forwarded to an EtherChannel, which is configured as a trunk, and enters the SCE 2000 platforms.

As can be seen from this drawing, the subscriber side VLAN tags must be different from those on the network side, or the router will simply forward the traffic to the opposite port. This can be supported very simply by having the SCE platform replace the VLAN tags according to a preset configuration.

VLAN Translation Features and Limitations

Features

- Configuration of an increment or decrement constant.
- Configuration of the constant is global for the line card.
- The configured operation (either increment or decrement) is applied to the network side.
- The subscriber side automatically performs the opposite operation. That is, if the VLAN is incremented by X on the network side, it is decremented by X on the subscriber side.
- VLAN tagged packets are changed (incremented or decremented) before transmission.
- Non-tagged packets are not changed.
- This feature allows seamless processing with non-VLAN traffic.

Limitations

- LIC Bypass not supported – Translation is done in the transmission. Therefore, in LIC bypass, where there is no transmission, there is also no translation.
This means that in general, installations using the VLAN translation feature should rely on cutoff on failure and at upgrade (use redundant SCE platform).
- STP hazard – VLAN translation may interfere with Spanning Tree Protocol. This should be taken in consideration when deploying the solution.
- The maximum offset that can be configured is 2047. Note that there is no protection for wraparound.

How to Set the VLAN Translation Constant

Use this command to define the VLAN translation constant. Make sure that the same VLAN translation constant is configured for all SCE platforms in the system.

- [Options, page 6-11](#)
- [Setting the VLAN Translation Constant: Example, page 6-11](#)

Options

The following options are available:

- **increment | decrement** —Keywords indicating whether to increment or decrement the VLAN by the specified *value*.
- **value** — Integer value by which the VLAN is to incremented or decremented.

The configured translation is applied to the network port side. The reverse operation is performed at the subscriber side.

For example, if "increment 5" is defined, at the network port the VLAN is incremented by 5, and at the subscriber port the VLAN is decremented by 5.

In this case, the network side VLAN tags might be 105, 205, 305, and the subscriber side the VLAN tags would then be 100, 200, 300.

Default = 0

Maximum = 2047 (Note that there is no protection for wraparound of the VLAN value.)

Step 1 From the SCE(config if)# prompt, type `vlan translation increment|decrement value value` and press Enter.

Sets the VLAN translation constant.

Setting the VLAN Translation Constant: Example

The following example sets the translation constant to 10, decremented at the network side.

```
SCE(config if)#vlan translation decrement value 10
```

How to Disable VLAN Translation

-
- Step 1** From the SCE# prompt, type `no vlan translation` and press Enter
Disables VLAN translation.
-

How to Monitor VLAN Translation

-
- Step 1** From the SCE# prompt, type `show interface linecard 0 vlan translation` and press Enter.
Displays current VLAN translation configuration.
-

How to Configure Traffic Rules and Counters

- [Information About Traffic Rules and Counters, page 6-12](#)
- [How to Configure Traffic Counters, page 6-14](#)
- [How to Configure Traffic Rules, page 6-15](#)
- [How to Manage Traffic Rules and Counters, page 6-18](#)

Information About Traffic Rules and Counters

- [What are Traffic Rules and Counters?, page 6-12](#)
- [Traffic Rules, page 6-13](#)
- [Traffic Counters, page 6-14](#)

What are Traffic Rules and Counters?

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

Possible uses for traffic rules and counters include:

- Enabling the user to count packets according to various criteria. Since the traffic counters are readable via the SCE SNMP MIB, these might be used to monitor up to 32 types of packets, according to the requirements of the installation.
- Ignoring certain types of flows. When a traffic rules specifies an “ignore” action, packets matching the rule criteria will not open a new flow, but will pass through the SCE platform without being processed. This is useful when a particular type of traffic should be ignored by the SCE platform.

Possible examples include ignoring traffic from a certain IP range known to require no service, or traffic from a certain protocol.

- Blocking certain types of flows. When a traffic rules specifies a “block” action, packets matching the rule criteria (and not belonging to an existing flow) will be dropped and not passed to the other interface. This is useful when a particular type of traffic should be blocked by the SCE platform.

Possible examples include performing ingress source address filtering (dropping packets originating from a subscriber port whose IP address does not belong to any defined subscriber-side subnet), or blocking specific ports.

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

Traffic Rules

A traffic rule specifies that a defined action should be taken on packets processed by the SCE Platform that meet certain criteria. The maximum number of rules is 128, which includes not only traffic rules configured via the SCE platform CLI, but also any additional rules configured by external management systems, such as SCA BB. Each rule is given a name when it is defined, which is then used when referring to the rule.

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

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

The possible actions are:

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

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

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

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

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

Traffic Counters

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

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

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

How to Configure Traffic Counters

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

- [How to Create a Traffic Counter, page 6-14](#)
- [How to Delete a Traffic Counter, page 6-15](#)
- [How to Delete all Existing Traffic Counters, page 6-15](#)

How to Create a Traffic Counter

Options

The following options are available:

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

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

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

How to Delete a Traffic Counter

-
- Step 1** From the SCE(config if)# prompt, type `no traffic-counter name name` and press Enter.
Note that a traffic counter cannot be deleted if it is used by any existing traffic rule.
-

How to Delete all Existing Traffic Counters

-
- Step 1** From the SCE(config if)# prompt, type `no traffic-counter all` and press Enter.
Removes all traffic counters.
Note that a traffic counter cannot be deleted if it is used by any existing traffic rule.
-

How to Configure Traffic Rules

Use the following commands to create and delete traffic rules.

- [How to Create a Traffic Rule, page 6-15](#)
- [How to Delete a Traffic Rule, page 6-18](#)
- [How to Delete all Traffic Rules, page 6-18](#)

How to Create a Traffic Rule

- [Options, page 6-15](#)
- [Configuring Traffic Rules: Examples, page 6-16](#)

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/Other
```

tunnel id specification:

all|([all-but] tunnel id)

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

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

Note 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

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

Enables privileged EXEC mode.

Configuring Traffic Rules: Examples

- [Example 1, page 6-17](#)
- [Example 2, page 6-17](#)
- [Example 3, page 6-17](#)

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 = other
- Direction = both
- Traffic counter = counter1
- The only action performed will be counting

```
SCE(config if)# traffic-rule rule1 IP-addresses subscriber-side all network-side
10.10.10.10 protocol other 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
- 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 rule2 IP-addresses subscriber-side all network-side all-but
10.10.10.0/24 protocol tcp 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 rule3 IP-addresses all protocol IS-IS direction upstream
traffic-counter none action ignore
```

How to Delete a Traffic Rule

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

How to Delete all Traffic Rules

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

How to Manage Traffic Rules and Counters

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

You can also reset a specific counter or all counters.

- [How to View a Specified Traffic Rule, page 6-18](#)
- [How to View all Traffic Rules, page 6-18](#)
- [How to View a Specified Traffic Counter, page 6-19](#)
- [How to View all Traffic Counters, page 6-19](#)
- [How to Reset a Specified Traffic Counter, page 6-19](#)
- [How to Reset all Traffic Counters, page 6-19](#)

How to View a Specified Traffic Rule

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

How to View all Traffic Rules

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

How to View a Specified Traffic Counter

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

Viewing a Traffic Counter: Example

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

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

How to View all Traffic Counters

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

Viewing the Traffic Counters: Example

The following example displays information for all existing traffic counters.

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

How to Reset a Specified Traffic Counter

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

How to Reset all Traffic Counters

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

TOS Marking

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

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

For information on configuring TOS marking, please refer to the *Cisco Service Control Application for Broadband User Guide, Rel 3.1.5*.

**Note**

TOS marking in release 3.1.5 is not backwards compatible with any previous SCOS releases.

How to Display the TOS Marking Configuration

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

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

How to Count Dropped Packets

- [About Counting Dropped Packets, page 6-21](#)
- [How to Disable the Hardware Packet Drop, page 6-21](#)

About Counting Dropped Packets

By default, the SCE platform hardware drops red 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-red-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-red-packets-by-hardware mode is enabled.

How to Disable the Hardware Packet Drop

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

By default hardware packet drop is enabled.

**Note**

Disabling this feature may have both delay and performance implications.


Step 1

From the SCE(config if)# prompt, type `no accelerate-packet-drops` and press Enter.

Disables hardware packet drop.

- To enable hardware packet drop, use the following command:

`accelerate-packet-drops`



CHAPTER 7

Configuring the Connection

- [How to Configure the Connection Mode, page 7-1](#)
- [How to Monitor the Connection Mode, page 7-3](#)
- [How to Configure the Link Mode, page 7-4](#)
- [How to Configure Asymmetric Routing Topology, page 7-5](#)
- [How to Configure a Forced Failure, page 7-7](#)
- [How to Configure the Failure Recovery Mode, page 7-8](#)
- [How to Configure the SCE Platform/SM Connection, page 7-9](#)
- [How to Enable and Disable Link Failure Reflection, page 7-10](#)

How to Configure 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 SCE platform.

- [Options, page 7-2](#)
- [Configuring the Connection Mode: Examples, page 7-3](#)

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 SCE platform:
 - Inline — single SCE platform inline
 - Receive-only — single SCE platform receive-only
 - Inline-cascade — two cascaded SCE platforms inline
 - Receive-only-cascade — two cascaded SCE platforms receive-only
- Default — **inline**
- **Physically-connected-links** — In cascaded topologies, defines which link is connected to this SCE platform.
Possible values are 'link-0' and 'link-1'.
Not applicable to single SCE platform topologies.
- **Priority** — This parameter defines which is the primary SCE platform. It is applicable only in a two SCE platform topology.
Possible values are 'primary' and 'secondary'
Not applicable to single SCE platform topologies.
- **On-failure** — This parameter determines whether the system cuts the traffic (cutoff) or bypasses it (bypass) when the SCE platform either has failed or is booting.
Default — **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 `physically-connected-links` [**link 0|link 1**] `priority` [**primary|secondary**] `on-failure` [**bypass|cutoff**] and press Enter.
Configures the connection mode.

Configuring the Connection Mode: Examples

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

Example 1

This example defines the primary device in a two-SCE platform redundant, inline topology. Link 0 is connected to this device, and the link mode on failure is bypass.

```
SCE(config if)# connection-mode inline-cascade physically-connected-links link-0 priority
primary on-failure bypass
```

Example 2

This example defines a single-SCE platform, dual link, receive-only topology. Neither link mode on failure, nor physically-connected-links, nor priority is applicable.

```
SCE(config if)# connection-mode receive-only
```

How to Monitor the Connection Mode

-
- Step 1** From the SCE>prompt, type `show interface linecard 0 connection-mode` and press Enter. Displays the connection mode configuration.
-

Monitoring the Connection Mode: Example

The following example shows how to display the current configuration of the connection mode.

```
SCE>show interface linecard 0 connection-mode
Connection mode is inline
slot failure mode is bypass
Redundancy status is standalone
SCE>
```

How to Configure the Link Mode

- [About the Link Mode, page 7-4](#)
- [Options, page 7-4](#)

About the Link Mode

The SCE platform has an internal hardware card used to maintain the links even when the SCE platform fails. This hardware card has four possible modes of operation:

- bypass
- forwarding
- cutoff
- sniffing

Normally, the link mode is selected by the 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 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 on the specified link to the SCE platform for processing.
- **Bypass** — stops all forwarding of traffic on the specified link to the SCE platform. Traffic still flows on the link, but is not processed in any way by the SCE platform.

This does not affect the redundancy states.

- **Sniffing** — allows the SCE platform to forward traffic on the specified link through the bypass mechanism while still analyzing the traffic passively.

Sniffing is permitted to be configured for all links, only (use the **all-links** option).

- **Cutoff** — completely cuts off flow of traffic through the specified link.
- **link#** — the number of the link being configured

Use the **all-links** option to apply the configuration to all links

Recommendations and restrictions

Note the following recommendations and restrictions:

- Since the SCE 1000 platform has only one link, the link is not specified.
- Since the SCE 2000 platforms have more than one link, it is required to specify the link. The link designations are different for the GBE and FE platforms, as follows:
 - SCE 2000 4xGBE — GBE1-GBE2/GBE3-GBE4
 - SCE 2000 4/8xFE — LINK1/LINK2
- Use the **all-links** option to configure the link mode for all links (SCE 2000 platforms only).
- It is recommended that both links be configured together. Use the **all-links** option.
- Link mode is relevant only to inline topologies.
- It is recommended that in cascaded topologies, both SCE platforms be configured for the same link mode, otherwise the service will be unpredictable.
- Sniffing can only be configured for all links, therefore, to configure sniffing, the all-links option is required, not just recommended.
- 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.

Step 1 From the SCE(config if)# prompt, type `link-mode [linknumber | all-links] [forwarding | bypass | sniffing | cutoff]` and press Enter.

Configures the link mode for the specified link.

How to Configure Asymmetric Routing Topology

- [About Asymmetric Routing Topology, page 7-6](#)
- [Asymmetric Routing and Other Service Control Capabilities, page 7-6](#)
- [Enabling Asymmetric Routing, page 7-6](#)
- [How to Monitor Asymmetric Routing, page 7-7](#)

About Asymmetric Routing Topology

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 SCE platforms, resulting in each 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 SCE platforms in cascade mode (or through an MGSCP cluster), thereby making sure that both directions of a flow run through the same SCE platform. However, this is sometimes not feasible, due to the fact that the 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 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 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, please see the *Cisco Service Control Application for Broadband User Guide* .

How to Monitor 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 SCE platform was last reloaded (or since the counters were last reset).

Step 1 From the SCE>prompt, type `show interface linecard 0 asymmetric-routing-topology` and press Enter.

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   :   2%
Traffic Processor 2   :   7%
Traffic Processor 3   :   0%
The statistics are updated once every two minutes
SCE>
```

How to Configure 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. (See [How to Manage Application Files](#), page 3-10.)

- [How to Force a Virtual Failure](#), page 7-7
- [How to Exit from a Virtual Failure](#), page 7-8

How to Force a Virtual Failure

Step 1 From the SCE(config if)# prompt, type `force failure-condition` and press Enter.

Forces a virtual failure condition.

How to Exit from a Virtual Failure

-
- Step 1** From the SCE(config if)# prompt, type `no force failure-condition` and press Enter.
Exits from the virtual failure condition.
-

How to Configure the Failure Recovery Mode

The **failure-recovery operation-mode** command defines the behavior of the system after boot resulting from failure.

- [Options, page 7-8](#)
- [Configure the Failure Recovery Mode: Examples, page 7-8](#)

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**.

-
- Step 1** From the SCE(config)# prompt, type `failure-recovery operation-mode operational|non-operational` and press Enter.
Specify 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
```

How to Configure the SCE Platform/SM Connection

- [About the SCE Platform/SM Connection, page 7-9](#)
- [How to Configure the Behavior of the SCE Platform in Case of Failure of the SM, page 7-9](#)
- [How to Configure the SM-SCE Platform Connection Timeout, page 7-10](#)

About the SCE Platform/SM Connection

The user can configure the behavior of the 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 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.

How to Configure the Behavior of the SCE Platform in Case of Failure of the SM

Options

The following options are available:

- **force-failure** — Force failure of SCE platform. The SCE platform then acts according to the behavior configured for the failure state.
- **remove-mappings** — Remove all current subscriber mappings.
- **shut** — The SCE platform shuts down and quits providing service.
- **none** (default) — Take no action.

Step 1 From the SCE(config if)# prompt, type `subscriber sm-connection-failure action [force-failure|none|remove-mappings|shut]` and press Enter.

Configures the action of the SCE platform in the case of failure of the connection to the SM.

How to Configure 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, page 7-10](#)

Options

The following option is available:

- **interval** — the timeout interval in seconds

-
- Step 1** From the SCE(config if)# prompt, type `subscriber sm-connection-failure action timeout interval` and press Enter.
- Configures the connection timeout.
-

How to Enable and Disable Link Failure Reflection

- [About Link Failure Reflection, page 7-10](#)
- [How to Enable Link Failure Reflection, page 7-10](#)
- [How to Disable Link Failure Reflection, page 7-11](#)
- [How to Enable and Disable Link Failure Reflection on All Ports, page 7-11](#)
- [How to Configure Link Failure Reflection in Linecard-Aware Mode \(SCE 2000 only\), page 7-12](#)

About Link Failure Reflection

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

-
- Step 1** From the SCE(config if)# prompt, type `link failure-reflection` and press Enter.
- Enables link failure-reflection.
-

How to Disable Link Failure Reflection

-
- Step 1** From the SCE(config if)# prompt, type `no link failure-reflection` and press Enter.
Disables link failure-reflection.
-

How to Enable and Disable Link Failure Reflection on All Ports

- [About Link Failure Reflection on All Ports, page 7-11](#)
- [Options, page 7-11](#)
- [How to Enable Link Failure Reflection on All Ports, page 7-12](#)
- [How to Disable Link Failure Reflection on All Ports, page 7-12](#)

About Link Failure Reflection on All Ports

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 SCE platform that the device is in a failure state, and therefore cannot be used.

**Note**

The **link reflection on all ports** feature cannot be used in a cascade mode, because in this mode one of the links is used to provide redundancy.

In **link reflection on all ports** mode, all ports of the 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

-
- Step 1** From the SCE(config if)# prompt, type `link failure-reflection on-all-ports` and press Enter.
Enables failure reflection to all ports.
-

How to Disable Link Failure Reflection on All Ports

-
- Step 1** From the SCE(config if)# prompt, type `no link failure-reflection` and press Enter.
Disables failure reflection to all ports.
-

How to Configure Link Failure Reflection in Linecard-Aware Mode (SCE 2000 only)

- [About Link Failure Reflection in Linecard-Aware Mode, page 7-12](#)
- [How to Enable Linecard-Aware Mode, page 7-12](#)
- [How to Disable Linecard-Aware Mode, page 7-13](#)

About Link Failure Reflection in Linecard-Aware Mode

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 of the SCE 2000 platform are connected to the same linecard in the router.

This mode reflects a failure of one port to the other three ports of the SCE 2000 differently, depending on different failure conditions, as follows:

- One interface of the SCE 2000 is down: Link failure is reflected to the all other SCE platform ports.
- Two reciprocal ports of the SCE 2000 are down simultaneously, indicating a possible problem in the linecard of the router to which the 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 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

-
- Step 1** From the SCE(config if)# prompt, type `link failure-reflection on-all-ports linecard-aware-mode` and press Enter.
Enables failure reflection to all ports with linecard aware mode.
-

How to Disable Linecard-Aware Mode

Step 1 From the SCE(config if)# prompt, type `no link failure-reflection linecard-aware-mode` and press Enter.

Disables linecard aware mode.

Note that this command does not disable link failure reflection on all ports.



CHAPTER 8

Raw Data Formatting: The RDR Formatter and NetFlow Exporting

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.

- [Information About the RDR Formatter and NetFlow Exporting Support, page 8-1](#)
- [How to Configure Data Destinations and Categories, page 8-6](#)
- [How to Configure the RDR Formatter, page 8-13](#)
- [How to Configure the NetFlow Exporting Support, page 8-14](#)
- [How to Configure Dynamic Mapping of RDRs to Categories, page 8-15](#)
- [How to Display Data Destination Configuration and Statistics, page 8-17](#)
- [How to Disable the Linecard from Sending RDRs, page 8-20](#)

Information About the RDR Formatter and NetFlow Exporting Support

- [The RDR Formatter, page 8-2](#)
- [NetFlow, page 8-2](#)
- [Data Destinations, page 8-3](#)

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 8-2](#)
- [NetFlow Exporting Support, page 8-3](#)

NetFlow Terminology

- **Exporter**
A device (in this case, the RDR formatter component in the 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 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 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

- [About Data Destinations, page 8-4](#)
- [Categories, page 8-5](#)
- [Priority, page 8-5](#)
- [Setting DSCP for NetFlow, page 8-5](#)
- [Forwarding Modes, page 8-6](#)
- [Protocol, page 8-6](#)
- [Transport Type, page 8-6](#)

About Data Destinations

The 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 or UDP)

The destination is assigned a priority for each category to which it is assigned.


The following figure illustrates the simplest data destination topology, with only one category and one destination.

Figure 8-1 Data Destination Topology: One Category and One Destination



The following figure illustrates a complex topology using two categories and four destinations. Each category can send data to three of the four destinations.

Figure 8-2 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 SCE platform.

The system supports up to four categories. Therefore, the 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, the user may define meaningful names for the categories. This generally reduces confusion and prevents errors.

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 be 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 NetFlow V9 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 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

How to Configure Data Destinations and Categories

- [How to Configure a Data Destination, page 8-7](#)
- [How to Configure the Data Categories, page 8-8](#)
- [How to Configure the Forwarding Mode, page 8-12](#)

How to Configure a Data Destination

- [About Configuring the Data Destinations, page 8-7](#)
- [Options, page 8-7](#)
- [Configuring the Data Destinations: Examples, page 8-7](#)

About Configuring the Data Destinations

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 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 [How to Configure a Destination and Assign Categories, page 8-9](#).

Step 1 From the SCE(config)# prompt, type `rdr-formatter destination ip-address port portnumber [priority priority] protocol protocol[transport transport]`, and press Enter.

Defines the destination. When no category is specified, as in the above example, the specified priority is assigned to all categories.

Configuring the Data Destinations: Examples

- [Example 1, page 8-7](#)
- [Example 2, page 8-8](#)

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
```

How to Configure the Data Categories

- [How to Configure a Destination and Assign Categories, page 8-9](#)
- [Options, page 8-8](#)

There are two steps in defining the data categories:

- Define the category names (optional).
- Configuring 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

Step 1 From the SCE(config)# prompt, type `rdr-formatter category number category-number name category-name`, and press Enter.

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.

How to Configure a Destination and Assign Categories

- [Options, page 8-9](#)
- [Configuring the Data Destinations with Categories: Examples, page 8-9](#)

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 be 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.

Step 1 From the SCE(config)# prompt, type `rdr-formatter destination ip-addressport portnumbercategory [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]`, and press Enter.

Defines the destination and assigns categories with optional priorities.

Configuring the Data Destinations with Categories: Examples

- [Example 1, page 8-10](#)
- [Example 2, page 8-10](#)
- [Example 3, page 8-11](#)
- [Example 4, page 8-11](#)
- [Example 5, page 8-12](#)


Example 1

The following example defines a name for one category, and then configures two destinations, assigning each to a different category (see diagram).

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 8-3 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. 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 8-4 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 8-5 *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 multi-cast
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:

- 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 8-6 Configuring Destinations: Three Categories and Two Protocols



```
SCE(config)# rdr-formatter forwarding-mode multi-cast
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
```

How to Configure 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, page 8-13](#)
- [Configuring the Forwarding Mode: Example, page 8-13](#)

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.

-
- Step 1** From the SCE(config)# prompt, type `rdr-formatter forwarding-mode mode`, and press Enter. 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
```

How to Configure the RDR Formatter

- [Options, page 8-13](#)
- [How to Enable the RDR Formatter, page 8-13](#)
- [How to Disable the RDR Formatter, page 8-14](#)
- [How to Configure the Size of the RDR Formatter History Buffer, page 8-14](#)

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. Note, however, that the size of the history buffer must be zero bytes (the default value). Other values may cause duplication of RDRs.
- Dynamic mapping of RDRs to categories (see [How to Configure Dynamic Mapping of RDRs to Categories, page 8-15](#))

How to Enable the RDR Formatter

-
- Step 1** From the SCE(config)# prompt, type `service rdr-formatter` , and press Enter. Enables the RDR formatter.
-

How to Disable the RDR Formatter

-
- Step 1** From the SCE(config)# prompt, type `no service rdr-formatter`, and press Enter.
Disables the RDR formatter.
-

How to Configure the Size of the RDR Formatter History Buffer

Options

The buffer size must be set to 0.

-
- Step 1** From the SCE(config)# prompt, type `rdr-formatter history-size 0` , and press Enter.
Sets the size of the RDR formatter history buffer.
-

How to Configure the NetFlow Exporting Support

- [Options, page 8-14](#)
- [How to Configure a DSCP Value for NetFlow, page 8-15](#)
- [How to Configure the Template Refresh Interval, page 8-15](#)

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)

Step 1 From the SCE(config)# prompt, type `rdr-formatter protocol NetFlowV9 dscp dscp-value`, and press Enter.

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 .)

Step 1 From the SCE(config)# prompt, type `rdr-formatter destination ip-addressport port-numberprotocol NetFlowV9 template data timeout timeout-value`, and press Enter.

Sets the template refresh interval.

How to Configure Dynamic Mapping of RDRs to Categories

- [About Dynamic Mapping of RDRs to Categories, page 8-16](#)
- [How to Configuring Mappings, page 8-16](#)

About 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.

How to Configuring Mappings

Use these command to add or remove a mapping.

- [Options, page 8-16](#)
- [How to Add a Mapping to a Category, page 8-16](#)
- [How to Remove a Mapping from a Category, page 8-16](#)
- [How to Restore the Default Mapping for a Specified RDR Tag, page 8-17](#)

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.

How to Add a Mapping to a Category

Step 1 From the SCE(config)# prompt, type `rdr-formatter rdr-mapping (tag-id tag-numbercategory-number category-number)` and press Enter.

If the table already contains a mapping with the same tag and category number, an error is issued and nothing is done.

How to Remove a Mapping from a Category

Step 1 From the SCE(config)# prompt, type `no rdr-formatter rdr-mapping (tag-id tag-numbercategory-number category-number)` and press Enter.

How to Restore the Default Mapping for a Specified RDR Tag

-
- Step 1** From the SCE(config)# prompt, type `default rdr-formatter rdr-mapping tag-id tag-number` and press Enter.
-

How to Display Data Destination Configuration and Statistics

- [How to the Display the Current RDR Formatter Configuration, page 8-17](#)
- [How to the Display the Current RDR Formatter Statistics, page 8-18](#)

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`

Refer to the *Cisco Service Control Engine CLI Command Reference* for a complete description of the other `show rdr-formatter` commands.

How to the Display the Current RDR Formatter Configuration

The system can display the complete data destination configuration, or just specific parameters.

-
- Step 1** From the SCE>prompt, type `show rdr-formatter` , and press Enter.
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 Address / | | |-----|
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
Last time these counters were cleared: 20:23:05 IST WED March 14 2007
```

How to the Display the Current RDR Formatter Statistics

-
- Step 1** From the SCE>prompt, type `show rdr-formatter statistics` , and press Enter.
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
```

How to Disable 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.

Step 1 From the SCE(config if)# prompt, type `silent` and press Enter.

Stops the linecard from producing RDRs.

To enable the linecard to produce data records, use the following command:

- `no silent`
-



CHAPTER 9

Managing Subscribers

The 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

- [What is a Subscriber?, page 9-2](#)
- [Subscriber Modes in Service Control Solutions, page 9-3](#)
- [Aging Subscribers, page 9-4](#)
- [Synchronizing Subscriber Information in a Cascade System, page 9-4](#)
- [Anonymous Groups and Subscriber Templates, page 9-5](#)
- [Information About Subscriber Files, page 9-5](#)

What is a Subscriber?

In the Service Control solution, a subscriber is defined as a managed entity on the subscriber side of the SCE Platform to which accounting and policy are applied individually.

The following table lists several examples of subscribers in Service Control solutions.

Table 9-1 **Subscriber Examples**

| The Subscriber | Subscriber Characteristics | Subscriber (Entity) Identified By |
|---|--|--|
| | Managed Entity | |
| 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 | <ul style="list-style-type: none"> • A range of IP addresses • A group of VLAN tags |
| A VPN | <ul style="list-style-type: none"> • MPLS/VPN-based subscriber • VLAN-based subscriber | <ul style="list-style-type: none"> • A set of IP addresses or ranges in a certain VPN • A VLAN tag or range of VLAN tags • All the IP addresses of a CE router, defined by a BGP community over a VPN |
| SCMP subscriber | SCMP subscriber | IP address or range Manager-Id of the SCMP peer device and 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 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 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 SCE platform. The 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 SCE platform. This is achieved by using the 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 SCE platforms.

MPLS/VPN-based subscribers are supported only in the dynamic subscriber aware mode. The system must dynamically map the internal MPLS label and the MAC address of the PE to the correct VPN subscriber.

Aging Subscribers

Subscribers can be aged automatically by the 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 SCE platform and are no longer occupying resources. Aging time can be configured individually for introduced subscribers and for anonymous subscribers.

Introduced subscriber aging is not supported when using VPN-based subscribers.

Synchronizing Subscriber Information in a Cascade System

In a hot standby, cascade setup with full redundancy, the external provisioning server updates only the active SCE platform. However, the standby SCE platform must always be updated with the latest subscriber-related information (login, logout). This is required to minimize information loss if fail-over. In general, the only entity that is allowed to change subscriber information in the standby SCE platform is the active SCE platform. The standby 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 SCE platform can change subscriber information of the default subscriber.
- Standby SCE platform can perform subscriber aging.

Therefore, when working as a pair, the active SCE platform constantly updates the standby SCE platform with external data information. In addition, the standby SCE platform constantly requests external data information from the active SCE platform. The synchronization is bi-directional to ensure that the subscriber databases in both 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 SCE platform if a fail-over.

The following subscriber information is considered as external data:

- subscriber name
- mappings (all supported types – IPs, VLANs, MPLS/VPN)
- tunables
- manager name
- aging time
- lease time
- is-static flag

Only the active SCE platform communicates with the SM. The SM is aware of the active/standby state of each SCE platform, and is also aware of a fail-over.

Specifically, this means the following:

- In push mode, the SM pushes events to the active SCE platform, which updates the standby SCE platform.
- In pull mode, only the active SCE platform pulls subscribers from the SM.
- The standby SCE platform can create anonymous subscribers based on the updates received from the active SCE platform, but does not generate pull-response for them.
- If SCE-SM connection failure, the SM handles the SCE recovery of the active SCE platform only. The active SCE platform propagates the information to the standby 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 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.

The SCE platform can support a maximum of 1000 anonymous groups. 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.

Information About Subscriber Files

- [Subscriber Files, page 9-6](#)
- [Subscriber default csv file format, page 9-6](#)
- [Subscriber anonymous groups csv file format, page 9-7](#)

Subscriber Files

**Note**

MPLS/VPN-based subscribers cannot be defined, imported, or exported by a subscriber file.

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 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 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, 10.1.12.32, 1
sub9, 5, 2
sub10, 13-17, 2
sub11, 39;41, 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, template-index, manager-name(optional)

- **name** — is the anonymous group name
- **IP-range** — dotted decimal, followed by the amount of significant bits. Example: 10.3.0.0/16
- **template-index** — is the index of the subscriber template to be used by subscribers belonging to this anonymous group.
- **manager-name** (optional) — 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, SCMP1
anon2, 10.1.2.0/24, 2, SCMP2
anon3, 10.1.3.0/32, 3, SCMP3
anon4, 10.1.4.0/24, 3, SCMP3
anon5, 10.1.5.0/31, 2, SM
anon6, 10.1.6.0/30, 1, SM
anon7, 0.0.0.0/0, 1, SM
```

How to Import and Export Subscriber Information

- [Options, page 9-8](#)
- [How to Import Subscriber Information, page 9-8](#)
- [How to Export Subscriber Information, page 9-8](#)
- [How to Import a Subscriber Template, page 9-8](#)
- [How to Export a Subscriber Template, page 9-8](#)

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
- 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.



Note

MPLS/VPN-based subscribers cannot be defined, imported, or exported by a subscriber file.

Options

The following option is available:

- **filename** — the name of the csv file.

How to Import Subscriber Information

-
- Step 1** From the SCE(config if)# prompt, type `subscriber import csv-file filename` and press Enter.
- Imports the subscriber information from the specified file.
- Imported subscriber information is added to the existing subscriber information. It does not overwrite the existing data.
- If the information in the imported file is not valid, the command will fail during the verification process before it is actually applied.
-

How to Export Subscriber Information

-
- Step 1** From the SCE(config if)# prompt, type `subscriber export csv-file filename` and press Enter.
- Exports the subscriber information to the specified file.
-

How to Import a Subscriber Template

-
- Step 1** From the SCE(config if)# prompt, type `subscriber template import csv-file filename` and press Enter.
- Imports the subscriber template from the specified file.
-

How to Export a Subscriber Template

-
- Step 1** From the SCE(config if)# prompt, type `subscriber template export csv-file filename` and press Enter.
- Exports the subscriber template to the specified file.
-

How to Remove Subscribers and Templates

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 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 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

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

How to Remove All Introduced Subscribers

-
- Step 1** From the SCE(config if)# prompt, type `no subscriber all` and press Enter.
Removes all introduced subscribers.
-

How to Remove a Specific Anonymous Subscriber

Options

The following option is available:

- **group-name** — the name of the anonymous subscriber group to be removed

Step 1 From the SCE(config if)# prompt, type `no subscriber anonymous-group name group-name` and press Enter.

Removes the specified anonymous subscriber group.

How to Remove All Anonymous Subscriber Groups

Step 1 From the SCE(config if)# prompt, type `no subscriber anonymous-group all` and press Enter.

Removes all anonymous subscriber groups.

How to Remove All Anonymous Subscribers

Step 1 From the SCE# prompt, type `clear interface linecard 0 subscriber anonymous all` and press Enter.

Removes all anonymous subscribers.



Note The `clear subscriber anonymous` command is a **Privileged Exec** command.

How to Remove All Subscriber Templates

Step 1 From the SCE(config if)# prompt, type `default subscriber template all` and press Enter.

Removes all subscriber templates.

All anonymous subscribers will be assigned to the default subscriber template.

How to Remove VPN-based Subscribers

About 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 SCE platform), use this command.

This command allows you to switch out of MPLS/VPN mode without reload when the SM is down.

**Note**

Use this command **ONLY** when the SCE platform is disconnected from the SM.

- Step 1** From the SCE(config if)# prompt, type `no subscriber all with-vpn-mappings` and press Enter. Clears all VPN-based subscribers.
-

How to Remove Subscribers by Device

- [How to Remove Subscribers from the SM, page 9-11](#)
- [How to Remove Subscribers from a Specified SCMP Peer Device, page 9-11](#)

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

-
- Step 1** From the SCE(config if)# prompt, type `no subscriber sm all` and press Enter. Clears all subscribers from the SM.
-

How to Remove Subscribers from a Specified SCMP Peer Device

Options

The following option is available:

- **peer-device-name** — the name of the SCMP peer device from which to clear the subscribers.

-
- Step 1** From the SCE(config if)# prompt, type `no subscriber scmp name peer-device-name all` and press Enter. Clears all subscribers from the specified SCMP peer device.
-

How to Import and Export Anonymous Groups

- [How to Import Anonymous Groups, page 9-12](#)
- [How to Export Anonymous Groups, page 9-12](#)

How to Import Anonymous Groups

Options

The following option is available:

- **filename** — name of the csv file.

Step 1 From the SCE(config if)# prompt, type `subscriber anonymous-group import csv-file filename` and press Enter.

Creates anonymous groups by importing anonymous subscribers from the specified csv file.

Imported anonymous subscriber information is added to the existing anonymous subscriber information. It does not overwrite the existing data.

The SCE platform can support a maximum of 1000 anonymous groups.

How to Export Anonymous Groups

Options

The following option is available:

- **filename** — name of the csv file.

Step 1 From the SCE(config if)# prompt, type `subscriber anonymous-group export csv-file filename` and press Enter.

Exports all existing anonymous groups to the specified csv file.

How to Monitor Subscribers

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 SCE platform via the 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

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

- Step 1** From the SCE# prompt, type `show interface linecard 0 subscriber db counters` and press Enter. 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: 2 used out of 99999 max.
Introduced subscribers: 2.
Anonymous subscribers: 0.
Subscribers with mappings: 2 used out of 99999 max.
SINGLE non-VPN IP mappings: 1.
non-VPN IP Range mappings: 1.
IP Range over VPN mappings: 1.
Single IP over VPN mappings: 3.
MPLS-based subscribers are enabled.
MPLS/VPN mappings: 2 used out of 57344 max.
MPLS-based subscribers with VPN mappings: 2 used out of 2015 max.
Subscribers with open sessions: 0.
Subscribers with TIR mappings: 0.
Sessions mapped to the default subscriber: 0.
Peak values:
=====
Peak number of subscribers with mappings: 2
Peak number occurred at: 14:56:55 ISR MON June 9 2007
Peak number cleared at: 15:29:39 ISR MON June 9 2007
Event counters:
=====
Subscriber introduced: 2.
Subscriber pulled: 0.
Subscriber aged: 0.
Pull-request notifications sent: 0.
State notifications sent: 0.
Logout notifications sent: 0.
Subscriber mapping TIR contradictions: 0
```

Clearing the Subscriber Database Counters

- Step 1** From the SCE# prompt, type `clear interface linecard 0 subscriber db counters` and press Enter. Clears the “**total**” and “**maximum**” counters..

How to Display Subscribers

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 (may be within a specified VPN).
- Mapped to a specified VLAN ID.
- Mapped to a specified MPLS/VPN.

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'`
- `show interface linecard 0 subscriber mapping MPLS-VPN PE-ID 'pe-id' BGP-label 'bgp-label'`

Displaying Subscribers: All Current Subscriber Names

You can display the names of all subscribers currently in the SCE subscriber database.

-
- Step 1** From the SCE>prompt, type `show interface linecard 0 subscriber all-names` and press Enter. Displays the names of all subscribers currently in the 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 9-16](#)
- [How to display subscribers that are greater than or less than a specified value of a subscriber property, page 9-16](#)
- [How to display subscribers that match a specified prefix, page 9-17](#)
- [How to display subscribers that match a specified suffix, page 9-17](#)
- [How to display the number of subscribers that match a specified value of a subscriber property, page 9-17](#)
- [How to display the number of subscribers that are greater than or less than a specified value of a subscriber property, page 9-17](#)
- [How to display the number of subscribers that match a specified prefix, page 9-18](#)

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

Step 1 From the SCE>prompt, type `show interface linecard 0 subscriber property propertyname equals property-val` and press Enter.

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

Step 1 From the SCE>prompt, type `show interface linecard 0 subscriber property propertyname greater-than|less-than property-val` and press Enter.

How to display subscribers that match a specified prefix

Options

The following options are available:

- **prefix** — subscriber prefix to match

Step 1 From the SCE>prompt, type `show interface linecard 0 subscriber prefix prefix` and press Enter.

How to display subscribers that match a specified suffix

Options

The following options are available:

- **suffix** — subscriber suffix to match

Step 1 From the SCE>prompt, type `show interface linecard 0 subscriber suffix suffix` and press Enter.

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

Step 1 From the SCE>prompt, type `show interface linecard 0 subscriber amount property propertyname equals property-val` and press Enter.

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

Step 1 From the SCE>prompt, type `show interface linecard 0 subscriber amount property propertyname greater-than|less-than property-val` and press Enter.

How to display the number of subscribers that match a specified prefix

Options

The following options are available:

- **prefix** — subscriber prefix to match

Step 1 From the SCE>prompt, type `show interface linecard 0 subscriber amount prefix prefix` and press Enter.

How to Display Subscribers: By Mapping (IP Address, VPN, VLAN ID, or MPLS/VPN)

- [How to display subscribers that are mapped to a specified IP address, or range of IP addresses, page 9-18](#)
- [How to display subscribers that are mapped to IP addresses that are included in a given IP address or IP range, page 9-19](#)
- [How to display subscribers that are mapped to a specified VLAN ID, page 9-19](#)
- [How to display subscribers that are mapped to a specified MPLS/VPN, page 9-19](#)
- [How to display subscribers with no mapping, page 9-19](#)
- [How to display the number of subscribers that are mapped to a specified VLAN ID, page 9-20](#)
- [How to display the number of subscribers with no mapping, page 9-20](#)

You can display the subscribers who are mapped to any of the following:

- A specified IP address, or range of IP addresses (may be within a specified VPN)
- IP addresses intersecting a given IP address or IP range (may be within a specified VPN)
- A specified VLAN ID
- A specified MPLS/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

Step 1 From the SCE>prompt, type `show interface linecard 0 subscriber mapping IP ip-range [VPN vpn-name]` and press Enter.

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

-
- Step 1** From the SCE>prompt, type `show interface linecard 0 subscriber mapping included-in IP ip-range [VPN vpn-name]` and press Enter.
-

How to display subscribers that are mapped to a specified VLAN ID**Options**

The following options are available:

- **VLAN-id** — VLAN ID to match

-
- Step 1** From the SCE>prompt, type `show interface linecard 0 subscriber mapping VLAN-id VLAN-id` and press Enter.
-

How to display subscribers that are mapped to a specified MPLS/VPN**Options**

The following options are available:

- **pe-id** — loopback IP address of the relevant PE router
- **bgp-label** — label of the relevant BGP LEG

-
- Step 1** From the SCE>prompt, type `show interface linecard 0 subscriber mapping MPLS-VPN PE-ID pe-id bgp-label BGP-label` and press Enter.
-

How to display subscribers with no mapping

-
- Step 1** From the SCE>prompt, type `show interface linecard 0 subscriber mapping none` and press Enter.
-

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

Step 1 From the SCE>prompt, type `show interface linecard 0 subscriber amount mapping VLAN-id VLAN-id` and press Enter.

How to display the number of subscribers with no mapping

Step 1 From the SCE>prompt, type `show interface linecard 0 subscriber amount mapping none` and press Enter.

How to Display Subscriber Information

You can display the following information about a specified subscriber:

- values of the various subscriber properties
- mappings (IP address, VLAN-ID or MPLS/VPN)
- 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`
- `show interface linecard 0 subscriber name 'name' vas-servers`

How to display a listing of subscriber properties

Step 1 From the SCE>prompt, type `show interface linecard 0 subscriber properties` and press Enter.

How to display complete information for a specified subscriber

Use this command to display complete information for a specified subscriber, including all values of subscriber properties and mappings.

Options

The following options are available:

- **name** — subscriber name

Step 1 From the SCE>prompt, type `show interface linecard 0 subscriber name name` and press Enter.

How to display values of subscriber properties for a specified subscriber

Options

The following options are available:

- **name** — subscriber name

Step 1 From the SCE>prompt, type `show interface linecard 0 subscriber name nameproperties` and press Enter.

How to display mappings for a specified subscriber

Options

The following options are available:

- **name** — subscriber name

Step 1 From the SCE>prompt, type `show interface linecard 0 subscriber name namemappings` and press Enter.

How to display OS counters for a specified subscriber

Options

The following options are available:

- **name** — subscriber name

Step 1 From the SCE>prompt, type `show interface linecard 0 subscriber name namecounters` and press Enter.

How to Display Anonymous Subscriber Information

You can display the following information regarding the anonymous subscriber groups:

- aging (see [How to Display Aging for Anonymous Group Subscribers, page 9-33](#))
- 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

-
- Step 1** From the SCE>prompt, type `show interface linecard 0 subscriber anonymous-group all` and press Enter.
-

How to display currently configured templates for anonymous groups

-
- Step 1** From the SCE>prompt, type `show interface linecard 0 subscriber templates` and press Enter.
-

How to display current configuration for a specified anonymous group

Options

The following options are available:

- **group-name** — name of the anonymous subscriber group

-
- Step 1** From the SCE>prompt, type `show interface linecard 0 subscriber anonymous-group name group-name` and press Enter.
-

How to display subscribers in a specified anonymous group

Options

The following options are available:

- **group-name** — name of the anonymous subscriber group

Step 1 From the SCE>prompt, type `show interface linecard 0 subscriber anonymous name group-name` and press Enter.

How to display all subscribers currently in anonymous groups

Step 1 From the SCE>prompt, type `show interface linecard 0 subscriber anonymous` and press Enter.

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

Step 1 From the SCE>prompt, type `show interface linecard 0 subscriber amount anonymous name group-name` and press Enter.

How to display the total number of subscribers in all anonymous groups

Step 1 From the SCE>prompt, type `show interface linecard 0 subscriber amount anonymous` and press Enter.

Subscriber Traffic Processor IP Ranges

- [Information About Traffic Processor IP Ranges, page 9-24](#)
- [How to Reserve Rules for TIRs, page 9-26](#)
- [How to Configure TIRs, page 9-26](#)
- [How to Remove TIRs and Subscriber Mappings, page 9-27](#)
- [How to Import and Export TIRs, page 9-28](#)
- [How to Monitor TIRs, page 9-30](#)

Information About Traffic Processor IP Ranges

- [Subscriber Mapping Modes, page 9-25](#)
- [Subscriber Mapping Conflicts, page 9-25](#)
- [Subscriber Rules for TIRs, page 9-25](#)

**Note**

Traffic Processor IP Range functionality is relevant only for IP-based subscribers. This functionality is not relevant for VLAN or MPLS/VPN subscribers.

In a Cable environment, the SCE platform supports the capability of associating all CPE machines in a single home network (i.e. behind a single cable modem) to a single subscriber-context and applying a single policy to this subscriber context. This is also relevant for cases where each CPE uses multiple global IP addresses (as opposed to a residential gateway that performs NAT and allows all CPE machines to share an IP address).

The SCE platform places no limit on the number of subscribers that can have multiple IP addresses. To achieve this, all IP addresses used by each CPE must use a common pool of addresses (usually that assigned with their downstream CMTS device/blade), and the subscriber that uses all these CPEs should be configured to a single traffic processor (a single PPC in the SCE platform).

Assigning subscribers to a specific traffic processor can be implemented in either of the following scenarios:

- All the IP ranges of a given CMTS/BRAS are configured to be processed by the same traffic processor. This can only be done if one SCE platform is handling several CMTS/BRAS (otherwise there is a load-balancing issue).
- The service provider can control the IP range from which the subscriber IP address is allocated based on additional criteria such as the subscriber type. In this case, the range can be used by the SCE platform to assign subscribers to a particular traffic processor, independent of the definition of the subscriber network ID.

In such cases the SCE platform (based on management configuration) can ensure that the various IP addresses (either ranges or single IPs) of each subscriber will actually be handled by the same traffic processor. This is accomplished by assigning a subscriber IP range (or specific IP address) to a configured Traffic Processor IP Range (TIR). Since each TIR is assigned to a traffic processor, the relevant subscriber IP range is also assigned to the matching traffic processor. Note that all ranges and single IPs of a specific subscriber must be assigned to the same traffic processor at any given time.

It is assumed that editing TIR configuration (addition or removal) is done infrequently. Also, that it is generally done either before the relevant IP ranges are in use or after they are no longer in use.

Subscriber Mapping Modes

The introduction of the TIR functionality provides two possible modes of subscriber mapping:

- Legacy subscriber mapping — ensures that all mappings of a single subscriber reach the same traffic processor by internal means, using a hash on the subscriber IP and/or using specific subscriber rules on the IP/range when required
- TIR subscriber mapping — generally (regarding any relevant subscribers) configures all mappings in a specific range to reach the same traffic processor, reducing the need for internal specific rule resources per subscriber.

TIRs functionality is not necessarily applicable to all subscribers. Therefore, while the user may choose to assign relevant subscribers to traffic processors via TIRs (TIRs subscriber mapping), the remaining subscribers are processed as usual (legacy subscriber mapping).

Subscriber Mapping Conflicts

It is important to note that while both subscriber mapping modes can co-exist in one deployment, any individual subscriber can be processed only in one mode or the other. The same subscriber cannot be processed partially using TIRs subscriber mappings and partially using legacy subscriber mappings. The resulting conflicting subscriber mappings will be rejected.

Another cause of conflicting subscriber mappings is when a subscriber is assigned a new range or single IP that is associated with a traffic processor, different from that with which the subscriber is already associated.

Conflicting mapping are rejected (any other subscriber mappings are accepted as is) in both cases below:

- Conflict between mappings of a single mapping request.
- Additive subscriber mappings that conflict with existing mappings.

Subscriber Rules for TIRs

The number of reserved rules for potential TIRs is configurable, and is at the expense of explicit subscriber rules. The total number of subscriber rules available is approximately 8000.

- The maximum number of allowed reserved rules is 4096. The remaining rules are reserved for explicit subscriber mappings usage (used by the SCE platform to enable the legacy internal OS allocation of subscribers to traffic processors).
- By default 0 (zero) rules are reserved for TIRs.
- Updating this configuration is a major system event and can only be performed when no subscriber mappings or TIRs are configured.

How to Reserve Rules for TIRs

Use these commands to reserve rules for TIRs and to restore default subscriber rule allocation.

Options

The following options are available:

- **#rules** — number of subscriber rules to be reserved for TIRs

Step 1 From the SCE(config if)# prompt, type `subscriber TP-mappings max-TP-IP-ranges #rules` and press Enter.

Allocates the specified number of rules for TIRs, decreasing the number of explicit subscriber rules available.

To restore all rules as explicit subscriber rules, with no rules reserved for TIRs (default), use the following command:

- **default subscriber TP-mappings**
-

How to Configure TIRs

- [Options, page 9-26](#)
- [How to Create or Update a TIR, page 9-27](#)
- [How to update a TIR even if subscriber mappings exist, page 9-27](#)

Options

The following options are available:

- **tir-name** — meaningful name assigned to this traffic processor IP range
- **ip-address/mask** — IP address and mask defining the IP range
- **traffic-processornumber** — number of the traffic processor to which this TIR is to be assigned
- **remove-subscriber-mappings** — Editing TIRs is permitted only if there are no subscriber mappings within the relevant IP ranges. Therefore, by default, if subscriber mappings already exist for the either an updated or an existing IP range, the command will fail. However, you can use this keyword to specify that any existing subscriber mappings in the IP range should be removed.

In this case the command will execute successfully even if subscriber mappings exist.

How to Create or Update a TIR

-
- Step 1** From the SCE(config if)# prompt, type `subscriber TP-IP-range name tir-nameIP-range ip-address/masktarget-TP traffic-processornumber` and press Enter.
- **Creating** — Creates a TIR with the specified name and IP range and assigns it to the specified traffic processor
 - **Updating** — Updates the IP range and/or assigned traffic processor for the specified TIR.
 - **Updating the IP range** — If subscriber mappings exist for either the new or the old IP range, the command will fail.
-

How to update a TIR even if subscriber mappings exist

-
- Step 1** From the SCE(config if)# prompt, type `subscriber TP-IP-range name tir-nameIP-range ip-address/masktarget-TP traffic-processornumberremove-subscriber-mappings` and press Enter.
- If subscriber mappings exist for either the new or the old IP range, they will be removed and the command will execute successfully.
-

How to Remove TIRs and Subscriber Mappings

- [How to Remove a Specified TIR, page 9-28](#)
- [How to Remove All TIRs, page 9-28](#)
- [How to Remove Mappings from a Specified TIR, page 9-28](#)
- [How to Remove Mappings from a Specified IP Range, page 9-28](#)

Use these commands to remove existing TIRs and subscriber mappings. You can perform the following operations:

- Remove a specified TIR
- Remove all TIRs
- Remove all subscriber mappings assigned to a specified TIR
- Remove all subscriber mappings assigned to a specified IP range

As with updating a TIR, by default, if subscriber mappings already exist for the specified IP range, the command will fail. However, you can specify that any existing subscriber mappings in the IP range should be removed (`remove-subscriber-mappings` keyword). In this case the command will execute even if subscriber mappings exist.

How to Remove a Specified TIR

Step 1 From the SCE(config if)# prompt, type `no subscriber TP-IP-range name tir-name [remove-subscriber-mappings]` and press Enter.

If subscriber mappings exist for this IP range, the command will fail. Specify **remove-subscriber-mappings** to remove any existing subscriber mappings for this IP range, and the command will execute successfully.

How to Remove All TIRs

Step 1 From the SCE(config if)# prompt, type `no subscriber TP-IP-range all [remove-subscriber-mappings]` and press Enter.

If subscriber mappings exist for this IP range, the command will fail. Specify **remove-subscriber-mappings** to remove any existing subscriber mappings for this IP range, and the command will execute successfully.

How to Remove Mappings from a Specified TIR

Step 1 From the SCE(config if)# prompt, type `no subscriber mappings included-in TP-IP-range name tir-name` and press Enter.

The **remove-subscriber-mappings** option is not applicable to this command.

How to Remove Mappings from a Specified IP Range

Step 1 From the SCE(config if)# prompt, type `no subscriber mappings included-in IP-range ip-address/mask` and press Enter.

The **remove-subscriber-mappings** option is not applicable to this command.

How to Import and Export TIRs

Use these commands to import TIR definitions from a csv file and to export TIR definitions to a csv file.

- [About TIR csv Files, page 9-29](#)
- [Options, page 9-29](#)
- [How to Import TIRs from a csv File, page 9-29](#)
- [How to Export TIRs to a csv File, page 9-29](#)

About TIR csv Files

Following is the format of the csv file:

range name, ip-address/mask-length, target-TP

- **range name** — The name of the to which the IP addresses will be assigned
- **ip-address/mask-length** — individual IP address of range of IP addresses indicated by IP address/mask
- **target-TP** — traffic processor to which the specified range will be assigned
- **remove-subscriber-mappings** — When importing TIR definitions, by default, if subscriber mappings already exist for any specified IP range, those IP ranges will not be updated by the import. However, you can use this keyword to specify that any existing subscriber mappings in the IP range should be removed.

In this case the command will execute successfully even if subscriber mappings exist.

Options

The following option is available:

- **filename** — the name of the csv file.

How to Import TIRs from a csv File

Step 1 From the SCE(config if)# prompt, type `subscriber TP-IP-range import csv-file file-name [remove-subscriber-mappings]` and press Enter.

Imports the TIR definitions from the specified csv file.

If the **remove-subscriber-mappings** keyword is specified, if subscriber mappings exist for any specified IP range, they will be removed and the command will execute successfully. Otherwise, if subscriber mappings exist for any IP range, those IP ranges will not be updated.

How to Export TIRs to a csv File

Step 1 From the SCE(config if)# prompt, type `subscriber TP-IP-range export csv-file file-name` and press Enter.

Exports current TIR definitions to the specified csv file.

The **remove-subscriber-mappings** option is not applicable to this command.

How to Monitor TIRs

Use these commands to monitor TIRs and subscriber mappings. You can view the following:

- Traffic processor mappings state, including the partitioning between subscriber and TIR mappings, and the utilization of each.
- Configuration of a specified TIR Configuration of all TIRs
- All subscriber mappings related to a specified TIR
- Number of subscribers with mappings related to a specified TIR
- Information for a specified subscriber, including assigned TIR, where applicable
- All subscriber mappings in a specified IP range
- Number of subscribers with mappings in a specified IP range

How to Display Traffic Processor Mappings State

-
- Step 1** From the SCE>prompt, type `show interface linecard 0 subscriber TP-mappings statistics` and press Enter.
-

How to Display Configuration of a Specified TIR

-
- Step 1** From the SCE>prompt, type `show interface linecard 0 subscriber TP-IP-range name tir-name` and press Enter.
-

How to Display Configuration of All TIRs

-
- Step 1** From the SCE>prompt, type `show interface linecard 0 subscriber TP-IP-range all` and press Enter.
-

How to Display Mappings Related to a Specified TIR

-
- Step 1** From the SCE>prompt, type `show interface linecard 0 subscriber mapping included-in TP-IP-range name tir-name` and press Enter.
-

How to Display the Number of Subscribers with Mappings Related to a Specified TIR

-
- Step 1** From the SCE>prompt, type `show interface linecard 0 subscriber amount mapping included-in TP-IP-range name tir-name` and press Enter.
-

How to Display Complete Subscriber Information

-
- Step 1** From the SCE>prompt, type `show interface linecard 0 subscriber name name` and press Enter. Displays complete subscriber information, including which TIR the subscriber belongs to (if applicable).
-

How to Display All Subscribers Mapped to a Specified IP Range

-
- Step 1** From the SCE>prompt, type `show interface linecard 0 subscriber mapping included-in TP-IP-range ip-address/mask` and press Enter.
-

How to Display the Number of Subscribers Mapped to a Specified IP Range

-
- Step 1** From the SCE>prompt, type `show interface linecard 0 subscriber amount mapping included-in TP-IP-range ip-address/mask` and press Enter.
-

How to Configure Subscriber Aging

As explained previously ([Aging Subscribers, page 9-4](#)), 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

-
- Step 1** From the SCE(config if)# prompt, type `subscriber aging anonymous` and press Enter.
-

How to Enable Aging for Introduced Subscribers

**Note**

Introduced subscriber aging is not supported when using VPN-based subscribers

Step 1 From the SCE(config if)# prompt, type `subscriber aging introduced` and press Enter.

How to Disable Aging for Anonymous Group Subscribers

Step 1 From the SCE(config if)# prompt, type `no subscriber aging anonymous` and press Enter.

How to Disable Aging for Introduced Subscribers

Step 1 From the SCE(config if)# prompt, type `no subscriber aging introduced` and press Enter.

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.

Step 1 From the SCE(config if)# prompt, type `no subscriber aging anonymous timeout aging-time` and press Enter.

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.

Step 1 From the SCE(config if)# prompt, type `no subscriber aging introduced timeout aging-time` and press Enter.

How to Display Aging for Anonymous Group Subscribers

Step 1 From the SCE>prompt, type `show interface linecard 0 subscriber aging anonymous` and press Enter.

How to Display Aging for Introduced Subscribers

Step 1 From the SCE>prompt, type `show interface linecard 0 subscriber aging introduced` and press Enter.

How to Configure the SCE Platform/SM Connection

The user can configure the behavior of the 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 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.

Options

The following options are available:

- **action** — the action to take in case of failure of the SCE platform/SM connection
 - **force-failure** — force failure of SCE platform. The SCE platform then acts according to the behavior configured for the failure state.
 - **remove-mappings** — remove all current subscriber mappings.
 - **shut** — the SCE platform shuts down and quits providing service.
 - **none** (default) — take no action.
- **interval** — the time interval, in minutes, after which the connection is timed out

How to Configure the Behavior of the SCE Platform in Case of Failure of the SM

-
- Step 1** From the SCE(config if)# prompt, type `subscriber sm-connection-failure action force-failure|none|remove-mappings|shut` and press Enter.
-

How to Configure the SM-SCE Platform Connection Timeout

-
- Step 1** From the SCE(config if)# prompt, type `subscriber sm-connection-failure action timeout interval` and press Enter.
-



CHAPTER 10

Redundancy and Fail-Over

- [Information About Redundancy and Fail-Over, page 10-1](#)
- [How to Configure Forced Failure, page 10-5](#)
- [Hot Standby and Fail-over, page 10-5](#)
- [Recovery, page 10-9](#)
- [CLI Commands for Cascaded Systems, page 10-10](#)
- [System Upgrades, page 10-13](#)

Information About Redundancy and Fail-Over

- [Terminology and Definitions, page 10-2](#)
- [Redundant Topologies, page 10-2](#)
- [In-line Dual Link Redundant Topology, page 10-3](#)
- [Failure Detection, page 10-4](#)

This module presents the fail-over and redundancy capabilities of the SCE platform. It first defines relevant terminology, as well as pertinent theoretical aspects of the redundancy and fail-over 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 SCE platform deployments. When fail over is required in a deployment, a topology with two cascaded SCE platforms is used. This cascaded solution provides both network link fail over, and fail over of the functionality of the SCE platform, including updated subscriber state.



Note

The information in this chapter applies to the SCE 2000 4xGBE and SCE 2000 4/8xFE platforms only.

Terminology and Definitions

Following is a list of definitions of terms used in the chapter as they apply to the Cisco fail-over solution, which is based on cascaded SCE platforms.

- **Fail-over** — A situation in which the SCE platform experiences a problem that makes it impossible for it to provide its normal functionality, and a second SCE platform device immediately takes over for the failed SCE platform.
- **Hot standby** — When two SCE platforms are deployed in a fail over topology, one SCE platform is active, while the second SCE platform is in standby, receiving from the active 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 SCE platform. The Primary 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 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 fail-over** — A fail over solution in which subscriber state is saved.

Redundant Topologies

All Cisco SCE platforms include an internal electrical bypass module, which provide the capability of preserving the network link in case the SCE platform fails. The SCE platform, which can handle two data links, includes two such bypass modules. However, in some cases, the service provider wishes to preserve the SCE platform functionality in case of a failure, in addition to preserving the network link.

Cisco provides a unique solution for this scenario, through deploying two cascaded SCE platforms on these two data links.

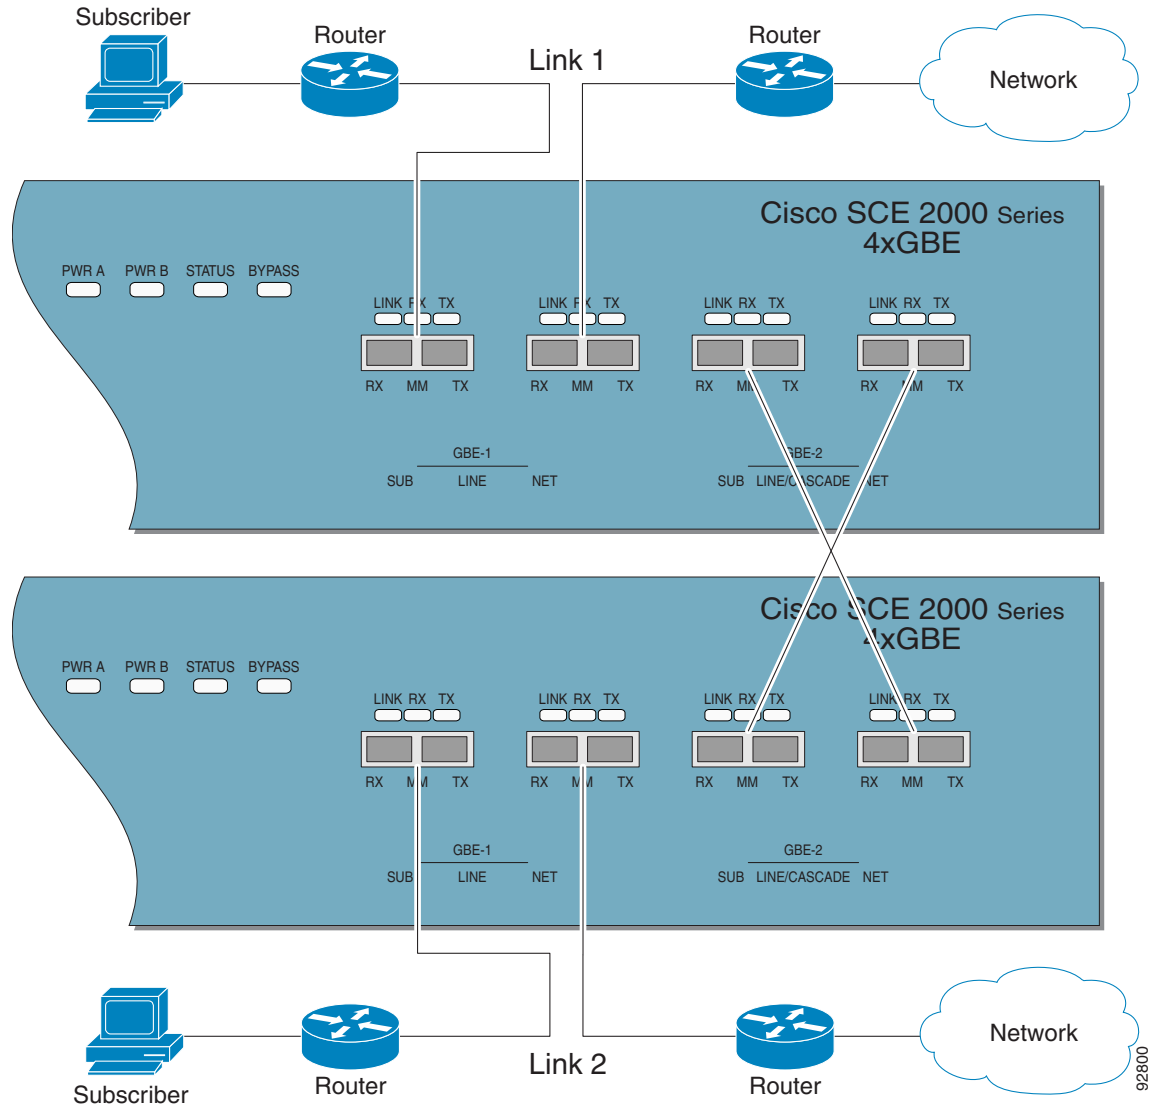
The cascading is implemented by connecting the two SCE platforms using two of the data links. This fail over solution applies to both inline and receive-only topologies.

In each 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 SCE platforms. (See the Cisco SCE 2000 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.

In-line Dual Link Redundant Topology

This topology serves inline deployments where the SCE platform functionality should be preserved in case of a failure, in addition to preserving the network link.

Figure 10-1 In-line Dual Link Redundant Topology



Failure Detection

The SCE platform has several types of mechanisms for detecting failures:

- Internal failure detection — The SCE platform monitors for hardware and software conditions such as overheating and fatal software errors.
- Inter-device failure detection — The SCE platform sends periodic keep-alive messages via the cascade ports
- SCE platform-Subscriber Manager (SM) communication failure detection — A failure to communicate with the SM may be regarded as a cause for fail over. However, this communication failure is not necessarily a problem in the SCE platform. If the connection to the SM of the active SCE platform has failed, while the connection to the SM of the standby SCE platform is alive, a fail over process will be initiated to allow the SCE platform proper exchange of information between the 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 SCE platform.
 - Cascade port link failure — Traffic cannot flow between the 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 SCE platform.

This type of failure, in most cases, does not require reboot of the SCE platform. When the connection with the SM is re-established the SCE platform is again ready for hot standby. If both 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 SCE platform.

Link Failure Reflection

The SCE platforms are transparent at Layers 2 and 3. The SCE platform operates in promiscuous mode, and the network elements on both sides of the 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 SCE platform to identify the link failures as quickly as possible, the SCE platform supports a functionality of reflecting to the other side of the SCE platforms events of link failure. When the link on one side of the 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 SCE platform with two data links, link failure is reflected between the two ports of each link.

When working with two cascaded SCE platforms, link failure is reflected in two cases:

- Reflection between the traffic ports of each SCE platform.
- If there is a failure in the cascade port link, the two SCE platforms can no longer support proper processing of the two links, since the traffic flowing on the standby SCE platform's link must be forwarded to the active SCE platform for processing. In this case the link failure is reflected from the cascade ports to the traffic ports of the standby SCE platform, in order to force the network to switch all the traffic only through the link of the active SCE platform.

Link failure reflection is supported both when the SCE platform is operational and when it is in failure/boot status.

Link reflection, like fail-over, is dependent on the bypass mechanism of the SCE platform

How to Configure Forced Failure

Use the following commands to force a virtual failure condition, and to exit from the failure condition when performing an application upgrade. (See [How to Manage Application Files, page 3-10.](#))

- [How to Force a Virtual Failure Condition, page 10-5](#)
- [How to Exit a Virtual Failure Condition, page 10-5](#)

How to Force a Virtual Failure Condition

-
- Step 1** From the SCE(config if)# prompt, type `force failure-condition` and press Enter.
Forces the SCE platform into a virtual failure state.
-

How to Exit a Virtual Failure Condition

-
- Step 1** From the SCE(config if)# prompt, type `no force failure-condition` and press Enter.
Exits from the virtual failure state.
-

Hot Standby and Fail-over

The fail over solution requires two SCE platforms connected in a cascade manner.

- [Hot Standby, page 10-6](#)
- [Fail-over, page 10-6](#)
- [Failure in the Cascade Connection, page 10-7](#)
- [Installing a Cascaded System, page 10-8](#)

Hot Standby

In fail over solution, one of the SCE platforms is used as the active SCE platform and the other is used as the standby. Although traffic enters both the active and the standby SCE platforms, all traffic processing takes place in the SCE platform which is currently the active one. The active SCE platform processes the traffic coming on both links, its own link and the link connected to the standby SCE platform, as follows

- All traffic entering the active SCE platform through its traffic ports is processed in that SCE platform and then forwarded to the line.
- All traffic entering the standby SCE platform through its traffic ports is forwarded through the cascade ports to the active SCE platform where it is processed, and then returned to the standby SCE platform through the cascade ports to be forwarded to the original line from which it came.

Since only one 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 fail-over, both SCE platforms hold subscriber states for all parties, and subscriber state updates are exchanged between the active SCE platform and the standby. This way, if the active SCE platform fails, the standby SCE platform is able to start serving the line immediately with a minimum loss of subscriber-state.

The two SCE platforms also use the cascade channel for exchanging periodic keep-alive messages.

Fail-over

In fail over solution, the two SCE platforms exchange keep alive messages via the cascade ports. This keep alive mechanism enables fast detection of failures between the SCE platforms and fast fail over to the standby SCE platform when required.

If the active SCE platform fails, the standby SCE platform then assumes the role of the active SCE platform.

The failed 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 SCE platform, and to forward processed traffic back to the link. The previously standby SCE platform now processes all the traffic of this other link that is forwarded to it by the previously active SCE platform in addition to the traffic of its own link.

When the failed SCE platform recovers, it will remain in standby, while the previously standby SCE platform remains active. Switching the SCE platforms back to their original roles may be performed manually, if required, after the failed SCE platform has either recovered or been replaced.

If the failure is in the standby SCE platform, it will continue to forward traffic to the active SCE platform and back to its link, while the active 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 SCE platforms and the effect of fail-over on the subscriber databases, see [Synchronizing Subscriber Information in a Cascade System, page 9-4](#).

There are two user-configurable options that are relevant in a situation when an SCE platform fails:

- **Bypass** — Maintain the link in bypass mode (continue sending traffic to the other SCE platform, and then continue forwarding the processed traffic back to the link). The incoming traffic in the failed SCE platform is forwarded to the working SCE platform, where it is processed and then sent back to the original SCE platform and back to the link.
 - **Effect on the network link** — negligible.
 - **Effect on the SCE platform functionality** — The effect on the SCE platform functionality is dependent on the failed SCE platform.
 - If the failure is in the standby SCE platform — the active SCE platform continues providing its normal functionality, processing the traffic of the two links.
 - If the failure is in the active SCE platform — the standby SCE platform takes over processing the traffic, and becomes the active SCE platform.
- **Cutoff** — Change the link of the failed SCE platform to cutoff (layer 1) forcing the network to switch all traffic through the line of the working SCE platform. This will, of course, decrease the network capacity by 50%, but may be useful in some cases.
 - **Effect on the network** — The network loses 50% of its capacity (until the failed SCE platform has recovered).
 - **Effect on the SCE platform functionality** — The effect on the SCE platform functionality is dependent on the failed SCE platform:
 - If the failure is in the standby SCE platform — SCE platform continues providing its normal functionality, processing the traffic of its own link.
 - If the failure is in the active SCE platform — the standby SCE platform takes over processing the traffic, and becomes the active SCE platform. This option is available for use in special cases, and requires specific configuration.

Failure in the Cascade Connection

The effect of a failure in the cascade connection between the two SCE platforms depends on whether one or both connections fail:

- Only one cascade connection is down — In this case, both 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 SCE platform. Therefore, split flow is avoided, but at the expense of half line capacity.
- Both cascade links are down — In this case, neither SCE platform knows anything about the status of the peer. Each platform then works in standalone mode, which means that each 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 SCE platforms.

Refer to the *Cisco SCE 2000 Installation and Configuration Guide* for information on topologies and connections.

Refer to the *Cisco Service Control Engine (SCE) CLI Command Reference* for details of the CLI commands.



Note

When working with two SCE platforms with split-flow and redundancy, it is extremely important to follow this installation procedure.

- Step 1** Install both SCE platforms, power them up, and perform the initial system configuration.
 - Step 2** Connect both SCE platforms to the management station.
 - Step 3** Connect the cascade ports. The cascade ports must be connected directly in Layer 1 (dark fibers), not through a switch. This means that two SCE platforms configured in this mode must look like they are physically connected, with no switch or router between them.
 - Step 4** Set topology configurations for each SCE platform via the connection-mode options. (See [Topology-Related Parameters for Redundant Topologies, page 10-11](#))
 - Step 5** Make sure that the SCE platforms have synchronized and active SCE platform was selected. Use the **show interface linecard 0 connection-mode** command.
 - Step 6** If you want to start with bypass/sniffing, change the link mode to your required mode in both SCE platforms on both links. The bypass mode will be applied only to the active SCE platform. (See [About the Link Mode, page 7-4.](#))
 - Step 7** Make sure that the link mode is as you required. (See [How to Monitor the System, page 10-12.](#)) Use the **show interface linecard 0 link mode** command.
 - Step 8** Connect the traffic port of SCE platform #1. This will cause a momentary down time until the network elements from both sides of the SCE platform auto-negotiate with it and start working (when working inline).
 - Step 9** Connect the traffic port of SCE platform #2, this will cause a momentary down time until the network elements from both sides of the SCE platform auto-negotiate with it and start working (when working inline).
 - Step 10** When full control is needed, change the link mode on both SCE platforms on both links to 'forwarding'. It is recommended to first configure the active SCE platform and then the standby. (See [About the Link Mode, page 7-4.](#))
 - Step 11** You can now start working with the Subscriber Manager.
-

Recovery

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 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 SCE platform.

Replacing the 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 10-9
- [Automatic steps \(in parallel with the manual steps, requires no user intervention\)](#); page 10-10

Manual steps:

-
- | | |
|---------------|--|
| Step 1 | Disconnect the failed SCE platform from the network |
| Step 2 | Connect a new SCE platform to the management link and the cascade links (leave network ports disconnected.) |
| Step 3 | Configure the SCE platform. |
| Step 4 | Basic network configurations done manually (first time). |
| Step 5 | Load application software (<i>Service Control Application for Broadband</i>) to the 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-SCE platform communication. |
| Step 2 | Synchronization with the SM. |
| Step 3 | Copying updated subscriber states from the active SCE platform to the standby. |
-

Reboot only (fully automatic recovery)

-
- | | |
|---------------|--|
| Step 1 | Reboot of the SCE platform. |
| Step 2 | Basic network configurations. |
| Step 3 | Establishment of inter-SCE platform communication. |
| Step 4 | Selection of the active SCE platform. |
| Step 5 | Synchronization of the recovered SCE platform with the SM. |
| Step 6 | Copying updated subscriber states from the active SCE platform to the standby. |
-

CLI Commands for Cascaded Systems

This section presents CLI commands relevant to the configuration and monitoring of a redundant system.

Use the following commands to configure and monitor a redundant system:

- connection-mode
- [no] force failure-condition
- show interface linecard 'number' connection-mode
- show interface linecard 'number' physically-connected links

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 SCE platforms. Therefore the connection mode for both SCE platforms may be either:
 - Inline-cascade
 - Receive-only-cascade
- **Physically-connected-links** — For each of the cascaded SCE platforms, this parameter defines the number of the link (Link 0 or Link 1) connected to this SCE platform.
- **Priority** — For each of the cascaded SCE platforms, this parameter defines whether it is the primary or secondary device.
- **On-failure** — For each of the cascaded SCE platforms, this parameter determines whether the system cuts the traffic or bypasses it when the 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
- physically connected links
- behavior upon failure of the SCE platform
- primary/secondary

To configure the connection mode, use the following command.

Step 1 From the SCE 2000 (config if)# prompt, type **connection-mode inline-cascade|receive-only-cascade [physically-connected-links {link-0|link-1}][priority {primary|secondary}] [on-failure {bypass|cutoff}]** and press Enter.

Examples

EXAMPLE 1

Use the following command to configure the primary SCE platform in a two-SCE platform inline topology. Link 1 is connected to this SCE platform and the behavior of the SCE platform if a failure occurs is *bypass*.

```
SCE 2000(config if)#connection-mode inline-cascade physically-connected-links link-1
priority primary on-failure bypass
```

EXAMPLE 2

Use the following command to configure the SCE platform that might be cascaded with the SCE platform in Example 1. This SCE platform would have to be the secondary SCE platform, and Link 0 would be connected to this 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 SCE platform if a failure occurs is also bypass.

```
SCE 2000(config if)# connection-mode inline-cascade physically-connected-links link-0
priority secondary on-failure bypass
```

How to Monitor the System

Use the following commands to view the current connection mode and link mode parameters.

- [How to View the Current Connection Mode, page 10-12](#)
- [How to View the Current Link Mode, page 10-12](#)
- [How to View Current Link Mappings, page 10-12](#)

How to View the Current Connection Mode

Step 1 From the SCE 2000# prompt, type `show interface linecard 0 connection-mode` and press Enter.

How to View the Current Link Mode

Step 1 From the SCE 2000# prompt, type `show interface linecard 0 link mode` and press Enter.

How to View Current Link Mappings

Step 1 From the SCE 2000# prompt, type `show interface linecard 0 physically-connected-links` and press Enter.

System Upgrades

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 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 SCE platforms (open the package and copy configuration).
 - Step 2** Reload the standby 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 SCE platform.
 - Step 6** After the former active SCE platform reboots and is ready to work manually, it may be left as standby or we can manually switch the SCE platforms to their original state.
-

Application Upgrade

-
- Step 1** Unload the application in the standby SCE platform.
 - Step 2** Load new application to the standby SCE platform.
 - Step 3** Make sure that the connection mode configurations are correct.
 - Step 4** Wait until the standby SCE platform finishes synchronizing and is ready to work.
 - Step 5** Force failure condition in the active SCE platform.
 - Step 6** Upgrade the application in the former active SCE platform.
 - Step 7** Remove the force failure condition in that platform.
 - Step 8** After the former active 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 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 SCE platform.
This makes the updated SCE platform the active one, and it begins to give the NEW service.
- Step 3** Repeat step 1 for the (now) standby SCE platform.
Since this includes a reboot, it is not necessary to undo the force failure command.
-



CHAPTER 11

Identifying and Preventing Distributed-Denial-Of-Service Attacks

This module describes the ability of the SCE platform to identify and prevent DDoS attacks, and the various procedures for configuring and monitoring the Attack Filter Module.

- [Information About Attack Filtering and Attack Detection, page 11-1](#)
- [How to Configure Attack Detectors, page 11-7](#)
- [How to Configure Subscriber Notifications, page 11-17](#)
- [Preventing and Forcing Attack Detection, page 11-19](#)
- [How to Monitor Attack Filtering, page 11-21](#)

Information About Attack Filtering and Attack Detection

- [Attack Filtering, page 11-2](#)
- [Specific Attack Filtering, page 11-2](#)
- [Attack Detection, page 11-3](#)
- [Attack Detection Thresholds, page 11-4](#)
- [Attack Handling, page 11-5](#)
- [Hardware Filtering, page 11-6](#)

Attack Filtering

The 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 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, page 11-13.](#))

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 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 than 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 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 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 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 SCE platform has two ways of handling an attack: by software or by hardware. Normally, attacks are handled by software. This enables the 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 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 [How to Monitor Attack Filtering, page 11-21](#)):

- 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.

How to Configure Attack Detectors

- [About Configuring Attack Detectors, page 11-7](#)
- [How to Enable Specific-IP Detection, page 11-9](#)
- [How to Configure the Default Attack Detector, page 11-10](#)
- [Specific Attack Detectors, page 11-13](#)
- [Sample Attack Detector Configuration, page 11-17](#)

About Configuring Attack Detectors

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 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 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 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 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 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 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 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 *numbers* suspected-flows-rate *numbers* 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*

How to Enable Specific-IP Detection

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

-
- Step 1** From the SCE(config if)# prompt, type `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)]` and press Enter.
-

How to Enable Specific-IP Detection for the TCP Protocol Only for all Attack Directions

-
- Step 1** From the SCE(config if)# prompt, type `attack-filter protocol TCP` and press Enter.
-

How to Enable Specific-IP Detection for the TCP Protocol for Port-based Detections Only for Dual-sided Attacks

-
- Step 1** From the SCE(config if)# prompt, type `attack-filter protocol TCP dest-port specific attack-direction dual-sided` and press Enter.
-

How to Disable Specific-IP Detection for Protocols Other than TCP, UDP, and ICMP for all Attack Directions

-
- Step 1** From the SCE(config if)# prompt, type `no attack-filter protocol other` and press Enter.
-

How to Disable Specific-IP Detection for ICMP for Single-sided Attacks Defined by the Source IP

-
- Step 1** From the SCE(config if)# prompt, type `no attack-filter protocol ICMP attack-direction single-side-source` and press Enter.
-

How to Configure the Default Attack Detector

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 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 `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) [action (report|block)] [open-flows-rate numbersuspected-flows-rate ratesuspected-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.

Step 1 From the SCE(config if)# prompt, type `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)` and press Enter.

Reinstates the system defaults for the defined attack types.

How to Reinstate the System Defaults for All Attack Types

Step 1 From the SCE(config if)# prompt, type `default attack-detector default` and press Enter.

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 11-13
- [How to Enable a Specific Attack Detector and Assign it an ACL](#), page 11-14
- [How to Define the Action and Optionally the Thresholds for a Specific Attack Detector](#), page 11-14
- [How to Define the Subscriber Notification Setting for a Specific Attack Detector](#), page 11-15
- [How to Define the SNMP Trap Setting for a Specific Attack Detector](#), page 11-15
- [How to Define the List of Destination Ports for TCP or UDP Protocols for a Specific Attack Detector](#), page 11-15
- [How to Delete User-Defined Values](#), page 11-16
- [How to Disable a Specific Attack Detector](#), page 11-16
- [How to Disable All Non-default Attack Detectors](#), page 11-16
- [How to Disable All Attack Detectors](#), page 11-16

Options

A specific attack detector may be configured for each possible combination of protocol, attack direction, and side. The 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 Access Control Lists.)
 - 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 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

-
- Step 1** From the SCE(config if)# prompt, type `attack-detector numberaccess-list (aclnumber|none)`
`[comment comment]` and press Enter.

Enables the attack detector and assigns it the specified ACL.

How to Define the Action and Optionally the Thresholds for a Specific Attack Detector

-
- Step 1** From the SCE(config if)# prompt, type `attack-detector numberprotocol ((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 numbersuspected-flows-rate ratesuspected-flows-ratio ratio]` and press Enter.

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.

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

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

-
- Step 1** From the SCE(config if)# prompt, type `attack-detector number TCP-port-list|UDP-port-list (all| (port1 [port2, port3...])` and press Enter.
- 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.

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

Defines the action of the specified attack detector.

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.

-
- Step 1** From the SCE(config if)# prompt, type `default attack-detector number` and press Enter.

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.

-
- Step 1** From the SCE(config if)# prompt, type `default attack-detector all-numbered` and press Enter.

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.

-
- Step 1** From the SCE(config if)# prompt, type `default attack-detector all` and press Enter.

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 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 UDP-ports-list 53 comment "DNS servers"` and press Enter.
Enables attack detector #1, assigns ACL #3 to it, and defines the list of UDP destination ports with one port, port 53.
- Step 4** From the SCE(config if)# prompt, type `attack-detector 1 protocol UDP dest-port specific attack-direction single-side-destination action report open-flow-rate 1000000 suspected-flows-rate 1000000` and press Enter.
Defines the thresholds and action for attack detector #1
- Step 5** 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 6** From the SCE(config if)# prompt, type `exit` and press Enter.
Exits the linecard interface configuration mode.
- Step 7** 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`
-

How to Configure Subscriber Notifications

- [About Subscriber Notifications, page 11-18](#)
- [How to Configure the Subscriber Notification Port, page 11-18](#)
- [How to Remove the Subscriber Notification Port, page 11-18](#)

About Subscriber Notifications

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 SCE platform, as explained in the appropriate Service Control Application user guide.

In the current solutions, the 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.

How to Configure 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 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

Step 1 From the SCE(config if)# prompt, type `attack-filter subscriber-notification ports portnumber` and press Enter.

How to Remove the Subscriber Notification Port

Step 1 From the SCE(config if)# prompt, type `no attack-filter subscriber-notification ports` and press Enter.

Preventing and Forcing Attack Detection

- [Options, page 11-19](#)
- [Preventing Attack Filtering, page 11-20](#)
- [Forcing Attack Filtering, page 11-20](#)

After configuring the attack detectors, the 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 SCE platform attack-detectors properly. For example:

- The 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 **don't-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 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 **don't-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 don't-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 `don't-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 don't-filter`).

- [How to Configure a don't-filter Setting for a Specified Situation, page 11-20](#)
- [How to Remove a don't-filter Setting from a Specified Situation, page 11-20](#)
- [How to Remove All don't-filter Settings, page 11-20](#)

How to Configure a don't-filter Setting for a Specified Situation

Step 1 From the SCE(config if)# prompt, type `attack-filter don't-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)` and press Enter.

How to Remove a don't-filter Setting from a Specified Situation

Step 1 From the SCE(config if)# prompt, type `no attack-filter don't-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)` and press Enter.

How to Remove All don't-filter Settings

Step 1 From the SCE(config if)# prompt, type `no attack-filter don't-filter all` and press Enter.

Forcing Attack Filtering

Attack filtering can be forced for a specified IP address/protocol. If filtering is already in process, it will be stopped. Forced attack filtering will continue until undone by an explicit CLI command (either `no force-filter` or `don't-filter`).

- [How to Configure a force-filter Setting for a Specified Situation, page 11-21](#)
- [How to Remove a force-filter Setting from a Specified Situation, page 11-21](#)
- [How to Remove All force-filter Settings, page 11-21](#)

How to Configure a force-filter Setting for a Specified Situation

Step 1 From the SCE(config if)# prompt, type `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) [notify-subscriber]` and press Enter.

How to Remove a force-filter Setting from a Specified Situation

Step 1 From the SCE(config if)# prompt, type `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)` and press Enter.

How to Remove All force-filter Settings

Step 1 From the SCE(config if)# prompt, type `no attack-filter force-filter all` and press Enter.

How to Monitor Attack Filtering

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

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 11-24](#)
- [Example, page 11-25](#)

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

Step 1 From the SCE>prompt, type `show interface linecard 0 attack-detector number` and press Enter.

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 rate | Ddos-Suspected flows rate | flows 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

Step 1 From the SCE>prompt, type `show interface linecard 0 attack-detector default` and press Enter.

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

Step 1 From the SCE>prompt, type `show interface linecard 0 attack-detector all` and press Enter.

How to display filter state (enabled or disabled)

Step 1 From the SCE>prompt, type `show interface linecard 0 attack-filter` and press Enter.

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.

Step 1 From the SCE>prompt, type `show interface linecard 0 attack-filter query ((single-sided ip ip-address)|(dual-sided source-IP source-ip-addressdestination-IP dest-ip-address))[dest-port portnumber] configured` and press Enter.

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
|      |      |      |      |      |      |      |      |      |      |      |
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) |      |      |      |      |      |      |      |      |
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) |      |      |      |      |      |      |      |      |
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.

-
- Step 1** From the SCE>prompt, type `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` and press Enter.
-

How to display all currently handled attacks

-
- Step 1** From the SCE>prompt, type `show interface linecard 0 attack-filter current-attacks` and press Enter.
-

How to display all existing force-filter settings

-
- Step 1** From the SCE>prompt, type `show interface linecard 0 attack-filter force-filter` and press Enter.
-

How to display all existing don't-filter settings

-
- Step 1** From the SCE>prompt, type `show interface linecard 0 attack-filter don't-filter` and press Enter.
-

How to display the list of ports selected for subscriber notification

-
- Step 1** From the SCE>prompt, type `show interface linecard 0 attack-filter subscriber-notification ports` and press Enter.
-

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 |      |      |
-----|-----|-----|-----|-----|-----|-----|-----
1|10.1.1.1   |Subscriber |      523 |      4045|Report|No    |No
|           |*|TCP      |          |          |      |      |
-----|-----|-----|-----|-----|-----|-----|-----

```

Viewing the Attack Log

- [The Attack Log, page 11-30](#)
- [How to View the Attack Log, page 11-31](#)
- [How to Copy the Attack Log to a File, page 11-31](#)

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

Step 1 From the SCE# prompt, type `more line-attack-log` and press Enter.

How to Copy the Attack Log to a File

Step 1 From the SCE# prompt, type `more line-attack-log redirect filename` and press Enter.
Writes the log information to the specified file.



CHAPTER 12

Value Added Services (VAS) Traffic Forwarding

This module provides an overview of VAS traffic forwarding, explaining what it is and how it works. It also explains the various procedures for configuring and monitoring the VAS traffic forwarding.

- [Information About VAS Traffic Forwarding, page 12-1](#)
- [How VAS Traffic Forwarding Works, page 12-2](#)
- [VAS Redundancy, page 12-9](#)
- [VAS Status and VAS Health Check, page 12-11](#)
- [VAS Traffic Forwarding Topologies, page 12-12](#)
- [SNMP Support for VAS, page 12-15](#)
- [Interactions Between VAS Traffic Forwarding and Other SCE Platform Features, page 12-15](#)
- [Configuring VAS Traffic Forwarding, page 12-16](#)
- [How to Monitor VAS Traffic Forwarding, page 12-27](#)
- [VAS over 10G, page 12-31](#)

Information About VAS Traffic Forwarding

This module provides an overview of VAS traffic forwarding, and explains how to configure and monitor VAS traffic forwarding. It also explains how to configure VAS over 10G installations.

With every new SCA BB release, the classification and control of new services is supported. The VAS integration capability enables classification and control of services not currently supported by SCA BB. The concept behind this capability enables the solution to use an external “expert system” for classification and control of the service traffic. Using this capability, the service provider can choose to forward selected flows to an external, third-party solution for per-subscriber value-added functionality. Possible use cases for this functionality would be intrusion detection and content-filtering. These value added services are provided on top of the services and functions of the SCA BB solution.

The VAS feature enables the user to divert a specified part of the traffic streams to an individual VAS server or appliance, or a cluster of them. The diversion is based on the subscriber package, flow type and the availability of the VAS servers. This capability also delivers a load balancing function for even distribution of the load on the various VAS servers.

The solution provides support for multiple VAS service types using different VAS Server Groups. Several servers of the same type can be deployed to increase the total capacity and resiliency.

The 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.

There is also a VAS over 10G solution, which is a special case of the Cisco Multi-Gigabit Service Control Platform (MGSCP) solution, supporting only one external 10G link and using a Cisco 6500/7600 Series router as a dispatcher to distribute the external 10G link and as the switch towards the VAS servers

VAS Service Goals

The VAS traffic forwarding functionality allows the Service Control solution to meet several important service goals:

- Allows service providers to provide a range of Value Added Services to their subscribers, thus increasing customer satisfaction.
- Allows the SCE platform to forward part of the traffic to third party devices that can provide additional, complementary services.

The 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
- Based on the policy that was configured
- Allows the Service Control solution to include Value Added Servers that cannot be deployed inline for various reasons (for example, cannot support throughput or are not carrier grade for inline insertion).
- Provides easy interoperability and flexibility for setting different services.

Since the VAS feature emulates a regular IP network for the third party devices, no special support is required on their part.

How VAS Traffic Forwarding Works

- [Requirements for VAS Servers, page 12-4](#)
- [VAS Traffic Forwarding and SCA BB, page 12-5](#)
- [VLAN Tags for VAS Traffic Forwarding, page 12-5](#)
- [Service Flow, page 12-6](#)
- [Data Flow, page 12-6](#)
- [Load Balancing, page 12-8](#)

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, 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 were classified to a VAS service get the usual SCA BB service, as well as being forwarded to the VAS servers for additional service.

Logically, “VAS engine” is upstream of “SCA engine”, meaning upstream traffic is first processed by the SCA BB application, downstream is first processed by the VAS server.

Routing the traffic to the VAS servers is done by using VLAN tags.

Figure 12-1 Typical VAS Traffic Forwarding Installation



Important information:

- A single SCE platform can support up to eight VAS servers.
- A maximum of 512 SCE platforms can be connected.
- The same VAS server may be used by more than one SCE platform.
- The VAS traffic forwarding feature is supported on the SCE 2000 4xGBE platform only.

The following sections provide a more detailed description of how VAS traffic forwarding works.

Requirements for VAS Servers

Since the VAS devices are installed behind the SCE platform, they should follow the network behavior of the SCE platform.

The two main guidelines that VAS device should follow are:

- They should be equipped with two separate interfaces – subscriber side and network side.
Traffic toward the subscribers should be sent from the subscriber interface and for the Internet from the network interface.
- Transparency in layer-2 – This means that the VAS servers act like layer-2 switches from the networking perspective. They are not allowed to change anything in the traffic headers or to generate new traffic on behalf of themselves

Layer-2 Transparency

Note the following guidelines for handling VAS services non-management traffic:

- 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 (injecting), the VAS devices must preserve the original layer-2 headers containing the MAC addresses and the VLAN tag. The VAS devices may not change the MAC addresses (destination or source) or the VLAN tags.
 - New traffic can be injected in the context of existing flow only.
 - The VAS device is not permitted to initiate new flows.
 - When injecting traffic, the layer-2 information (MAC addresses, VLAN tags, and the IP / TCP parameters) must be taken from the context of the flow into which the traffic is being injected.
- VAS devices may not generate any network transactions on behalf of themselves or replay 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 management traffic will be sent directly to the router and not through the 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 since the classification is used to select the routing of the packet to a VAS server or to the subscriber/network.

VAS traffic forwarding rules are configured through 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 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 GUI, this definition is given per package

VLAN Tags for VAS Traffic Forwarding


The traffic is routed between the SCE platform and the VAS servers by VLANs. There is a unique VLAN tag for each SCE platform/VAS server combination.

Before being forwarded to the VAS servers, the SCE platform adds the VLAN tag to the original traffic. When the traffic returns to the SCE platform, the SCE platform removes the VLAN tag it previously added, and then forwards the traffic on its original link.

The VLAN tag to be used for each VAS server is configured by the user. To preserve consistency of the traffic flow, the VAS solution requires the user to configure a unique VLAN tag for each SCE platform/VAS server combination.

The VLAN tag format is shown in the figure below.

Figure 12-2 VLAN Tag Format



The VLAN tag has twelve bits, divided as follows:

- The VLAN tag has twelve bits, divided as follows:
 - The lower three bits identify the VAS server.
 - The higher nine bits identify the SCE platform.

- For example:

0x20 = 100 000 = SCE #4, VAS #0

0x21 = 100 001 = SCE #4, VAS #1

0x58 = 1101 000 = SCE #13, VAS #0

Observe the following for the nine bits that identify the SCE platform:

- These nine bits must be the same for all VAS servers attached to a specific SCE platform.
- These nine bits must be different for VAS servers attached to different SCE platforms.

Examples of valid VLAN tag ranges for an SCE platform:

0x20, 0x21 – 0x27, but not 0x33

0x58, 0x59 – 0x5F, but not 0x26

The SCE platform enforces that the VLAN tags configured by the user keep this format, meaning that the lower three bits match the VAS server number for which the VLAN tag is configured and that the higher nine bits match the higher nine bits previously configured for other VAS servers on this SCE platform. However, the SCE platform is not aware of the configuration of other SCE platforms, and therefore it is the responsibility of the user to configure a unique nine bits (SCE id) for each SCE platform.

Note that the use of VLAN tags is an integral part of the VAS solution, and therefore requires the VAS device to be able to work in 802.1q trunk while preserving the VLAN information.

Service Flow

The mapping of traffic portions to VAS Server Groups is done through the standard SCA GUI, this definition is given per package.

The 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 SCE platform performs load sharing between multiple VAS servers belonging to the same Server Group; the balance is based on subscriber load. In other words, the SCE platform ensures that the subscribers are evenly distributed between the VAS servers in the same group. Note that the mapping of subscriber to a VAS server (per group) is kept 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 will change only if the same server changes its status.

The following sections explain in more detail when and how the mapping is changed.

Data Flow

In a deployment using VAS traffic forwarding, there are two types of data flows:

- Non-VAS flow
- VAS flow

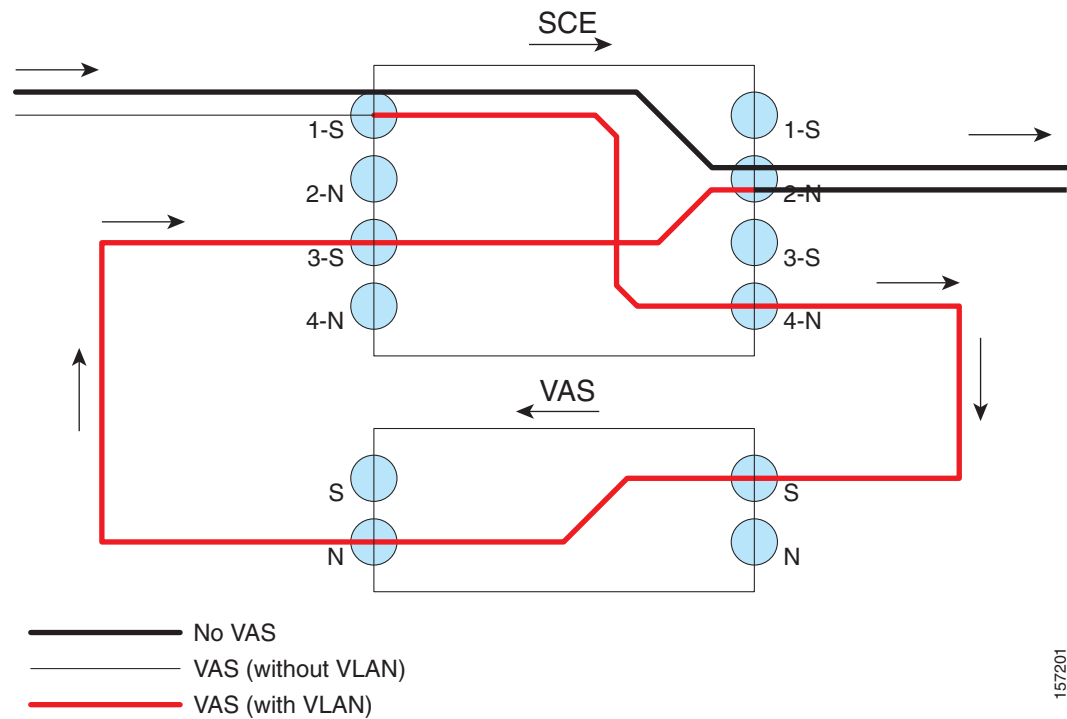
The following figure depicts the two types of data flows running through a single SCE platform and a single VAS server.

- Ports are illustrated as two uni-directional half ports, RX (on the left side) and the TX (on the right side).
 - The SCE platform has four ports.
 - The VAS server has two ports.
- For the sake of illustration, the SCE platform traffic flow direction is left to right while the VAS traffic flow is 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

This figure illustrates the data flow from the subscriber to the network. Data flow from the network to the subscriber works in exactly the same way, but is received on the network port (N) and transmitted on the subscriber port (S).

Figure 12-3 Data Flow in a VAS System



Non-VAS Data Flow

The flow steps for a non-VAS flow are:

- A subscriber packet is received at the SCE platform port 1 (S).
- The SCE platform classifies the flow as non-VAS flow.
- The packet is sent to the network on port 2 (N).

VAS Data Flow

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 SCE platform flow, but before it is transmitted to its original destination, it flows through the VAS server.

The flow steps are:

- A subscriber packet is received at the SCE platform port 1 (S).
- The SCE platform classifies the flow as VAS flow.
- The SCE platform adds a VLAN tag to the packet.
- The VLAN tag is used by the Ethernet switch to route the packet to the proper VAS server.

Note that at this point the packet has a VLAN tag, which is indicated by the red line.

- The packet is sent to the VAS subscriber port from SCE platform port 4 (N).
- The VAS server processes the packets and either drops the packet or sends it back to the SCE platform; from the VAS network port to the SCE platform subscribers port 3 (S).

Note that 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 SCE platform.

- The SCE platform receives the packet on port 3 (S), drops the VLAN tag, and passes the packet towards the network through port 2 (N).

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. Let us say that the system requires five VAS servers, three to provide FTP caching, and two to provide virus filtering. Defining two VAS Server Groups, FTP caching and virus filtering, permits load sharing across the servers for each Server Group.

The VAS Server Group to which the flow should be attached is determined by the package of the subscriber. The selection of a specific VAS server from the VAS servers within this 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 VAS server may be used by more than one SCE platform. Remember that the SCE platform performs load balancing only on the traffic that it sends to the VAS server; it is not aware how much of a load the VAS server may be bearing from a different SCE platform. It is the responsibility of the user to allocate available VAS servers to the SCE platforms in a way that ensures proper total 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, meaning that 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 the server can make subscriber based decisions.

The 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 will only apply on new flows. Flows that were already mapped to a server before it became active will remain attached to it.

The mapping of subscriber to VAS servers not is saved across logouts or SCE platform reload

Load Balancing and Subscriber Mode

Since the load balancing is subscribers based, this solution will not work properly in subscriberless mode, as the entire traffic load would be carried only by one VAS server per group.

**Note**

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 will be processed as defined by the default template. Therefore, the default template should provide a proper package, so these flows will get VAS service.

VAS Redundancy

- [VAS Server Failure, page 12-10](#)
- [VAS Server Group Failure, page 12-10](#)
- [Ethernet Switch Failure, page 12-10](#)
- [Disabling a VAS Server, page 12-11](#)

The services provided by the VAS servers should be highly available. The failure of a single VAS server should not degrade the 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 by which the system guarantees the performance and availability of the VAS services:

- Load sharing — The SCE platform distributes the subscribers between all the active VAS servers within a server group.
- Monitoring — The 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.

VAS Server Failure

The system monitors the health of a VAS server by periodically checking the connectivity between the SCE platform and the VAS server. When the SCE platform fails to establish or maintain a connection to the server within a configurable window of time, the state of the server is considered to be **Down**.

The implications of such state are:

- New logged-in subscribers will be distributed between the other active servers in the group.
- Subscribers that are mapped to this server will be mapped to a new server if they initiate a new flow.
- The Server Group may move to a Failure state if this failure caused the number of active servers in the group go below the minimum 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 will continue 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, the user can configure the following:

- The minimum number of active servers necessary.
- The action to take in case the actual number of active servers goes below this number.

Note that if the minimum number equals the total number of configured servers, it means there is no redundancy and failure of one server will cause the failure of the whole server group.

When the 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 will then be applied to all new flows mapped for that VAS Server Group (existing flows will not be affected.)

There are two possible actions when the VAS Server Group has failed:

- **Block** — all new flows assigned to the failed VAS Server Group will be blocked by the 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 (that is, they will get 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 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 the 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) will be 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 will bring down the VAS Server Group since a disabled VAS server is equal to a VAS server in a Down state.

Health check is not performed on disabled VAS servers.

VAS Status and VAS Health Check

To manage the VAS redundancy, the SCE platform needs to know the state of each VAS server. The SCE platform performs periodic health checks for all the configured VAS servers. These checks are the basis for VAS redundancy control; they enable the SCE platform to identify and react to VAS server failure, and to check the connectivity of the SCE platform-VAS server before enabling the server to handle traffic.

The health check is performed over the VAS link, the link that connects the SCE platform with the VAS servers. It validates the traffic flow between the SCE platform and the VAS server in both directions through special health check packets generated by the SCE platform.

The health check mechanism does not require special interaction with the VAS device, since the VAS server does not need to answer the health check packets, only to pass them as they are back to the SCE platform. As long as the packets are received by the SCE platform, the VAS server is considered to be alive. Failing to receive the packets back from the VAS server within a pre-defined time window is considered by the SCE platform as a failure of the VAS server and the server status is changed to **Down**.

Important information about the health check packets:

- They are carried over UDP flows
- Their source and destination IP addresses are configurable by the user

IP addresses should be:

- Unique to the SCE platform
- Addresses that will not be used by the network traffic (such as private IPs)

The SCE platform uses default UDP ports between 63140 and 63155, unless the user has configured different ports for the health check.

The SCE platform adds its own layer 7 data on top of the UDP transport layer. This data is used by the 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 VLAN tag
- Pseudo IPs are configured for the GBE 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).

To check the connectivity with the VAS server before enabling it to handle traffic, the server should not be assigned to any 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 does not drop traffic unless it is specifically configured to do so. Therefore, if the VAS server-SCE platform connectivity is operative, the health check packets should reach the SCE platform safely.

Alternatively, it should be able 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, not bypass, the traffic (cut the link), so that the SCE platform will be able to identify the failure.

VAS Server States

When determining whether a VAS server is active, the system considers the following two parameters:

- Admin mode as configured by the user — 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 SCE Platform, Multiple VAS Servers, page 12-12](#)
- [Multiple SCE Platforms, Multiple VAS Servers, page 12-14](#)
- [VAS over 10G, page 12-31](#), which is a special case of Cisco Multi-Gigabit Service Control Platform (MGSCP) solution, supporting only one external 10G link and using a Cisco 6500/7600 Series router as a dispatcher to distribute the external 10G link and as the switch towards the VAS servers.



Note

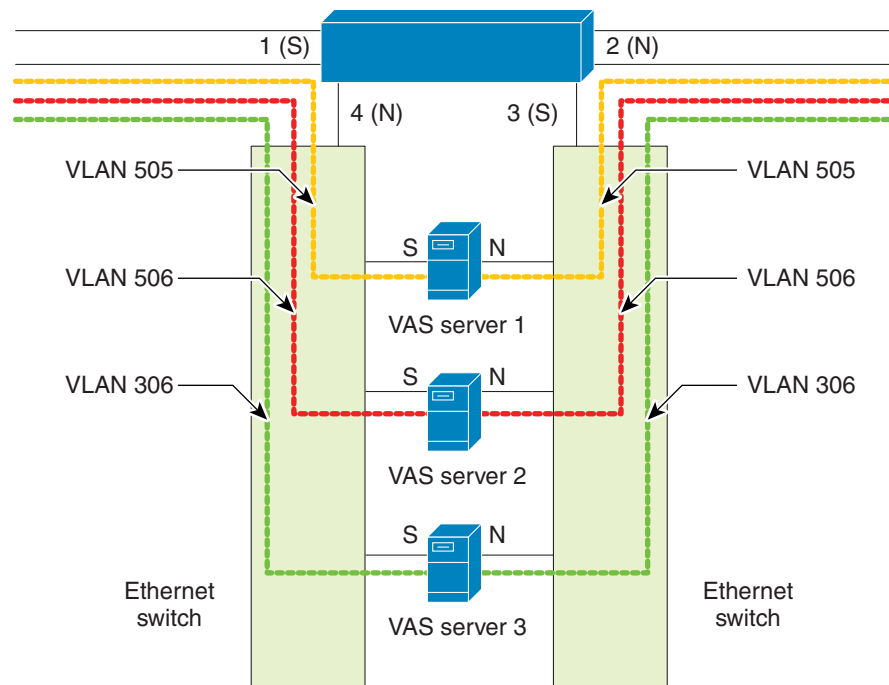
A topology in which a VAS server is directly connected to the SCE platform is not supported. If a topology of a single SCE platform connected to a single VAS server is desired, a switch should still be used between the SCE platform and the VAS server.

Single SCE Platform, Multiple VAS Servers

In this topology, a single SCE platform is forwarding VAS traffic to one or more VAS servers through two Ethernet switches.

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 12-4 Single SCE Platform, Multiple VAS Servers



157199

Data Flow

The data flow is:

- A subscriber packet is received at port #1 (Subscriber).
- The SCE platform opens a flow and classifies the flow as either a non-VAS (blue) flow or as a VAS flow (red).
- If the flow is non-VAS (blue), the SCE platform passes the packet to the network. The VAS server is not involved in this case.
- If the flow is a VAS flow (red), the 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).
- 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).
- The VAS server processes the packet and either drops or forwards it without changing the VLAN tag.
- The packet is forwarded by the Ethernet switch to the SCE platform according to its VLAN tag (the port towards the SCE platform should be the only port with this VLAN tag allowed).
- The 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 SCE Platforms, Multiple VAS Servers

In this topology, multiple SCE platforms are connected to multiple VAS servers. Note that at least one VAS server receives traffic from more than one SCE platform; if the VAS servers are each in an exclusive relationship to a particular SCE platform, it would simply be several single SCE platform/multiple VAS server topologies grouped together.


In the following figure, the top SCE platform forwards traffic to VAS servers 1 and 2, while the bottom SCE platform forwards to VAS servers 2 and 3. A unique VLAN tag must designate each SCE-platform-to-VAS-server path. This topology is illustrated with two SCE platforms, but a maximum of 512 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 previous topology.

Note that SCE platform redundancy on the cascade ports is not supported in this topology.

Figure 12-5 Multiple SCE Platforms, Multiple VAS Servers



SNMP Support for VAS

The following items in the “PCUBE-SE-MIB” proprietary MIB support VAS traffic forwarding:

- SCE-MIB object: `vasTrafficForwardingGrp` 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 SCE Platform Features

- [Incompatible SCE Platform Features, page 12-15](#)
- [VAS Traffic Forwarding and DDoS Processing, page 12-15](#)
- [VAS Traffic Forwarding and Bandwidth Management, page 12-16](#)

Incompatible SCE Platform Features

There are certain 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.

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

VAS Traffic Forwarding and DDoS Processing

VAS traffic forwarding has minor effects on the DDoS mechanisms.

- [Specific IP DDoS Attack Detection, page 12-15](#)
- [Specific IP Attack filter, page 12-16](#)

Specific IP DDoS Attack Detection

The specific IP DDoS mechanism uses software counters. The second pass VAS packets do not reach the software, 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 also are 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 to give.
- Bypass – Traffic will be bypassed and NO SCA BB or VAS services are given.

VAS Traffic Forwarding and Bandwidth Management

The complexity of the VAS traffic forwarding results in the modification of some SCE platform bandwidth management capabilities when using this feature:

- VAS flows are not subject to global bandwidth control.
- The number of global controllers available to regular flows has been decreased from 64 to 48.

Certain changes are required in the configuration of the global controllers to support these two restrictions.

Global Controllers and VAS flows

When VAS traffic forwarding is enabled, the global controllers function slightly differently.

- Only 48 global controllers are available to the user.
- Global controllers 49-63 are used to count VAS traffic.
- The reserved global controllers cannot be configured.
- On VAS flows, the flow does not get its global controller from the traffic controller to which it belongs. Rather, its global controller is set according to VAS rules.

Configuring VAS Traffic Forwarding

There are three broad aspects to VAS traffic forwarding configuration in the 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 Traffic Forwarding Settings* 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 SCE platform— define the servers and the server groups, configure Pseudo IP for the GBE interfaces, 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 [How to Monitor VAS Traffic Forwarding, page 12-27](#)).
3. Configure which traffic goes to which Server Group through the SCA BB console (see [Configuring VAS Traffic Forwarding from the SCA BB Console, page 12-17](#)).

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 SCE platform CLI:

- 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, the user must enable it.

For instructions on how to disable VAS traffic forwarding, see [Disabling VAS Traffic Forwarding, page 12-18](#).

There are certain other 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.

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

Step 1 From the SCE(config if)# prompt, type `VAS-traffic-forwarding` and press Enter.

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 SCE platform while the applied SCA BB policy instructs the 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 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 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 `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.

How to Configure the VAS Traffic Link

- [Options, page 12-19](#)
- [How to Select the Link for VAS Traffic, page 12-19](#)
- [How to Revert to the Default Link for VAS Traffic, page 12-19](#)

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.

To configure the link for VAS over 10G, see [VAS over 10G, page 12-31](#)).

**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

-
- Step 1** From the SCE(config if)# prompt, type `vas-traffic-forwarding traffic-link {link-0|link-1}` and press Enter.
-

How to Revert to the Default Link for VAS Traffic

-
- Step 1** From the SCE(config if)# prompt, type `no vas-traffic-forwarding traffic-link` and press Enter.
-

How to Configure a VAS Server

The user 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 following commands 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 that 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

Use this command to enable a VAS server.



Note

The server is not operational until a VLAN tag has also been defined.

- Step 1** From the SCE(config if)# prompt, type `vas-traffic-forwarding VAS server-id number enable` and press Enter.
-

How to Disable a VAS Server

- Step 1** From the SCE(config if)# prompt, type `vas-traffic-forwarding VAS server-id number disable` and press Enter.
-

How to Restore all VAS Server Properties to Default

- Step 1** From the SCE(config if)# prompt, type `no vas-traffic-forwarding VAS server-id number` and press Enter.
-

How to Assign a VLAN ID to a VAS Server

Use this command to assign the VLAN ID to a specified VAS server.

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

-
- Step 1** From the SCE(config if)# prompt, type `VAS-traffic-forwarding VAS server-id number VLAN vlan-id` and press Enter.
-

How to Remove the VLAN Tag Number from a Specified VAS Server

-
- Step 1** From the SCE(config if)# prompt, type `no VAS-traffic-forwarding VAS server-id number VLAN` and press Enter.

You can also use the default form of the command to remove the VLAN tag configuration.

default VAS-traffic-forwarding VAS server-id *number*VLAN

How to Configure the Health Check

- [About the Health Check, page 12-22](#)
- [Options, page 12-22](#)
- [How to Disable VAS Server Health Check, page 12-22](#)
- [How to Define the UDP Ports to be Used for Health Check, page 12-22](#)
- [How to Remove the UDP Ports Configuration, page 12-23](#)

About the Health Check

Use these commands to enable and disable the Health Check, and to define the ports it should use.

By default, the VAS server health check is enabled, however the user may disable it.

Note that 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 SCE platform GBE ports on the VAS traffic link
- VAS server is enabled
- Server has a VLAN tag
- Health check for the server is enabled

To configure VAS server health check for VAS over 10G, see also [How to Configure Health Check for VAS over 10G, page 12-42](#).

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)

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 Disable VAS Server Health Check

-
- Step 1** From the SCE(config if)# prompt, type `no VAS-traffic-forwarding VAS server-id number health-check` and press Enter.
-

How to Define the UDP Ports to be Used for Health Check

-
- Step 1** From the SCE(config if)# prompt, type `VAS-traffic-forwarding VAS server-id number health-check UDP ports source portnumber destination portnumber` and press Enter.
-

How to Remove the UDP Ports Configuration

Step 1 From the SCE(config if)# prompt, type `no VAS-traffic-forwarding VAS server-id number health-check UDP ports` and press Enter.

You can also use the default form of the command to remove the UDP port configuration.

VAS-traffic-forwarding VAS server-id number health-check UDP ports

How to Configure Pseudo IP Addresses for the Health Check Packets

- [About Pseudo IP Addresses, page 12-23](#)
- [Options, page 12-23](#)
- [How to Define the pseudo IP Address, page 12-24](#)
- [How to Delete the pseudo IP Address, page 12-24](#)

About Pseudo IP Addresses

Use this command to configure source and destination pseudo IP address for the health check packets. This command allows you to specify a unique IP address to be used by the health check packets.

This is a ROOT level command and is available under the GBE configuration interface mode. The interfaces that should be configured are those interfaces which connect the SCE platform with the VAS servers (by default interfaces GBE 0/3 and GBE 0/4).

The 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 SCE platform. Note that the 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.

-
- Step 1** From the SCE(config if)#>prompt, type `pseudo-ip ip-address [mask]` and press Enter.
-

How to Delete the pseudo IP Address

-
- Step 1** From the SCE(config if)#>prompt, type `no pseudo-ip ip-address [mask]` and press Enter.
-

How to Configure a VAS Server Group

- [About VAS Server Groups, page 12-24](#)
- [How to Add and Remove Servers, page 12-25](#)
- [How to Configure VAS Server Group Failure Parameters, page 12-25](#)

About VAS Server Groups

The user 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

Use the following commands to perform these operations for a VAS server group:

- Add or remove a VAS server to or from a specified group.
- Configure the minimum number of active servers for a specified group.
- Configure failure behavior for a specified group.

How to Add and Remove Servers

Use these commands to add and remove servers to or from a specified VAS server group.

- [Options, page 12-25](#)
- [How to Add a VAS Server to a Specified VAS Server Group, page 12-25](#)
- [How to Remove a VAS Server from a Specified VAS Server Group, page 12-25](#)
- [How to Remove all VAS Servers from a Specified VAS Server Group, page 12-25](#)

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

-
- Step 1** From the SCE(config if)# prompt, type `VAS-traffic-forwarding VAS server-group group-number server-id id-number` and press Enter.
-

How to Remove a VAS Server from a Specified VAS Server Group

-
- Step 1** From the SCE(config if)# prompt, type `no VAS-traffic-forwarding VAS server-group group-number server-id id-number` and press Enter.
-

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 their default value.

-
- Step 1** From the SCE(config if)# prompt, type `no VAS-traffic-forwarding VAS server-group group-number` and press Enter.
-

How to Configure VAS Server Group Failure Parameters

- [About VAS Server Group Failure Parameters, page 12-26](#)
- [Options, page 12-26](#)
- [How to Configure the Minimum Number of Active Servers for a Specified VAS Server Group, page 12-26](#)
- [How to Reset the Minimum Number of Active Servers for a Specified VAS Server Group to the Default, page 12-26](#)
- [How to Configure the Failure Action for a Specified VAS Server Group, page 12-26](#)
- [How to Configure the Failure Action for a Specified VAS Server Group to the Default, page 12-27](#)

About VAS Server Group Failure Parameters

Use the following commands to configure these failure parameters for the 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 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

-
- Step 1** From the SCE(config if)# prompt, type `VAS-traffic-forwarding VAS server-group group-number failure minimum-active-servers min-number` and press Enter.
-

How to Reset the Minimum Number of Active Servers for a Specified VAS Server Group to the Default

-
- Step 1** From the SCE(config if)# prompt, type `default VAS-traffic-forwarding VAS server-group group-number failure minimum-active-servers min-number` and press Enter.
-

How to Configure the Failure Action for a Specified VAS Server Group

-
- Step 1** From the SCE(config if)# prompt, type `VAS-traffic-forwarding VAS server-group group-number failure action {block | pass}` and press Enter.
-

How to Configure the Failure Action for a Specified VAS Server Group to the Default

Use this command to revert the failure action configuration for the specified VAS server group to the default value (pass).

-
- Step 1** From the SCE(config if)# prompt, type `default VAS-traffic-forwarding VAS server-group group-number failure action` and press Enter.
-

How to Monitor VAS Traffic Forwarding

Use these commands 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

- Bandwidth per VAS server and VAS direction (to VAS / from VAS)
- VAS health check counters

Sample outputs are included.

How to Display Global VAS Status and Configuration

-
- Step 1** From the SCE>prompt, type `show interface linecard 0 VAS-traffic-forwarding` and press Enter.
-

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

-
- Step 1** From the SCE>prompt, type `show interface linecard 0 VAS-traffic-forwarding VAS server-group id-number` and press Enter.
-

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

-
- Step 1** From the SCE>prompt, type `show interface linecard 0 VAS-traffic-forwarding VAS server-group all` and press Enter.
-

How to Display Operational and Configuration Information for a Specific VAS Server

-
- Step 1** From the SCE>prompt, type `show interface linecard 0 VAS-traffic-forwarding VAS server-id id-number` and press Enter.
-

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

- Step 1** From the SCE>prompt, type `show interface linecard 0 VAS-traffic-forwarding VAS server-id all` and press Enter.

How to Display the VAS Servers Used by a Specified Subscriber

- Step 1** From the SCE>prompt, type `show interface linecard 0 subscriber name subscriber-name VAS-servers` and press Enter.

How to Display Health Check Counters for a Specified VAS Server

- Step 1** From the SCE>prompt, type `show interface linecard 0 VAS-traffic-forwarding VAS server-id id-number counters health-check` and press Enter.

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

-
- Step 1** From the SCE>prompt, type `show interface linecard 0 VAS-traffic-forwarding VAS server-id all counters health-check` and press Enter.
-

How to Clear the Health Check Counters for a Specified VAS Server

-
- Step 1** From the SCE>prompt, type `clear interface linecard 0 VAS-traffic-forwarding VAS server-id id-number counters health-check` and press Enter.
-

How to Clear the Health Check Counters for All VAS Servers

-
- Step 1** From the SCE>prompt, type `clear interface linecard 0 VAS-traffic-forwarding VAS server-id all counters health-check` and press Enter.
-

How to Display Bandwidth per VAS Server and VAS Direction

Note that the bandwidth presented in this command is measured at the Transmit queues, therefore the first table in the example presents the bandwidth of traffic transmitted towards the VAS servers and the second table presents the bandwidth of traffic transmitted out of the SCE platform after being handled by the VAS servers.

The counting is based on L2 bytes.

-
- Step 1** From the SCE>prompt, type `show interface linecard 0 counters VAS-traffic-bandwidth` and press Enter.
-

Example

```
SCE>show interface linecard 0 counters VAS-traffic-bandwidth
Traffic sent to VAS processing TxBW [Kbps] (bytes are counted from Layer 2):
```

| Port 1 | Port 2 | Port 3 | Port 4 |
|------------------|--------|--------|--------|
| ----- | ----- | ----- | ----- |
| VAS server id 0: | | 0 | 0 |
| VAS server id 1: | | 0 | 0 |
| VAS server id 2: | | 0 | 0 |
| VAS server id 3: | | 0 | 0 |
| VAS server id 4: | | 0 | 0 |
| VAS server id 5: | | 0 | 0 |
| VAS server id 6: | | 0 | 0 |
| VAS server id 7: | | 0 | 0 |

Traffic after VAS processing TxBW [Kbps] (bytes are counted from Layer 2):

| Port 1 | Port 2 | Port 3 | Port 4 |
|------------------|--------|--------|--------|
| ----- | ----- | ----- | ----- |
| VAS server id 0: | | 0 | 0 |
| VAS server id 1: | | 0 | 0 |
| VAS server id 2: | | 0 | 0 |
| VAS server id 3: | | 0 | 0 |
| VAS server id 4: | | 0 | 0 |
| VAS server id 5: | | 0 | 0 |
| VAS server id 6: | | 0 | 0 |
| VAS server id 7: | | 0 | 0 |

VAS over 10G

- [About VAS over 10G, page 12-32](#)
- [Data Flow in VAS over 10G Topology, page 12-33](#)
- [Failover Support, page 12-37](#)
- [Health Check in VAS over 10G Topology, page 12-38](#)
- [Configuring VAS over 10G: General Guidelines, page 12-39](#)
- [How to Configure VAS over 10G, page 12-40](#)
- [How to Configure Health Check for VAS over 10G, page 12-42](#)
- [How to Enable the Health Check for VAS over 10G Topology, page 12-44](#)
- [VAS Over 10G Sample Configuration, page 12-45](#)


About VAS over 10G

A specific configuration of VAS traffic forwarding is VAS over 10G using a Cisco 6500/7600 Series router as a dispatcher. The VAS over 10G topology is a specific application of the Cisco Multi-Gigabit Service Control Platform (MGSCP) solution in which only one external 10G link is supported. The 7600 distributes the external 10G link and also functions as the switch for the VAS servers.

VAS functionality is supported over a dual 10G topology only. This topology provides a solution with no single point of failure.

In this topology, there are two external 10G links, each one connected to a separate 7600 platform and VAS server array. Only one set of VAS servers is used at a time, serving the VAS traffic of both 10G links. The other set of VAS servers is reserved for failover in case of either a switch failure or VAS server failure. The following figure illustrates the VAS over 10G topology.

Figure 12-6 Typical VAS over 10G Topology



157213

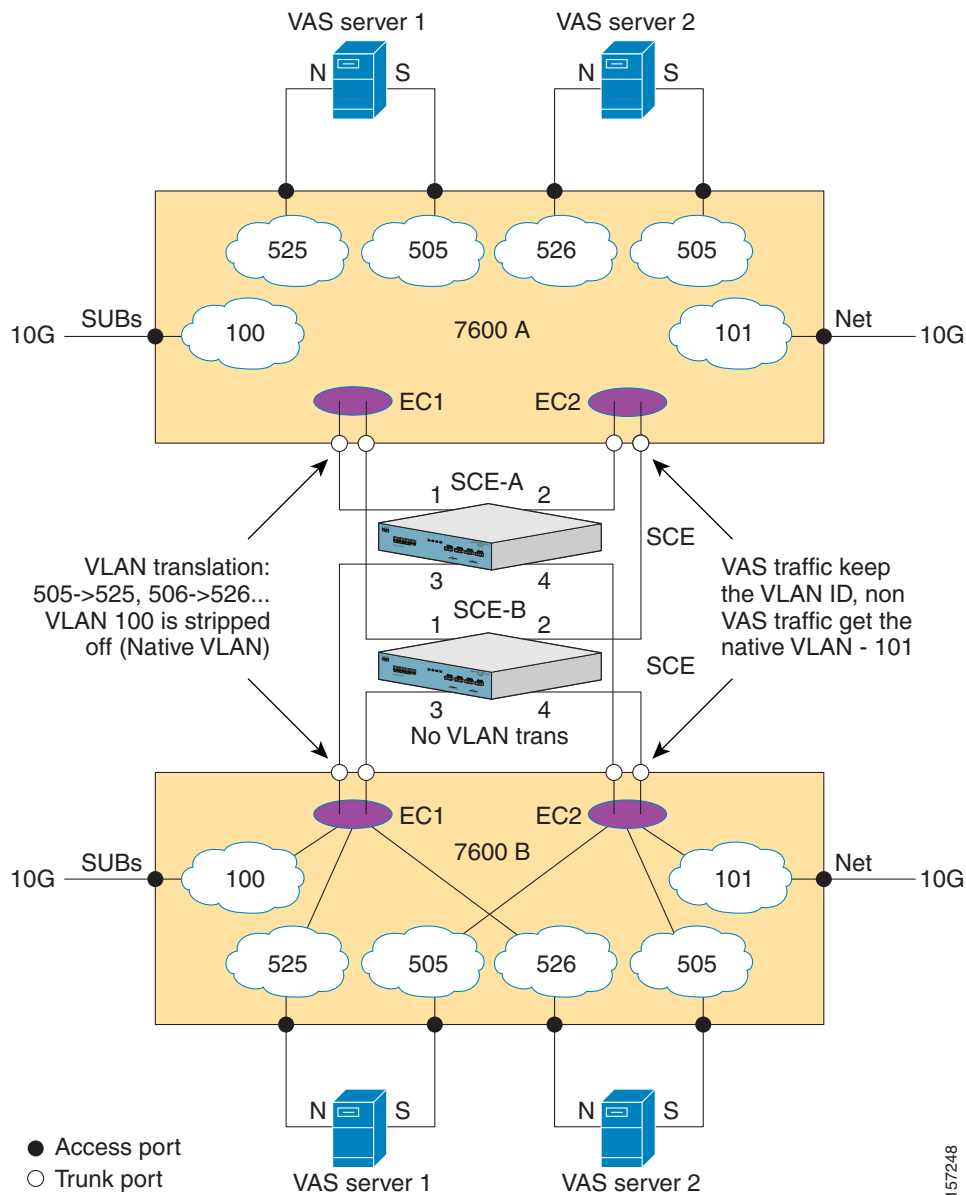
Data Flow in VAS over 10G Topology

- [VAS Data Flow: To the VAS Server, page 12-35](#)
- [VAS Data Flow: From the VAS Server, page 12-36](#)

Data flow in the VAS solution over 10G topology depends on the appropriate use of VLAN tags to route the packets through the system, from the 7600/6500 to the SCE platform, to the appropriate VAS server, back to the SCE platform and then back to the network through the 7600/6500.

The following figure illustrates the flow of VAS data in the VAS solution over 10G topology. Note that the path between the SCE platform and the VAS servers has the same VLAN tag for all SCE platforms in the same EtherChannel.

Figure 12-7 Data Flow in VAS over 10G Topology



157248

VAS flows enter the SCE platform without a VLAN tag, and are then transmitted from the SCE platform with the VLAN tag of the VAS server. They must return from the VAS server to the SCE platform with this tag, which is then stripped off.


In the solution using the 7600/6500, VLAN tags are also used to identify the external link. Note that this is not the same VLAN tag used for the VAS servers. This VLAN tag must be defined as native in the trunk ports towards the SCE platforms, so that the external traffic arrives at the SCE platform without a VLAN tag.

In this description of the VAS data flow, note the following important points:

- This section describes the data flow of packets originating at the subscriber side of the 10G link and sent towards the network. The network to subscriber flow exactly mirrors this flow.
- The description does not elaborate on the internal path inside the 7600/6500 device. It is intended to present the path between the SCE platform, the 7600/6500, and the VAS servers, and describe how the VLAN tag changes along that path.
- Although the figures show only one SCE platform, in actuality the VAS over 10G topology would usually consist of multiple SCE platforms on multiple ECs. In such a topology, the ports towards the VAS servers must be trunk ports, which allow the presence of multiple VLAN tags, since there will be a unique VLAN tag for each EC. (As noted previously, all SCE platforms on one EC must use the same VLAN tag per VAS server).
- The data flow is presented in two parts:
 - To the VAS servers
 - From the VAS servers
- The figures assume that the VAS link is link 1.

VAS Data Flow: To the VAS Server

Figure 12-8 Data Flow in VAS over 10G Topology: To the VAS Server




The following is the sequence of steps in the VAS data flow from the subscriber side to the VAS server:

1. A subscriber packet is received at the 7600/6500 external 10G link.
2. The packet is marked with VLAN 100 at the access port and is sent towards the EtherChannel (EC) configured with VLAN 100.
3. The EC selects port 1 of the SCE platform based on the subscriber side IP, and the VLAN tag is stripped off (100 is the native VLAN tag in the 7600/6500 trunk port).
4. The SCE platform classifies the flow as VAS flow and tags it with the VAS server VLAN tag 505.
5. The packet is sent to the VAS server from SCE platform port 2 (N) towards the 7600/6500 with VLAN tag 505.
6. The packet is received on the 7600/6500 trunk port and is sent to the access port configured with VLAN 505, which is the port connected to the VAS server subscriber side.

The packet has no VLAN tag when it arrives at the VAS server.

VAS Data Flow: From the VAS Server

Figure 12-9 Data Flow in VAS over 10G Topology: From the VAS Server



The following is the sequence of steps in the VAS data flow from the VAS server out to the network:

1. The VAS server processes the packet and sends it to the SCE platform from the VAS network port. Note that as the VAS server received the packet with no VLAN tag, it also transmits it with no VLAN tag.
2. The packet is received on the 7600/6500 access port, is assigned VLAN tag 525, and is sent towards the EC configured with VLAN 525.
3. At the trunk port, the VLAN tag 525 is translated to VLAN tag 505.
4. The packet is sent towards port 1 of the SCE platform based on the subscriber side IP.
5. SCE platform gets the packet on port 1 (S) and forwards it towards the network through port 2 (N). The SCE platform forwards the packet with NO VLAN tag.
6. The packet is received on the 7600/6500 trunk port, gets the native VLAN 101 and sent towards the access port configured with VLAN 101.
7. The packet is sent towards the network with no VLAN tag.

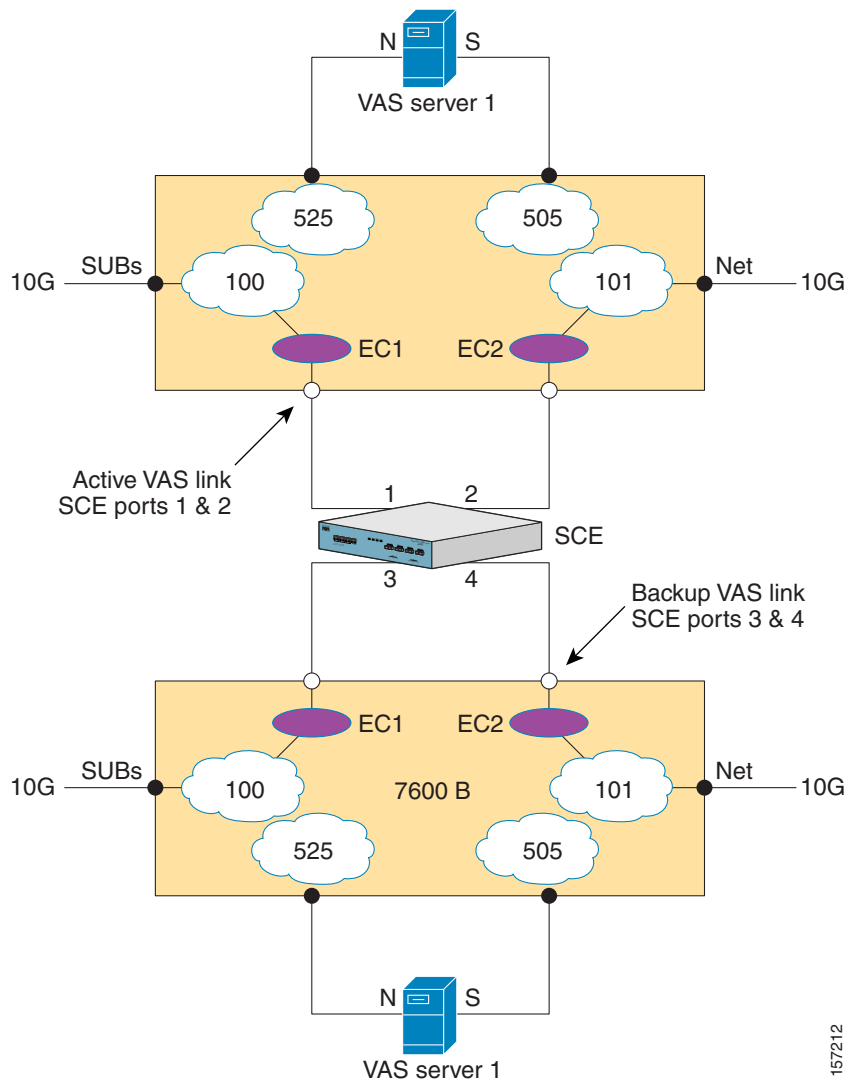
Failover Support

The SCE monitors the health of the connection to each VAS server on the active VAS link; failure of either the 7600/6500 or the VAS server will cause the server health check to fail. The failure of one or more server groups causes the VAS traffic to be forwarded to the redundant 7600/6500/VAS servers system.

How failover works:

- The VAS link (the link on which to send the VAS traffic) is dynamically selected. In failover, the SCE platform switches to its backup subscriber and network ports, so that the VAS traffic is forwarded to the redundant set of VAS devices, as illustrated in the following figure.

Figure 12-10 VAS over 10G Topology: Failover



- The VAS link does not revert automatically. It will be switched again only if required due to a VAS server group failure on this link.
- The system always checks only one set of the VAS servers; those on the active VAS link.
- There is a user-configurable parameter that controls the link-switch rate if there has not been a successful health check on the current VAS link. The default value of this parameter is less than the failure detection time. It is recommended to configure a larger value.

Once there is a successful health check on the VAS link, the link switches immediately upon failure (see [How to Configure the Minimum Time between Link Switches, page 12-41](#)).

In the VAS over 10G topology, a failure may occur at any of the following three points:

- Failure of an SCE platform

SCE platform failure is detected and handled by the 7600/6500 device. In such case, the EtherChannel will balance the traffic load between the other active SCE platforms. SCA BB and VAS service through the other SCE platforms continues uninterrupted with no change in the VAS link.

- Failure of a 7600/6500

The 10G link that goes through the failed 7600/6500 device is lost completely, however SCA BB and VAS services on the 10G link that goes through the second 7600/6500 device continue uninterrupted. In case of failure of the 7600/6500 that is connected to the active set of VAS servers, all the VAS traffic of the other 10G link is forwarded to the standby set of VAS servers

- Failure of a VAS server group

Failure of a VAS server group on the active VAS traffic link is detected by the VAS health check. Failure of any VAS server group triggers the switch of the entire link to the standby VAS servers. A server group failure is declared when the number of active VAS servers drops below the parameter “minimum active VAS servers in a group” (see [How to Configure VAS Server Group Failure Parameters, page 12-25](#)).

Both links preserve SCA BB and VAS services. However, during the transition period, the replacing VAS servers will see VAS flows in the middle and the VAS service may be temporarily damaged.

Health Check in VAS over 10G Topology

In the VAS over 10G topology, special attention must be paid to the selected IP addresses of the health check flows. A flow initiated by an SCE platform may not always be hashed correctly by the EtherChannel. Therefore, a health check packet can be sent out from one SCE platform, go through the 7600/6500 towards the VAS server, then be sent back from the VAS server through the 7600/6500 but hashed by the EtherChannel to a SCE platform different than the originator SCE platform.

To prevent this from happening, the SCE platform opens eight flows per VAS server. This ensures that at least one of the flows will be mapped to the correct SCE platform; the other SCE platforms disregard health check packets not initiated by them.

Configuring VAS over 10G: General Guidelines

- [Configuring the 7600/6500 for VAS over 10G, page 12-39](#)

When configuring VAS over 10G, the following changes must be made to the configuration process:

- Configuration of the VAS traffic link is different (see [Configuring the VAS Traffic Link \(VAS over 10G\)](#))
- You must configure a range of IP addresses as the source IP addresses for the health check (see [How to Configure the Health Check IP Address, page 12-43](#)).
- The health check must be specifically enabled to be compatible with the 10G topology (see [How to Enable Health Check Compatibility for VAS over 10G \(MGSCP\), page 12-44](#)).



Note

The VLAN tags and configuration of the two sets of VAS servers must be identical.



Note

Additional VAS traffic forwarding configuration and monitoring options are available from the SCA BB Console. See *Managing VAS Traffic Forwarding Settings* in the *Cisco Service Control Application for Broadband User Guide* .

Configuring the 7600/6500 for VAS over 10G

This section explains some important points to keep in mind when configuring the 7600/6500 as part of the VAS over 10G solution. For complete information on how to configure the 7600/6500, please refer to the appropriate Cisco documentation.

Please refer to the following guidelines when configuring the 7600/6500 as part of the VAS over 10G solution:

- The 7600/6500 device traffic distribution is based on the EtherChannel dispatching function. Specifically it is required that:
 - External traffic coming from the subscriber side of the 7600/6500 device must be hashed by the EtherChannel according to the source IP
 - External traffic coming from the network side must be hashed according to the destination IP.
- This requirement insures that the same SCE platform handles all the traffic of a subscriber. Since the hashing metric is configurable per line card, the external 10G link subscriber and network ports must be on different line cards. The VAS servers should be connected to the 7600/6500 following this convention as well:
 - Connect the VAS server subscriber leg to the same line card as the network 10G port, or to a line card that is configured with per destination IP dispatching function.
 - Connect the VAS server network leg to the same line card as the subscriber 10G port, or to a line card that is configured with per source IP dispatching function.
- In order for the native VLAN configuration to be effective, disable the "vlan dot1q tag native" configuration on the 7600/6500.
- Introduce the VLAN tags of the VAS servers and the external subscriber and Network ports by running the "vlan XXX" configuration command on the 7600/6500 device.

How to Configure VAS over 10G

- [How to Configure the VAS Traffic Link Auto-Select Parameters \(VAS over 10G\), page 12-40](#)
- [How to Configure the Minimum Time between Link Switches, page 12-41](#)
- [How to Set the Active VAS Link, page 12-42](#)

How to Configure the VAS Traffic Link Auto-Select Parameters (VAS over 10G)

- [How to Configure the Link for VAS over 10G, page 12-40](#)
- [How to Revert to the Default Link Configuration, page 12-40](#)

To enable switching the VAS traffic automatically upon a failover situation, the following options must be configured for VAS over 10G:

- Set the VAS traffic link to auto-select, so that the system can switch between the links in case of 7600/6500/VAS servers failure. .
- Specify the minimum time allowed between two consecutive link switches before any health check has succeeded.
- Specify the link on which to transmit VAS traffic initially after changing the configuration to ‘ **auto-select** ’ (in runtime or after reload) or the current VAS traffic link if ‘ **auto-select** ’ is already configured

How to Configure the Link for VAS over 10G

By default, the VAS traffic is transmitted on Link 1. However, for VAS over 10G, the VAS link should be set to auto-select, so that the system can switch to the backup link when necessary.

Options

The following option is available:

- **VAS traffic-link {link-0|link-1|auto-select}** — The link number on which to transmit VAS traffic
 - For VAS over 10G, specify **auto-select** .

Step 1 From the SCE(config if)# prompt, type `VAS-traffic-forwarding traffic-link auto-select` and press Enter.

How to Revert to the Default Link Configuration

By default, the VAS traffic is transmitted on Link 1.

Options

The following option is available:

- **VAS traffic-link {link-0|link-1|auto-select}** — The link number on which to transmit VAS traffic
 - For VAS over 10G, specify **auto-select** .

Step 1 From the SCE(config if)# prompt, type `no VAS-traffic-forwarding traffic-link` and press Enter.

How to Configure the Minimum Time between Link Switches

- [Options, page 12-41](#)
- [How to Configure the Delay between Link Switches, page 12-41](#)
- [How to Revert to the Default Delay Configuration, page 12-41](#)

You can configure the minimum time allowed between two consecutive link switches. This parameter applies only after a link switch and before any health check has succeeded.

Note that the system assumes the servers are UP while the health check initializes (initialization occurs as a result of every change in the configuration related to the health check or after a link switch, and it lasts until the first health check success or failure). This means that even if the servers are actually DOWN or not even connected, it is assumed that they are UP and user traffic is forwarded to them. Once the health check fails, the servers are declared to be Down, and user traffic is no longer forwarded to them.

In VAS over 10G topology, the default delay between two consecutive link switches (30 seconds) is less than the time it takes for the health check to fail. This means that once a VAS server group fails, the SCE platform switches immediately to the second link.

In cases where there is at least one failed VAS server group on both links, the SCE platform will flip continuously between the links, and as described above, most of this time the state of the servers will be UP.

To avoid the constant flip between the links in such a case, it is recommended to configure a link-switch-delay time greater than 3 minutes.

It is also recommended to monitor the SNMP traps conveying messages on changes in server status.

Options

The following option is available:

- **switch-time** — The time in seconds to hold between two consecutive link switches on initial health check state.
 - Default = 30 seconds

How to Configure the Delay between Link Switches

-
- Step 1** From the SCE(config if)# prompt, type `VAS-traffic-forwarding traffic-link auto-select link-switch-delay switch-time` and press Enter.
-

How to Revert to the Default Delay Configuration

-
- Step 1** From the SCE(config if)# prompt, type `no VAS-traffic-forwarding traffic-link auto-select link-switch-delay` and press Enter.

You can also use the **default** form of the command:

default VAS-traffic-forwarding traffic-link auto-select link-switch-delay

How to Set the Active VAS Link

- [Options, page 12-42](#)
- [How to Set the Active VAS Link, page 12-42](#)
- [How to Revert to the Default Active VAS Link Configuration, page 12-42](#)

Use this command to set the active VAS link, the link on which to transmit VAS traffic after a system reload and when working in auto-select mode.

When executed, this command triggers an immediate link switch if the currently active VAS traffic link used is different from the one specified in the command.

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 Set the Active VAS Link

-
- Step 1** From the SCE(config if)# prompt, type `VAS-traffic-forwarding traffic-link auto-select initial-selection {link-0 | link-1}` and press Enter.
-

How to Revert to the Default Active VAS Link Configuration

-
- Step 1** From the SCE(config if)# prompt, type `no VAS-traffic-forwarding traffic-link auto-select initial-selection` and press Enter.

You can also use the **default** form of the command:

default VAS-traffic-forwarding traffic-link auto-select initial-selection

How to Configure Health Check for VAS over 10G

When configuring the health check for VAS over 10G, you must perform the following steps:

- Configure the health check source and destination IP addresses
- Enable health check compatibility for VAS over 10G

How to Configure the Health Check IP Address

- [About the Health Check IP Address, page 12-43](#)
- [Options, page 12-43](#)

About the Health Check IP Address

Use this command to configure the IP addresses to be used for the VAS health check flows. Any traffic to the configured IP address will be handled as belonging to health check flows; it will not be processed as usual traffic and will be dropped by the SCE platform.

There are three important rules to follow when configuring the IP addresses of the VAS health check. Improper configuration will cause the health check to fail and may cause health check traffic to be forwarded outside of the 7600/6500.

- It is required to configure a range of IP addresses (at least eight IP addresses) for the source IP. This will insure that at least one out of the eight flows will be hashed correctly to the SCE platform. If no range is configured and the VAS over 10G mode (MGSCP) is selected, the health check will fail to run.
- The configured IP addresses must be unique to the SCE platforms and should not exist in the network. Any traffic to the configured IP addresses other than VAS health check traffic will be regarded as fault traffic and be dropped by the SCE platform.
- All the SCE platforms under the same EtherChannel must have the same IP address configuration. Using the same IP addresses allows the SCE platform to correctly identify health check flows coming from other SCE platforms (as a result of the EtherChannel hashing) and drop these flows before they are transmitted out of the SCE platform

Options

The following options are available:

- **ip-address** — Specify the IP addresses to be used for the health check:
- **source-ip** — health check source IP address The source-ip must include a range indication (A.B.C.D/E or A.B.C.D:0xMASK, where A,B,C,D are numbers between [0,255] and E is in the range [0,32]. MASK is the IP mask in 8 hexadecimal characters.)
- **dest-ip** — health check destination IP address

The configured IP addresses should not be in use in the network.

The same IP address should be used by all the SCE platforms under the same EtherChannel.

Use the no form to remove the configured IP addresses.

-
- Step 1** From the SCE(config if)# prompt, type `vas-traffic-forwarding health-check ip-address source source-ip destination dest-ip` and press Enter.
-

How to Remove the IP Address Configuration

Step 1 From the SCE(config if)# prompt, type `no vas-traffic-forwarding health-check ip-address` and press Enter.

You can also use the **default** form of the command:

default vas-traffic-forwarding health-check ip-address

How to Enable the Health Check for VAS over 10G Topology

Use this command to configure the health check to adjust to VAS over 10G (MGSCP) conditions (see [Health Check in VAS over 10G Topology, page 12-38](#)).

- [Options, page 12-44](#)
- [How to Enable Health Check Compatibility for VAS over 10G \(MGSCP\), page 12-44](#)
- [How to Remove the Health Check Compatibility Configuration, page 12-44](#)

Options

The following options are available:

- The keyword **MGSCP** is specified to enable health check compatibility because VAS over 10G is a special case of a MGSCP (Multi-Gigabit Service Control Platform) system.
 - By default, VAS over 10G compatibility is disabled.

How to Enable Health Check Compatibility for VAS over 10G (MGSCP)

Step 1 From the SCE(config if)# prompt, type `vas-traffic-forwarding health-check topology MGSCP` and press Enter.

How to Remove the Health Check Compatibility Configuration

Step 1 From the SCE(config if)# prompt, type `no vas-traffic-forwarding health-check topology MGSCP` and press Enter.


You can also use the **default** form of the command:

default vas-traffic-forwarding health-check topology MGSCP

VAS Over 10G Sample Configuration

Following is a sample illustrating the steps in configuring the VAS over 10G solution.

-
- Step 1** From the SCE(config)# prompt, type `interface LineCard 0` and press Enter.
Enter the LineCard Interface configuration mode.
- Step 2** From the SCE(config if)# prompt, type `VAS-traffic-forwarding health-check topology MGSCP` and press Enter.
Set VAS health-check to MGSCP mode.
- Step 3** From the SCE(config if)# prompt, type `VAS-traffic-forwarding health-check ip-address source 192.168.100.0:0xfffff00 destination 192.168.101.0` and press Enter.
Sets VAS health-check source IP address to the range 192.168.100.0/24 and destination to 192.168.101.0
- Step 4** From the SCE(config if)# prompt, type `VAS-traffic-forwarding traffic-link auto-select` and press Enter.
Sets VAS traffic link to *auto-select* so in case of a failure in any of the VAS servers group, the VAS traffic link will be automatically switched.
- Step 5** From the SCE(config if)# prompt, type `VAS-traffic-forwarding traffic-link auto-select link-switch-delay 240` and press Enter.
Configures the link switch delay to 4 minutes. The delay will be applied only if there was no successful health check on the current link.
- Step 6** From the SCE(config if)# prompt, type `#VAS-traffic-forwarding traffic-link auto-select initial-selection link-0` and press Enter.
Sets link-0 to be used as the initial VAS traffic link in auto-select mode.
- Step 7** Assign VAS servers 0-3 to VLAN 600-603 respectively.
`SCE(config if)#VAS-traffic-forwarding VAS server-id 0 VLAN 600`
`SCE(config if)#VAS-traffic-forwarding VAS server-id 1 VLAN 601`
`SCE(config if)#VAS-traffic-forwarding VAS server-id 2 VLAN 602`
`SCE(config if)#VAS-traffic-forwarding VAS server-id 3 VLAN 603`
- Step 8** Map VAS servers 0-1 and 2-3 to server groups 0 and 1 respectively, allowing server redundancy within each group:
`SCE(config if)#VAS-traffic-forwarding VAS server-group 0 server-id 0`
`SCE(config if)#VAS-traffic-forwarding VAS server-group 0 server-id 1`
`SCE(config if)#VAS-traffic-forwarding VAS server-group 1 server-id 2`
`SCE(config if)#VAS-traffic-forwarding VAS server-group 1 server-id 3`
- Step 9** From the SCE(config if)# prompt, type `#VAS-traffic-forwarding` and press Enter.
Sets the SCE platform to forward VAS traffic (enables VAS traffic forwarding).
-



CHAPTER 13

MPLS/VPN Support

This module provides an overview of the Service Control MPLS/VPN support. It also explains the various procedures for configuring and monitoring MPLS/VPN.

- [Service Control in the MPLS/VPN Environment, page 13-2](#)
- [Definitions and Acronyms, page 13-3](#)
- [What are the Challenges for Service Control for MPLS/VPN Support?, page 13-3](#)
- [How MPLS/VPN Support Works, page 13-4](#)
- [Service Control MPLS/VPN Concepts, page 13-7](#)
- [Service Control MPLS/VPN Requirements, page 13-9](#)
- [How to Configure MPLS/VPN Support, page 13-11](#)
- [How to Manage MPLS/VPN Support, page 13-17](#)

Service Control in the MPLS/VPN Environment

MPLS/VPN networks are very complex and utilize many routing protocols and many different levels of addressing and control. In addition, the various VPNs may use overlapping IP addresses (private IPs).

The SCE platform makes a distinction between identical IP addresses that come from different VPNs, and maps them into subscribers according to the MPLS labels attached to the packets. This involves various mechanisms in all levels of the system.

The following assumptions and requirements allow the SCE platform to operate in an MPLS/VPN environment:

- The MPLS/VPN architecture is according to RFC-2547.
- The specific type of encapsulation used is the MPLS shim header over Ethernet (described in RFC-3032).
- There are two levels of MPLS labels.
 - External labels — Used for transport over the service provider MPLS core network.

These labels are not mandatory for VPN classification, and some situations do not appear in the packet due to PHP or other reasons.

- Internal labels (BGP labels) — Used to identify the VPNs connected to each edge router, and typically controlled by the BGP protocol.

These labels are mandatory for VPN classification.

- The MPLS/VPN solution contains the SCE platform and the SM. The SM acts as a BGP peer for the PE routers in the service provider network, and communicates the BGP information to the SCE platform as subscriber information.

**Note**

The MPLS/VPN solution supports the existence of non-VPN-based subscribers concurrently with the MPLS/VPN-based subscribers (see [Non-VPN-Based Subscribers, page 13-7](#)).

Definitions and Acronyms

The following table defines important terms and acronyms.

Table 13-1 *MPLS/VPN Terms and Acronyms*

| Term or Acronym | Definition |
|---|---|
| PE (Provider Edge router) | A router at the edge of the service provider network. The PE routers are the ones that connect to the customers, and maintain the VPNs |
| P (Provider router) | A router in the core of the service provider network. P routers only forward MPLS packets, regardless of VPNs. |
| VPN (Virtual Private Network) | In the Service Control context, a VPN is the part of the VPN that resides in a specific site. It is a managed entity over which private IP subscribers can be managed. |
| BGP LEG | A software module that resides on the SM server and generates BGP-related login events. The BGP LEG communicates with the BGP routers (PEs) and passes the relevant updates to the SM software, which generates login events to the SCE platform for the updated VPN-based subscribers. |
| Upstream | Traffic coming from the PE router and going into the P router |
| Downstream | Traffic coming from the P router and going into the PE router |
| RD (Route Distinguisher) | Used to uniquely identify the same network/mask from different VRFs (such as, 10.0.0.0/8 from VPN A and 10.0.0.0/8 from VPN B) |
| RT (Route Target) | Used by the routing protocols to control import and export policies, to build arbitrary VPN topologies for customers |
| VRF (Virtual Routing and Forwarding instance) | Mechanism used to build per-interface routing tables. Each PE has several VRFs, one for each site it connects to. This is how the private IPs remain unique. |

What are the Challenges for Service Control for MPLS/VPN Support?

- Private IP addresses cause flows to look the same except for their MPLS labels.
- The MPLS labels are different in each direction, and must be matched.
- Detecting that a flow belongs to a certain VPN is complicated by the fact that in the downstream direction there is no external label. The SCE platform must be able to understand the VPN information from the internal label + the MAC address of the PE..

How MPLS/VPN Support Works

Service Control supports three mechanisms that make MPLS/VPN support work:

- Flow detection – This is the job of the SCE platform, to match upstream and downstream traffic to identify flows.
- VPN detection – Downstream VPN labels are identified by the SM. The SCE platform learns the upstream labels from the traffic to identify the VPN.
- Subscriber detection – The SM and the SCE platform function together to identify the IP range within a VPN that is defined as a single subscriber.

Flow Detection

Flow detection is the process of deciding which packets belong to the same flow. This relates to the first two challenges listed:

- Private IP addresses cause flows to look the same except for their MPLS labels.
- The MPLS labels are different in each direction, and must be matched.

Flow detection is based on the MPLS labels, extending the basic 5 tuple that SCOS uses to identify flows, and taking into account the fact that in MPLS, the packet is labeled differently in each direction.

Since MPLS traffic is unidirectional, each direction is classified separately by the SCE platform, using the following:

- Downstream – the BGP label and the MAC address of the PE (only one label that is relevant to the classification)

Downstream labels are learned from the control plane (through the SM BGP LEG).

- Upstream – the combination of the external label, the BGP label, and the MAC address of the P router (two labels that are relevant to the classification)

Upstream labels are learned from the data plane.

VPN Detection

The network configuration that provides the division into VPNs is controlled by the SM. The network-wide value that describes a VPN most closely is either the Route Target or the Route Distinguisher

- The administrator configures the SM to detect VPNS, according to selected attribute (RT or RD) (see [How to Configure MPLS/VPN Support, page 13-11.](#))
- The network operator provides the SCE platform with a mapping between RT values and VPN subscriber names. (See [How to Manage MPLS/VPN Support, page 13-17](#))

The relevant module in the Subscriber Manager server (SM) is the BGP-LEG. The BGP-LEG is added to the BGP neighborhood for obtaining the information on the MPLS labels. The local PEs are configured to add the BGP-LEG as a BGP peer.

The SCE platform detects that a flow belongs to a certain VPN according to the downstream label that the flow carries, and the MAC address of the PE that it is sent to.

One VPN may spread over more than one PE router, as long as all the sites of the VPN are connected to the subscriber side of the same SCE platform

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

Subscriber Detection

- [What is an MPLS/VPN-based Subscriber?, page 13-5](#)
- [Private IP Subscriber Support, page 13-5](#)

What is an MPLS/VPN-based Subscriber?

As in other modes of operation, in MPLS/VPN each flow belongs to a certain subscriber. A VPN-based subscriber is a customer part of the Service Provider, who pays for the VPN service; which, in the SCE platform, corresponds to a set of IP addresses that belong to a certain VPN that is managed separately.

An MPLS/VPN-based subscriber can be defined as either of the following

- A set of IP addresses or ranges in a certain VPN.
- All the IP addresses of a CE router, defined by a BGP community over a VPN.

The network configuration that provides the division into VPNs and VPN-based subscribers is controlled by the SM. (For more information, see the *Cisco Service Control Management Suite Subscriber Manager User Guide*)

Private IP Subscriber Support

VPN-based subscribers can have private IP mappings, which are a combination of an IP range and a VPN mapping. Since the source of such mappings is typically in the BGP protocol, and they are received automatically from the protocol by the BGP agent, the IP ranges may contain overlapping ranges. The semantics of such overlaps is that of a longest prefix match.

For example, if subscriber A receives the range 10.0.0.0/8@VPN1 and subscriber B receives the range 10.1.0.0/16@VPN1, then the system maps IPs that start with 10.1 to subscriber B, and any other address that begins with 10 to subscriber A. Traffic with other IP addresses on VPN1 will be mapped to the unknown subscriber.

For private IP subscribers, flows are distributed to traffic processors according to the VPN, not according to the IP address. This means that all traffic from any one VPN is mapped to the same traffic processor.

How the Service Control MPLS/VPN Solution Works

- [How the Service Control MPLS/VPN Solution Works: A Summary, page 13-6](#)
- [SCE Platform Tasks in the MPLS/VPN Solution, page 13-6](#)
- [BGP LEG Tasks in the MPLS/VPN Solution, page 13-6](#)
- [SM Tasks in the MPLS/VPN Solution, page 13-6](#)

How the Service Control MPLS/VPN Solution Works: A Summary

- The SM is configured with the VPNs and VPN-based subscribers that should be managed. A VPN is identified by the RD / RT and the PE.
- The BGP-LEG updates the SM with the MPLS labels and IP routes.
- The SM pushes the VPNs with their labels and the VPN-based subscriber to the SCE platform with the downstream MPLS labels of the VPN.
- The SCE platform resolves the PE MAC addresses and updates its tables with the new information.
- The SCE platform learns the upstream labels, including the P MAC address.
- The SCE platform provides the regular services to the VPN-based subscribers (BW management, reports, etc.)

SCE Platform Tasks in the MPLS/VPN Solution

- Matching upstream to downstream labels
 - Mappings of downstream labels to VPNs are received from the SM
 - Upstream labels are learned from the data
- The MAC addresses of the PEs are used to distinguish downstream labels of different PEs
- After the learning, each flow is classified as belonging to one of the VPNs.
- The SCE platform performs a longest prefix match on the IP address inside the VPN, and classifies each flow to the correct VPN-based subscriber
- The SCE platform runs the SCA-BB application for the network flows, which are classified to VPNs, thus providing subscriber aware service control and reporting

BGP LEG Tasks in the MPLS/VPN Solution

- The BGP LEG is a software module that runs on the SM server
- The LEG maintains a BGP session with a list of PEs
- After the sessions establishment, the LEG propagates MP-BGP route-updates from the PEs to the SM module

SM Tasks in the MPLS/VPN Solution

- The VPNs are stored in the SM database.
- Each VPN is defined by:
 - The IP address of the loopback interface of the PE router.
 - The RD or RT that identifies the VPN within the PE router.
- A VPN-based subscriber is defined by the IP range in a specified VPN or the BGP community (CE as subscriber).
- The SM receives updates from the BGP LEG, and updates the VPN information with the new MPLS labels.
- The relevant SCE platforms that will get the MPLS updates are defined by the VPN domain.

Service Control MPLS/VPN Concepts

- [Non-VPN-Based Subscribers, page 13-7](#)
- [Bypassing Unknown VPNs, page 13-7](#)
- [Additional MPLS Pattern Support, page 13-8](#)
- [VPN Identifier \(RD or RT\), page 13-8](#)

Non-VPN-Based Subscribers

The MPLS/VPN solution supports the existence of non-VPN-based (regular IP) subscribers concurrently with the MPLS/VPN-based subscribers, with the following limitations and requirements:

- The SM must work in "push" mode.
- Non-VPN-based subscribers cannot have IP in VPN mappings.
- VLAN-based subscribers are NOT supported at the same time as MPLS/VPN-based subscribers.

In typical MPLS/VPN networks, traffic that does not belong to any VPN is labeled with a single MPLS label in the upstream direction, which is used for routing. The downstream direction of such flows typically contains no label, due to penultimate hop popping.

The SCE platform uses the one or more labels upstream and no label downstream definition to identify non-VPN flows. Classification and traffic processor load balancing on these flows is performed according to the IP header, rather than the label.

This process requires learning of the upstream labels in use for such flows, and is done using the flow detection mechanism described above (see [Flow Detection, page 13-4](#)).

Bypassing Unknown VPNs

In an MPLS network, there may be many VPNs crossing the SCE platform, only a small number of which require service control functionality. It is necessary for the SCE platform to recognize which VPNs are not managed.

- The SCE platform automatically bypasses any VPN that is not configured in the SM
- The VPNs are bypassed by the SCE platform without any service

Note that the label limit (see [Limitations, page 13-10](#)) of 57,344 different labels includes labels from the bypassed VPNs.

Each bypassed VPN entry, both upstream and downstream, is removed from the database after a set period of time (10 minutes). If the entry is still used in the traffic, it will be re-learned. This allows the database to remain clean, even if the labels are reused by the routers for different VPNs.

show bypassed VPNs

In the **show bypassed VPNs** command, the age is indicated with each label - the length of time since it was learned.

Additional MPLS Pattern Support

The MPLS/VPN solution was designed to provide DPI services in MPLS/VPN network. These networks use BGP protocol as the control plane for the VPNs and LDP protocol for routing. There are complex networks where the MPLS infrastructure is used not only for VPN and routing, but also for other features such as traffic engineering (TE) and better fail-over. These features are usually enabled per VRF in the PE.

The Service Control MPLS/VPN solution does not support VPNs that use other MPLS-related features. Features such as MPLS-TE or MPLS-FRR (Fast Reroute) are not supported. VPNs for which these features are enabled can be automatically bypassed in the system, but are not allowed to be configured in the SM as serviced VPNs. Configuration of these VPNs in the SM might cause misclassification due to label aliasing.

The following list describes the labels combinations that are supported by the SCE platform and how each combination is interpreted by the platform:

- One or more labels upstream, no labels downstream:
Assumed to be non-VPN (see [Non-VPN-Based Subscribers, page 13-7](#)).
The SCE platform treats the following IP flows as non-VPN flows, and ignores their labels.
- One label upstream, one label downstream:
Assumed to be VPN traffic, in which the P router happens to be the last hop in the upstream.
The label in the downstream is treated as a BGP label, like the regular case. If the BGP label is known from the SM, then the flow is assigned to the correct subscriber, otherwise, it is treated as a bypassed VPN.
- Two labels upstream, one label downstream:
This is the typical configuration of the system. Of the two upstream labels, one is for BGP and one for LDP. The downstream label is for BGP only
- More than two labels upstream, or more than one label downstream:
These combinations occur when other MPLS-related features are enabled for the VPN. Such VPNs are not supported and should not be configured in the SM. However, they can be bypassed in the SCE platform without any service and without harming the service for other VPNs.

VPN Identifier (RD or RT)

Either the Route Distinguisher (RD) attribute or the Route Target (RT) attribute can be used to identify the VPN. It is required to decide which attribute best reflects the VPN partitioning, and configure the system accordingly. Note that the configuration is global for all the VPNs, that is, all VPNs must be identified by the same attribute.

The Route Distinguisher (RD) is generally used to distinguish the distinct VPN routes of separate customers who connect to the provider, so in most cases the RD is a good partition for the VPNs in the network. Since the RD is an identifier of the local VRF, and not the target VRF, it can be used to distinguish between VPNs that transfer information to a common central entity (for example a central bank, IRS, Port Authority, etc.).

The Route Target (RT) is used to define the destination VPN site. Though it is not intuitive to define the VPN based on its destination route, it might be easier in some cases. For example, if all the VPN sites that communicate to a central bank should be treated as a single subscriber, consider using the RT as the VPN identifier.

It is important to note that this configuration is global. Therefore, if at some point in time, any VPN would have to be defined by RD, then all the other VPNs must be defined by RD as well. This is a point to consider when designing the initial deployment.

Service Control MPLS/VPN Requirements

- [Topology, page 13-9](#)
- [Capacity, page 13-10](#)
- [Limitations, page 13-10](#)
- [Backwards Compatibility, page 13-11](#)

Topology

Following are the general topology requirements for MPLS/VPN support:

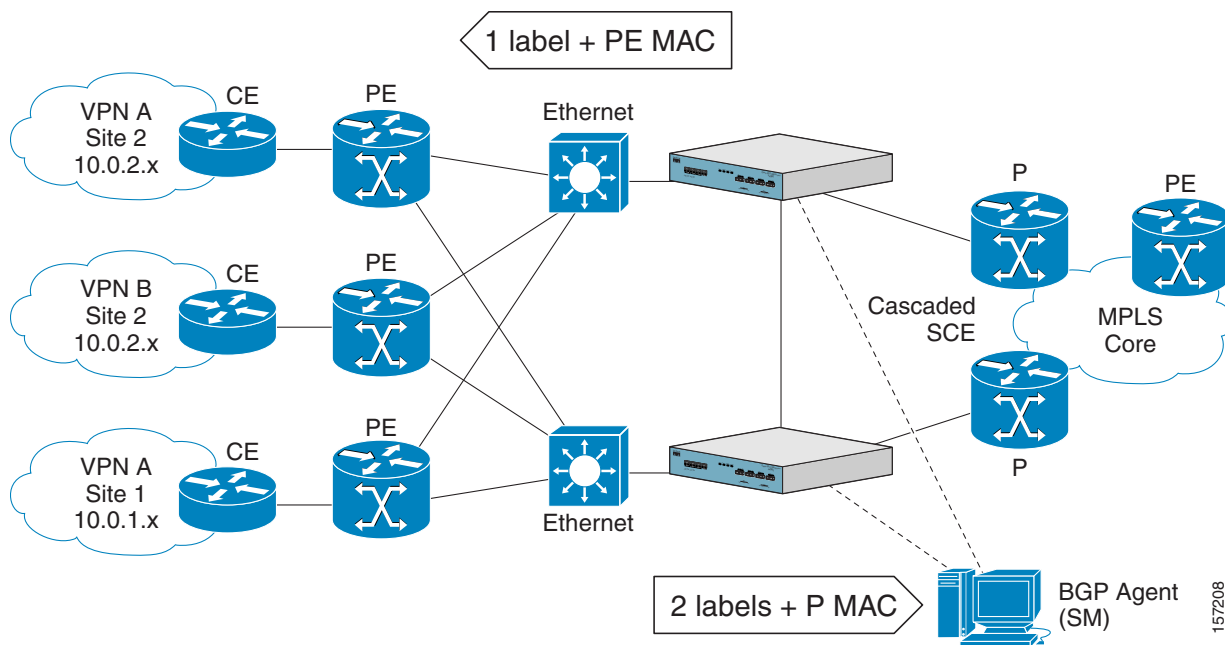
- The SCE platform is placed in the network between the P routers (Provider MPLS core) and the PE (Provider Edge) routers.
- The subscriber side of the SCE platform is connected toward the PE router.
- The network side of the SCE platform is connected toward the P router.
- The BGP LEG is installed on the SM, and is placed somewhere in the network. It speaks with the SCE platform through the management IP.

In a cascade installation:

- The two SCE platforms are connected to each other via the cascade interfaces.
- The data link between the P and the PE is connected via the other interfaces on each SCE platform, as described above:
 - Subscriber side of each SCE platform connected toward the PE router
 - Network side of each SCE platform connected toward the P router

The following drawing depicts a typical cascade installation.

Figure 13-1 Typical MPLS/VPN Installation



Capacity

The system supports:

- 2015 MPLS/VPNs
 - 80,000 IP mappings over VPNs
- 57,344 different labels (including upstream and downstream, and including the bypassed VPNs)
- 256 PEs per SCE platform
 - 4 interfaces per PE

Limitations

Mutually exclusive system modes

When the system works in MPLS/VPN mode, the following modes are not supported:

- The following tunneling modes:
 - MPLS traffic engineering skip
 - MPLS VPN skip
 - L2TP skip
 - VLAN symmetric classify
- TCP Bypass-establishment
- DDoS
- Value Added Services (VAS) mode

Number of MPLS labels

- The choice of the unique VPN site must be based on the BGP label only. The BGP label must be the innermost label.
- The MPLS/VPN solution supports various combinations of labels. See Additional MPLS Pattern Support .
- The system does not support VPNs for which other MPLS-related features, such as MPLS-TE or MPLS-FRR, are enabled.

Subscriber-related limitations

The following subscriber-related limitations exist in the current solution:

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

Topology-related limitations

- An asymmetrical routing topology in which the traffic may be unidirectional, is not supported, since the MPLS/VPN solution relies on the bidirectional nature of the traffic for various mechanisms.

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

- Two VPN sites must be aggregated into one VPN if the following conditions are both true:
 - They are both connected to the same SCE platform
 - They both communicate with a common remote site using the same upstream labels and P router.
- An MPLS/VPN-based subscriber MAY NOT have IP mappings over more than one VPN.

Backwards Compatibility

An SCE platform running SCOS V3.1.5 and up does not support MPLS/VPN subscribers of the type used in older versions. Instead of defining an MPLS/VPN subscriber, which reflects the whole VPN, the user must configure a VPN entity and a full range private IP subscriber within that VPN (0.0.0.0/0@VPN1)

When working with the combination of SM of a version before V3.1.5LA and an SCE with V3.1.5 and up, only regular IP subscribers are supported. VPN-based subscribers are not supported at all in this combination.

How to Configure MPLS/VPN Support

- [Configuring the MPLS Environment, page 13-12](#)
- [How to Configure the SCE Platform for MPLS/VPN Support, page 13-13](#)
- [How to Configure the SM for MPLS/VPN Support, page 13-16](#)

Configuring the MPLS Environment

In order for MPLS/VPN support to function, the environment must be configured correctly, specifically the following are required:

- All other tunneling protocols should be configured to the default mode.
- The MPLS auto-learning mechanism must be enabled.

How to Check the Running Configuration

Check the running configuration to verify no user-configured values appear for tunneling protocols or VLAN support, indicating that they are all in default mode.

-
- Step 1** From the SCE# prompt, type `show running-config` and press Enter.
Displays the running configuration.
- Step 2** Check that no VLAN or L2TP configuration appears.
-

How to Configure the MPLS Environment

If either VLAN or tunneling support is in default mode, skip the relevant step in the following procedure.

-
- Step 1** From the SCE(config if)# prompt, type `default vlan` and press Enter.
Configures VLAN support to default mode.
- Step 2** From the SCE(config if)# prompt, type `no IP-tunnel` and press Enter.
Disables all other tunneling protocol support.



Note All subscribers with VPN mappings must be cleared to change the tunneling mode. If the connection with the SM is down, use the **no subscriber all with-tunnel-mappings** CLI command.



Note In addition, all VPN mappings must also be removed. This can only be done via the SM CLU (which means that the connection with the SM must be up). See [How to Manage MPLS/VPN Support, page 13-17](#)

- Step 3** From the SCE(config if)# prompt, type `MPLS VPN auto-learn` and press Enter.
Enables the MPLS auto-learning mechanism.
-

How to Configure the SCE Platform for MPLS/VPN Support

- [About Configuring the SCE Platform for MPLS/VPN Support, page 13-13](#)
- [How to Define the PE Routers, page 13-13](#)
- [How to Configure the MAC Resolver, page 13-14](#)
- [How to Monitor the MAC Resolver, page 13-16](#)

About Configuring the SCE Platform for MPLS/VPN Support

There are three main steps to configure the SCE platform for MPLS/VPN support:

1. Correctly configure the MPLS tunneling environment, disabling all other tunneling protocols, including VLAN support. (see [How to Configure the MPLS Environment, page 13-12](#))
2. Define all PE routers, specifying the relevant interface IP addresses necessary for MAC resolution (see [How to Define the PE Routers, page 13-13](#).)
3. Configure the MAC resolver (see [How to Configure the MAC Resolver, page 13-14](#).)

How to Define the PE Routers

- [Options, page 13-13](#)
- [How to Add a PE Router, page 13-13](#)
- [How to Remove PE Routers, page 13-14](#)

Options

The following options are available:

- **PE-ID** — IP address that identifies the PE router.
- **interface-ip** — Interface IP address for the PE router. This is used for MAC resolution.
 - At least one interface IP address must be defined per PE router.
 - Multiple interface IP addresses may be defined for one PE router.
 - In the case where the PE router has multiple IP interfaces sharing the same MAC address, it is sufficient to configure just one of the PE interfaces
- **vlan** — A VLAN tag can optionally be provided for each interface IP.

Two interfaces cannot be defined with the same IP address, even if they have different VLAN tags. If such a configuration is attempted, it will simply update the VLAN tag information for the existing PE interface.

How to Add a PE Router

Each PE router that has managed VPNs behind it must be defined using the following CLI command.

Step 1 From the SCE(config if)# prompt, type `MPLS VPN PE-ID pe-id interface-IP interface-ip [vlan vlan][Interface-IP interface-ip [vlan vlan]]` and press Enter.

Defines the PE router with optional VLAN tag and optional additional IP addresses.

How to Remove PE Routers

- [About Removing PE Routers, page 13-14](#)
- [How to Remove a Specified PE Router, page 13-14](#)
- [How to Remove All PE Routers, page 13-14](#)
- [How to Remove a Specified Interface from a PE Router, page 13-14](#)

About Removing PE Routers

Use these commands to remove one or all defined PE routers.

Please note the following:

- You cannot remove a PE if it retains any MPLS mappings. You must logout the VPN and remove all mappings before removing the router it uses. (You must use the SM CLU to remove VPN mappings. See [How to Manage MPLS/VPN Support, page 13-17](#))
- Removing the last interface of a PE router removes the router as well. Therefore, you must logout the relevant VPN to remove the last interface.
- Likewise, all VPNs must be logged out before using the no PE-Database command below, since it removes all PE routers.

How to Remove a Specified PE Router

-
- Step 1** From the SCE(config if)# prompt, type `no MPLS VPN PE-ID pe-id` and press Enter.
Removes the specified PE router.
-

How to Remove All PE Routers

-
- Step 1** From the SCE(config if)# prompt, type `no MPLS VPN PE-Database` and press Enter.
Removes all configured PR routers.
-

How to Remove a Specified Interface from a PE Router

-
- Step 1** From the SCE(config if)# prompt, type `no MPLS VPN PE-ID pe-id interface-IP interface-ip` and press Enter.
Removes the specified interface from the PE router definition. The PE router itself is not removed.
-

How to Configure the MAC Resolver

- [About the MAC Resolver, page 13-15](#)
- [Options, page 13-15](#)
- [How to Add a Static IP Address, page 13-15](#)
- [How to Remove a Static IP Address, page 13-16](#)

About the MAC Resolver

The MAC resolver allows the SCOS to find the MAC address associated with a specific IP address. The MAC resolver must be configured when the SCE platform operates in MPLS/VPN mode, to translate the IP addresses of the provider edge router interfaces to their respective MAC addresses.

The MPLS/VPN mode needs the MAC resolver, as opposed to the standard ARP protocol, because ARP is used by the management interface, while MPLS/VPN uses the traffic interfaces of the SCE platform, which ARP does not include.

The MAC resolver database holds the IP addresses registered by the clients to be resolved. The IP addresses of the routers are added to and removed from the database in either of two modes:

- Dynamic mode (default)

In this mode, the system listens to ARP messages of the configured PE interfaces, and this way it stays updated with their MAC addresses. There is no configuration required when operating in dynamic mode.

- Benefit: it works even if the MAC address of the PE interface changes.

- Drawback: depending on the specific network topology, the MAC resolution convergence time may be undesirably long.

- Static mode

In this mode, the MAC address of each PE router must be explicitly defined by the user.

- Benefit: no initial delay until IP addresses converge

- Drawback: PE interface is not automatically updated via ARP updates; therefore it doesn't automatically support cases where the MAC address changes on the fly.

However, for statically configured MAC addresses, a user log message appears when the system detects that the MAC address changed. This can be used by the operator to configure the new address.

These two modes can function simultaneously; therefore selected PE routers can be configured statically, while the rest are resolved dynamically

For more information regarding the MAC resolver, refer to the *Cisco Service Control Engine Software Configuration Guide*.

Options

The following options are available:

- **ip address** — The IP address entry to be added to or removed from the database.
- **vlan tag** — VLAN tag that identifies the VLAN that carries this IP address (if applicable).
- **mac address** — MAC address assigned to the IP address, in xxxx.xxxx.xxxx format.

How to Add a Static IP Address

Step 1 From the SCE(config if)# prompt, type `mac-resolver arp ip_address[vlan vlan_tag] mac_address` and press Enter.

Adds the specified IP address and MAC address pair to the MAC resolver database.

How to Remove a Static IP Address

Step 1 From the SCE(config if)# prompt, type `no mac-resolver arp ip_address[vlan vlan_tag] mac_address` and press Enter.

Removes the specified IP address and MAC address pair to the MAC resolver database.

How to Monitor the MAC Resolver

Use this command to see a listing of all IP addresses and corresponding MAC addresses currently registered in the MAC resolver database.

Step 1 From the SCE# prompt, type `show interface linecard 0 mac-resolver arp` and press Enter.

Displays a listing of all IP addresses and corresponding MAC addresses currently registered in the MAC resolver database.

How to Configure the SM for MPLS/VPN Support

- [Configuring the SM for MPLS/VPN Support, page 13-16](#)
- [How to Edit the SM Configuration File, page 13-16](#)
- [How to Configure the SM to Allow IP Ranges, page 13-17](#)

Configuring the SM for MPLS/VPN Support

There are two main steps to configure the SM for MPLS/VPN support:

Step 1 Edit the `p3sm.cfg` configuration file to specify the field in the BGP messages that should be used by the SM for MPLS-VPN identification.

See [How to Edit the SM Configuration File, page 13-16](#)

Step 2 Install and configure the BGP LEG

Refer to the *Cisco SCMS SM MPLS/VPN BGP LEG Reference Guide*.

How to Edit the SM Configuration File

The SM configuration file, `p3sm.cfg`, must be configured for the following:

- To specify the field in the BGP messages that should be used by the SM for MPLS-VPN identification.
- To enable IP ranges

How to Configure the SM for MPLS/VPN Support

Step 1 Add the following section to the *p3sm.cfg* configuration file:

```
# The following section enables SM operation with MPLS-VPN support.
[MPLS-VPN]
# The following parameter defines the BGP attribute to use to identify VPN subscribers
# possible values: "rd" or "rt".
# (default: rt)
vpn_id=rt
```

How to Configure the SM for Troubleshooting MPLS/VPN Support

An optional parameter may be turned on to facilitate troubleshooting the BGP LEG installation. This parameter turns on detailed logging of messages received from the BGP LEG. It should only be turned on when necessary for troubleshooting and should always be turned off for normal operation of the system.

Step 1 Add the following parameter to the [MPLS-VPN] section of the *p3sm.cfg* configuration file:

```
# The following parameter turns on detailed logging of messages received from the BGP LEG
# should be changed to true only during troubleshooting
# (default: false)
log_all=true
```

How to Configure the SM to Allow IP Ranges

To setup the SM to work with MPLS/VPN, you must enable IP ranges by setting the **support_ip_ranges** in the configuration file.

Step 1 Set the **support_ip_ranges** parameter in the [Data Repository] section of the *p3sm.cfg* configuration file to 'yes', as in the following example.

```
support_ip_ranges=yes
```



Note

Resetting this parameter requires restarting the SM. This parameter is discarded on regular configuration loading (using CLU).

How to Manage MPLS/VPN Support

- [How to Manage MPLS/VPN Support via SNMP, page 13-18](#)
- [How to Monitor MPLS/VPN Support via SCE Platform CLI, page 13-18](#)
- [How to Manage MPLS/VPN Support via SM CLU, page 13-24](#)

How to Manage MPLS/VPN Support via SNMP

SNMP support for MPLS/VPN auto-learn is provided in two ways:

- MIB variables
- SNMP traps

MPLS/VPN MIB Objects

The `mplsVpnAutoLearnGrp` MIB object group (`pcubeSEObjs 17`) contains information regarding MPLS/VPN auto-learning.

The objects in the `mplsVpnAutoLearnGrp` provide the following information:

- maximum number of mappings
- allowed current number of mappings

For more information, see the "Proprietary MIB Reference" in the *Cisco Service Control Engine Software Configuration Guide*.

MPLS/VPN Traps

There is one MPLS/VPN-related trap:

- `mplsVpnTotalHWMappingsThresholdExceeded` (`pcubeSeEvents 45`)

To provide online notification of a resource deficiency, when the system reaches a level of 80% utilization of the hardware MPLS/VPN mappings, a warning message appears in the user log, and this SNMP trap is sent.

Both the warning and the trap are sent for each 100 mappings that are added after the threshold has been exceeded.

How to Monitor MPLS/VPN Support via SCE Platform CLI

The SCE platform CLI allows you to do the following:

- Display VPN-related mappings
- Monitor subscriber counters
- Monitor PE routers
- Monitor bypassed VPNs

How to Display 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
- The subscriber to whom a specified downstream mapping (PE loopback IP address & BGP label) is mapped. (This option is provided for backwards compatibility and has certain restrictions. See below [How to Display the Name of the Subscriber Mapped to a Specified VPN, page 13-21.](#))

How to Display Mappings for a Specified VPN

- [Options, page 13-19](#)
- [Displaying Mappings for a Specified VPN: Example, page 13-19](#)

Options

The following option is available:

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

Step 1 From the SCE>prompt, type `show interface linecard 0 VPN name vpn-name` and press Enter.

Displaying Mappings for a Specified VPN: Example

```
SCE# show interface linecard 0 VPN name vpn1
VPN name: Vpn1
Downstream MPLS Mappings:
PE-ID = 1.0.0.1 Mpls Label = 20
PE-ID = 1.0.0.1 Mpls Label = 30
=====>Total Downstream Mappings: 2
Upstream MPLS Mappings:
=====>Total Upstream Mappings: 0
Number of subscriber mappings: 0
```

How to Display a Listing of all VPNs

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

Step 1 From the SCE>prompt, type `show interface linecard 0 VPN all-names` and press Enter.

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

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.

Step 1 From the SCE>prompt, type show interface linecard 0 subscriber mapping included-in IP *ip-range*VPN *vpn-name* and press Enter.

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

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.

Step 1 From the SCE>prompt, type show interface linecard 0 subscriber amount mapping included-in IP *ip-range*VPN *vpn-name* and press Enter.

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'.
```


How to Display the Name of the Subscriber Mapped to a Specified VPN

If the MPLS/VPN is configured as a single subscriber mapped to 0.0.0.0/0 on the VPN that is mapped to the specified MPLS, this option displays that subscriber



Note

This command provides backward compatibility for MPLS/VPN subscriber configuration in SCOS versions previous to 3.1.5.

- Step 1** From the SCE# prompt, type `show interface linecard 0 subscriber mapping MPLS-VPN PE-ID pe-idBGP-label label` and press Enter.

Displaying the Subscriber Mapped to a Specified VPN: Example 1

```
SCE#>show interface lineCard 0 subscriber mapping MPLS-VPN PE-ID 1.0.0.1 BGP-label 30
BGP MPLS label 30 on PE 1.0.0.1 is mapped to VPN named 'Vpn1'
The VPN is NOT mapped to a single subscriber (0.0.0.0/0@Vpn1)
```

Displaying the Subscriber Mapped to a Specified VPN: Example 2

```
SCE#>show interface lineCard 0 subscriber mapping MPLS-VPN PE-ID 1.0.0.1 BGP-label 30
BGP MPLS label 30 on PE 1.0.0.1 is mapped to VPN named 'Vpn1'
Subscriber 'Sub10' is mapped to 0.0.0.0/0@Vpn1
```

How to Display the Mappings of Upstream Labels that Belong to Non-VPN Flows

- Step 1** From the SCE# prompt, type `show interface linecard 0 MPLS-VPN non-VPN-mappings` and press Enter.

How to Clear Upstream VPN Mappings

Use this command to remove all learned upstream labels of a specified VPN.

Options

The following option is available:

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

- Step 1** From the SCE# prompt, type `clear interface linecard 0 VPN name vpn-name upstream mpls all` and press Enter.

This command, in effect, causes early label aging. Clearing the mappings allows relearning; labels will probably be quickly relearned after they have been cleared. Therefore, this command is useful when you want to update the VPN mappings without waiting for the standard aging period.

How to Monitor Subscriber Counters

Use the following Viewer command to display subscriber counters, including those related to MPLS/VPN mappings.

About Subscriber Counters

When MPLS/VPN-based subscribers are enabled, the following related counters appear in addition to the basic subscriber counters:

- MPLS/VPN-based subscribers:
 - Current number of MPLS/-based subscribers that have VPN mappings.
 - Maximum number of MPLS/VPN-based subscribers
- MPLS/VPN-based subscribers are also counted in the general subscribers counters, but the general subscribers maximum number does not apply to MPLS/VPN-based subscribers, which have a smaller maximum number.
- MPLS/VPN mappings:
 - Current number of used MPLS/VPN mappings
 - Maximum number of MPLS/VPN mappings
- Note that these values reflect the total number of mappings, not just the mappings used by MPLS/VPN-based subscribers. Bypassed VPNs also consume MPLS/VPN mappings.

Step 1 From the SCE# prompt, type **show interface linecard 0 subscriber db counters** and press Enter.

Monitoring Subscriber Counters: Example

```
SCE#show interface linecard 0 subscriber db counters
Current values:
=====
Subscribers: 2 used out of 99999 max.
Introduced subscribers: 2.
Anonymous subscribers: 0.
Subscribers with mappings: 2 used out of 99999 max.
SINGLE non-VPN IP mappings: 1.
non-VPN IP Range mappings: 1.
IP Range over VPN mappings: 1.
Single IP over VPN mappings: 3.MPLS-based subscribers are enabled. MPLS/VPN mappings: 2
used out of 57344 max. MPLS-based subscribers with VPN mappings: 2 used out of 2015 max.
Subscribers with open sessions: 0.
Subscribers with TIR mappings: 0.
Sessions mapped to the default subscriber: 0.
Peak values:
=====
Peak number of subscribers with mappings: 2
Peak number occurred at: 14:56:55 ISR MON June 9 2007
Peak number cleared at: 15:29:39 ISR MON June 9 2007
Event counters:
=====
Subscriber introduced: 2.
Subscriber pulled: 0.
Subscriber aged: 0.
Pull-request notifications sent: 0.
State notifications sent: 0.
Logout notifications sent: 0.
Subscriber mapping TIR contradictions: 0
```

How to Monitor MPLS/VPN Counters

Use the following Viewer command to display MPLS/VPN information.

Step 1 From the SCE# prompt, type `show interface linecard 0 mpls vpn` and press Enter.

Monitoring MPLS/VPN Counters: Example

```
SCE#show interface linecard 0 mpls vpn
MPLS/VPN auto-learn mode is enabled.
MPLS based VPNs with subscriber mappings: 0 used out of 2015 max
Total HW MPLS/VPN mappings utilization: 0 used out of 57344 max
MPLS/VPN mappings are divided as follows:
  downstream VPN subscriber mappings: 0
  upstream VPN subscriber mappings: 0
  non-vpn upstream mappings: 0
  downstream bypassed VPN mappings: 0
  upstream bypassed VPN mappings: 0
```

How to Monitor the PE Routers

Use the following Viewer commands to monitor PE routers. These commands provide the following information:

- Configuration of all currently defined PE routers.
- Configuration of a specified PE router.

How to Display the Configuration of all Currently Defined PE Routers

Step 1 From the SCE# prompt, type `show interface linecard 0 MPLS VPN PE-Database` and press Enter.

How to Display the Configuration of a Specified PE Router

Step 1 From the SCE# prompt, type `show interface linecard 0 MPLS VPN PE-Database PE-ID pe-id` and press Enter.

How to Monitor Bypassed VPNs

- [How to Display the Currently Bypassed VPNs, page 13-24](#)
- [How to Remove all Learned Bypassed VPNs, page 13-24](#)

How to Display the Currently Bypassed VPNs

Step 1 From the SCE# prompt, type `show interface linecard 0 MPLS VPN Bypassed-VPNs` and press Enter.

How to Remove all Learned Bypassed VPNs

Step 1 From the SCE# prompt, type `clear interface linecard 0 MPLS VPN Bypassed-VPNs` and press Enter.

How to Monitor Non-VPN Mappings

- [How to Display Non-VPN Mappings, page 13-24](#)
- [How to Remove all Learned non-VPN Mappings, page 13-24](#)

How to Display Non-VPN Mappings

Step 1 From the SCE# prompt, type `show interface linecard 0 MPLS VPN non-VPN-mappings` and press Enter.

How to Remove all Learned non-VPN Mappings

Step 1 From the SCE# prompt, type `clear interface linecard 0 MPLS VPN non-VPN-mappings` and press Enter.

How to Manage MPLS/VPN Support via SM CLU

The SM CLU allows you to do the following:

- Add and remove VPNs
- Display VPN information
- Clear MPLS/VPN mappings

For more information, see the *Cisco Service Control Management Suite Subscriber Manager User Guide*.

Managing VPNs

Use the **p3vpn** utility to manage VPNs.

Options

The following options are available:

- **VPN-Name** — The name assigned to the VPN when it was added, or, if adding a VPN, the name to be assigned to it..
- **RT@PE-IP** — The mapping assigned to the VPN. Multiple mappings can be specified using a comma.
 - **RT** = the route target of the VPN, specified using the ASN:n notation or the IP:n notation

Note that the Route Distinguisher may be specified rather than the route target

- **PE-IP** = the loopback IP of the PE router connected to that VPN

How to Add a New MPLS-based VPN

-
- Step 1** From the shell prompt, type the following command: `p3vpn --add --vpn=VPN-Name --mpls-vpn=RT@PE, (RT@PE2, RT@PE3,`
-

How to Remove a VPN

-
- Step 1** From the shell prompt, type the following command: `p3vpn --remove --vpn= VPN-Name`
-

How to Display VPN Information

- [To List All Existing VPNs, page 13-25](#)
- [To List All Subscribers for a Specified VPN, page 13-25](#)
- [To Display the Mappings for a Specified VPN, page 13-26](#)

To List All Existing VPNs

-
- Step 1** From the shell prompt, type the following command: `p3vpn --show-all l`
-

To List All Subscribers for a Specified VPN

- [Listing All Subscribers for a Specified VPN: Example, page 13-26](#)

-
- Step 1** From the shell prompt, type the following command: `p3vpn --show-sub --vpn= VPN-Name l`
-

Listing All Subscribers for a Specified VPN: Example

```
p3vpn -show-sub --vpn=vpn1
sub1: 10.1.1.0/24@vpn1
sub2: 20.1.1.0/24@vpn1
Command terminated successfully
```

To Display the Mappings for a Specified VPN

Step 1 From the shell prompt, type the following command: `p3vpn --show --vpn= VPN-Name1`

Listing All Subscribers for a Specified VPN: Example

```
p3vpn --show --vpn=vpn1
Name:          vpn1
Domain:        subscribers
Mappings:
MPLS/VPN: 1:1000@10.0.0.1      (no BGP information)
MPLS/VPN: 1:1000@10.0.0.2      label: 10 IP range: 1.1.1.1/32
Command terminated successfully
```

How to Manage VPN Mappings**To Remove All Existing Mappings from a Specified VPN**

Step 1 From the shell prompt, type the following command: `p3vpn --remove-all-mappings --vpn= VPN-Name`

To Remove a Specified Mapping from a Specified VPN

Step 1 From the shell prompt, type the following command: `p3vpn --remove-mappings --vpn= VPN-Name--mpls-vpn= RT@PE, (RT@PE2, RT@PE3, ...)`

How to Add Mappings to VPN-based Subscribers

There are three types of mappings that can be added to an existing VPN-based subscriber:

- A set of IP addresses defined as IP@VPN
- A complete VPN (this is actually a special case of IP@VPN mappings, in which the mapping is defined as 0.0.0.0/0@VPN)
- All the IP addresses of a CE router, defined by a AS:value@VPN-NAME (BGP community)

How to Add IP Address Mappings

Options

The following options are available

- **SUB-NAME** — The name of the subscriber to be associated with the specified community attribute
- **IP1[/RANGE][,...]@VPN-NAME** — IP address or addresses to assign to the VPN
 - **IP** = the IP address. This may be any of the following
 - a single IP address (x.x.x.x)
 - a single range of IP addresses (x.x.x.x/y)
 - a list of IP addresses separated by commas (x.x.x.x, y.y.y.y, z.z.z.z)
 - a list of IP address ranges (x.x.x.x/a, y.y.y.y/b, z.z.z.z/c)
 - **VPN-NAME** = name of the VPN to which the community attribute will be assigned
- **--additive-mappings** — Use this option to add the new mapping(s) to any existing ones. (Without this option, any existing mappings are overwritten.)

Step 1 From the shell prompt, type the following command: `p3subs -add --subscriber= SUB-NAME--ip=IP1[/RANGE][,...]@VPN-NAME [--additive-mappings]`

How to Add VPN-based Mappings

This option is supported to provide backwards compatibility with MPLS/VPN-based subscribers in releases before 3.1.5.

Options

The following options are available

- **SUB-NAME** — The name of the subscriber to be associated with the specified community attribute
- **VPN-NAME** — The name of the VPN to which the subscriber will be mapped. (This option is equivalent to defining the mapping as 0.0.0.0/0@VPN)
- **--additive-mappings** — Use this option to add the new mapping(s) to any existing ones. (Without this option, any existing mappings are overwritten.)

Step 1 From the shell prompt, type the following command: `p3subs -add --subscriber SUB-NAME --vpn=VPN-NAME [--additive-mappings]`

How to Configure the Community Parameter

An optional parameter may be set defining a community attribute. The community attribute provides a mechanism for defining the BGP community as one subscriber, using the *community@VPN* specification.

The community attribute in the BGP protocol is used to dynamically map IP ranges to subscribers. The community attribute can be configured in the Provider Edge (PE) router or in the Customer Edge (CE) router.

The *community@VPN* specification is replaced by an *IP@VPN* specification by the BGP LEG.

Use the **p3subs** utility to configure the community parameter.

Options

The following options are available:

- **SUB-NAME** — The name of the subscriber to be associated with the specified community attribute
- **AS:value@VPN-NAME** — The community attribute to assign to the VPN
 - **AS** = autonomous system. Integer in the range 0-65535 assigned by the network administrator
 - **value** = the community attribute. Integer in the range 0-65535 assigned by the network administrator
 - **VPN-NAME** = name of the VPN to which the community attribute will be assigned

Step 1 From the shell prompt, type the following command: `p3subs -add --subscriber= SUB-NAME--community= AS:value@VPN-NAME`

How to Remove VPN Mappings from Subscribers

- [To Remove All Existing Mappings from a Specified Subscriber, page 13-28](#)
- [To Remove a Specified IP Mapping from a Specified Subscriber, page 13-28](#)
- [To Remove a Specified VPN Mapping from a Specified Subscriber, page 13-29](#)
- [To Remove a Specified Community-based Mapping from a Specified Subscriber, page 13-29](#)

To Remove All Existing Mappings from a Specified Subscriber

Step 1 From the shell prompt, type the following command: `p3subs --remove-all-mappings -subscriber= SUB-NAME`

To Remove a Specified IP Mapping from a Specified Subscriber

Step 1 From the shell prompt, type the following command: `p3psubs --remove-mappings -subscriber= SUB-NAME--ip= IP1[/RANGE] [, ...]@VPN-NAME`

To Remove a Specified VPN Mapping from a Specified Subscriber

-
- Step 1** From the shell prompt, type the following command: `p3psubs --remove-mappings -subscriber=
SUB-NAME--vpn= VPN-NAME`
-


To Remove a Specified Community-based Mapping from a Specified Subscriber

-
- Step 1** From the shell prompt, type the following command: `p3psubs --remove-mappings -subscriber=
SUB-NAME--community= AS:value@VPN-NAME`
-

How to Monitor Subscriber MPLS/VPN Mappings

Use the **p3subs** utility to manage VPNs.

-
- Step 1** From the shell prompt, type the following command: `p3subs --show-all-mappings --subscriber=
SUB-NAME`
-



CHAPTER 14

Managing the SCMP

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 14-1](#)
- [How to Configure the SCMP, page 14-8](#)
- [How to Monitor the SCMP Environment, page 14-15](#)

About SCMP

The Service Control Management Protocol (SCMP) is a protocol that integrates the SCE platform and the ISG (Intelligent Service Gateway) functionality of the Cisco routers, thereby providing a mechanism that allows the ISG and the SCE platform to manage subscriber sessions together without requiring coordination and orchestration by additional components.

- [SCMP Terminology, page 14-2](#)
- [Deployment Scenarios, page 14-3](#)
- [SCMP Peer Devices, page 14-7](#)
- [SCMP Subscriber Management, page 14-8](#)

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 SCE platform represents a session in the SCMP peer (as defined by the ISG terminology).

Connection management

The SCE platform initiates the connection to the peer device. On SCMP connection establishment, the 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 SCE platform:

- The SCMP peer provisions the session to the SCE platform when it is created in the peer device (push)
- The 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 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 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 SCE platform terminology. It is derived from the ISG terminology, since every 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 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

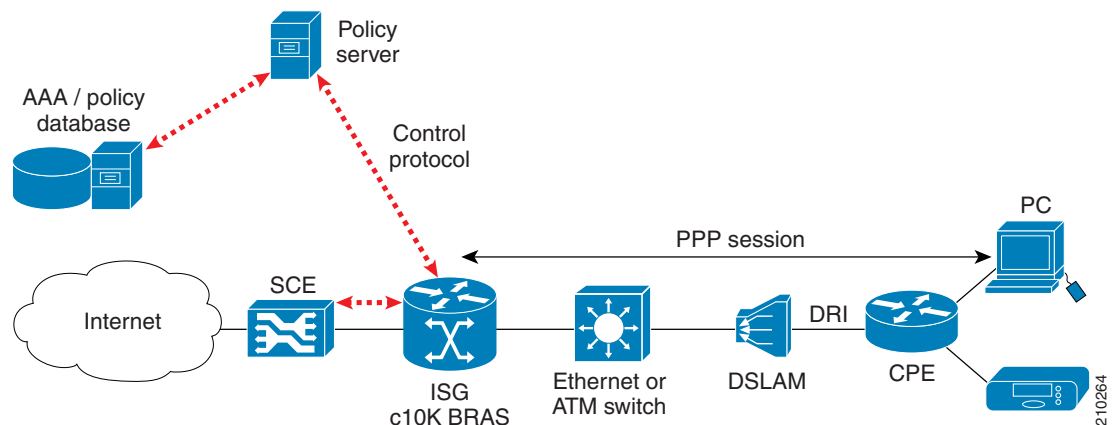
The following sections illustrate the basic types of SCMP deployment scenarios.

- 1xISG – 1xSCE
- 1xISG – 2xSCE (SCE cascade)
- NxISG – 2xSCE (SCE cascade)
- NxISG – MxSCE Via Load Balancing (MGSCP)

Single ISG Router with a Single SCE Platform (1xISG – 1xSCE)

The following diagram illustrates a deployment using one ISG router with a single SCE platform.

Figure 14-1 Single ISG Router with a Single SCE Platform



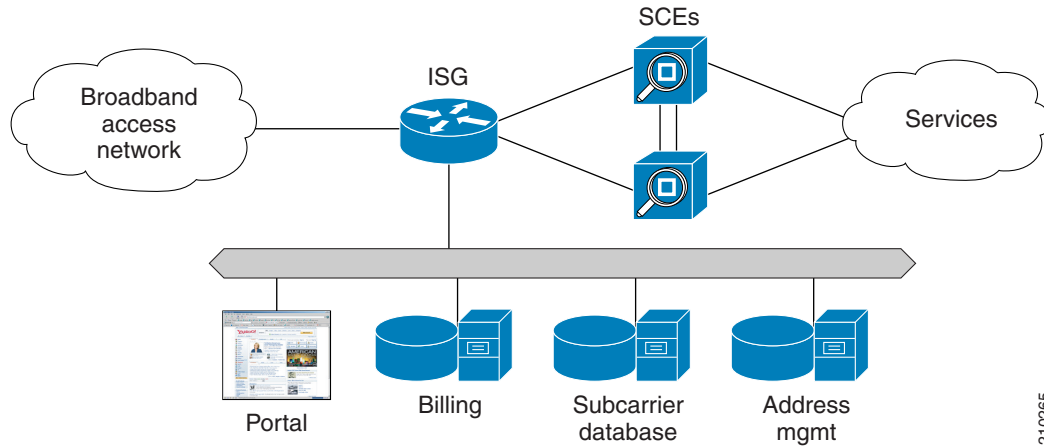
Note the following:

- 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 SCE platform results in a single point of failure, which is not generally acceptable in an actual deployment.

Single ISG Router with Two Cascaded SCE Platforms (1xISG – 2xSCE)

The following diagram illustrates a deployment using one ISG router with two cascaded SCE platforms.

Figure 14-2 *Single ISG Router with Two Cascaded 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 SCE platform has been added to provide redundancy. This redundancy scheme assumes that SCE platforms are connected in a cascade, with one active SCE platform and one backup.

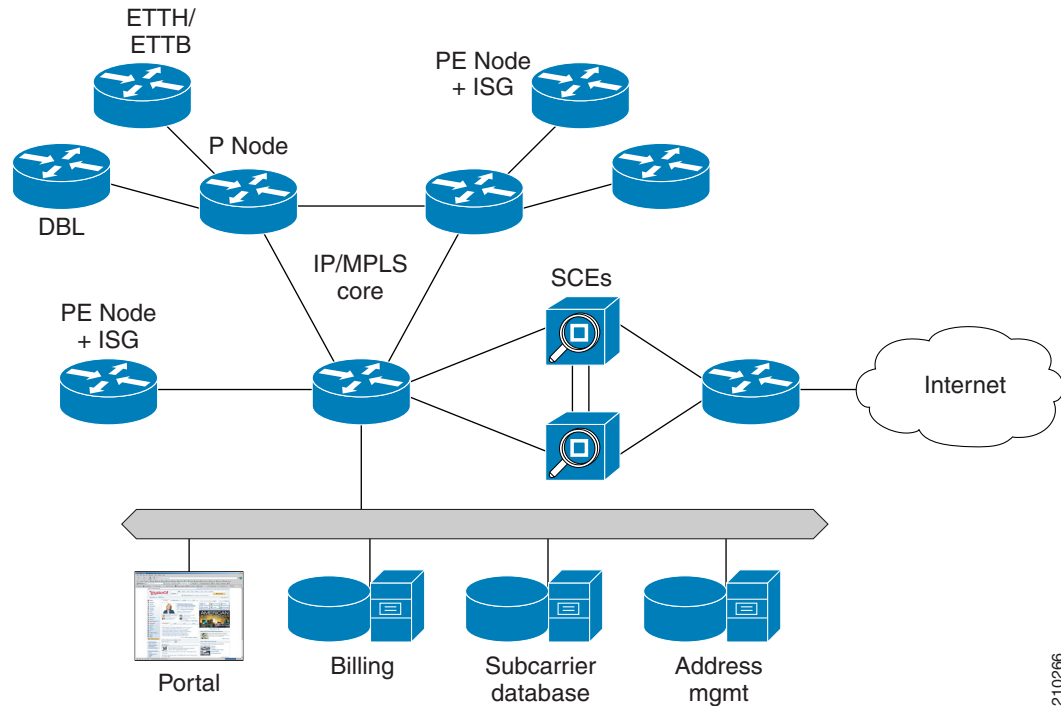
Please note the following:

- When cascaded SCE platforms are connected to one or more ISG devices, only the active SCE platform maintains a connection to the ISG devices.
- You can configure the cascaded SCE platforms to receive session info from the SCMP peer on session creation or pull the session info when the subscribers traffic traverses the SCE platform.
- An ISG device cannot push sessions to two SCE platforms at the same time

Multiple ISG Routers with Two Cascaded SCE Platforms (NxISG – 2xSCE)

The following diagram illustrates a deployment using multiple ISG routers with two cascaded SCE platforms.

Figure 14-3 Multiple ISG Routers with Two Cascaded SCE Platforms



210266


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 SCE platforms centrally, just before traffic requiring such services exits the SP network, since not all traffic needs to be processed by SCE platforms. Please note the following:

- When cascaded SCE platforms are connected to one or more ISG devices, only the active SCE platform maintains a connection to the ISG devices.
- You can configure the cascaded SCE platforms to receive session info from the SCMP peer on session creation or pull the session info when the subscribers traffic traverses the SCE platform.
- An ISG device cannot push sessions to two SCE platforms at the same time.

Multiple ISG Routers with Multiple SCE Platforms via Load Balancing (NxISG – MxSCE)

The following diagram illustrates a deployment using multiple ISG routers with multiple SCE platforms via load balancing. This is the scenario required for a MGSCP deployment.

Figure 14-4 Multiple ISG Routers with Multiple SCE Platforms via Load Balancing



210267

This scenario includes several SCE platforms connected to the 7600/6500 switch. For efficient control of subscriber flows, the same SCE platform must process both directions of each subscriber flow, since the SCE platform keeps the subscriber context. The 7600/6500 switch to which the SCE platforms are connected acts as a dispatching element, distributing subscriber flows between SCE platforms and guaranteeing that all flows of a specific subscriber will pass through the same 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 SCE platforms at the same time.
- You must configure multiple 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 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 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, page 9-5](#)).

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 SCE platform supports a maximum of 20 SCMP peer devices.

Connection Management

The SCMP attempts to maintain an open connection to each peer device.

The following figure illustrates the SCMP connection state functionality.

Figure 14-5 SCMP Connection State Functionality



The loss-of-sync timeout prevents the 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 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 SCE platform uses the GUID for all communication with the SCMP peer.

SCMP creates the 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.

How to Configure the SCMP

- [How to Configure SCMP Parameters, page 14-9](#)
- [How to Add an SCMP Peer Device, page 14-11](#)
- [How to Delete Subscribers Managed by an SCMP Peer Device, page 14-13](#)
- [How to Delete an SCMP Peer Device, page 14-13](#)
- [How to Define the Subscriber ID, page 14-13](#)
- [How to Configure the RADIUS Client, page 14-14](#)

How to Configure SCMP Parameters

You can configure the following options for the SCMP:

- Enable the SCMP
- Configure the SCMP peer device to push sessions to the SCE platform
- Allow the SCMP peer device to provision each subscriber to only one 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.

Step 1 From the SCE(config)# prompt, type `scmp` and press Enter.

How to Disable the SCMP

Step 1 From the SCE(config)# prompt, type `no scmp` and press Enter.

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 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).

Step 1 From the SCE(config)# prompt, type `scmp subscriber send-session-start` and press Enter.

How to Disable Pushing Sessions

Use this command to disable pushing sessions to the SCE platform. This means that the SCE platform will pull all sessions from the SCMP peer.

Step 1 From the SCE(config)# prompt, type `no scmp subscriber send-session-start` and press Enter.

How to Configure the SCMP Peer Device to Force Each Subscriber to Single 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 SCE platform.

Use this command to configure the SCMP peer device to verify that each subscriber is provisioned to only one SCE platform. If a subscriber was provisioned to a different SCE platform, the SCMP removes it from the previous SCE platform and provisions it to the new SCE platform. This configuration is required in MGSCP topology where, if a fail-over between SCE platforms, subscribers might move from one SCE platform to another. If transferred subscribers are not cleared from the previous SCE platform, it can cause capacity issues.

Use the **no** form of the command to allow SCMP to provision subscribers to more than one SCE platform.

This configuration takes effect only after the connection is re-established.

Default is disabled (subscribers can be provisioned to more than one SCE platform).

Step 1 From the SCE(config)# prompt, type `scmp subscriber force-single-sce` and press Enter.

How to Disable Forcing Each Subscriber to Single SCE Platform

Use this command to disable forcing each subscriber to only one SCE platform. This allow subscribers to be provisioned to more than one SCE platform.

Step 1 From the SCE(config)# prompt, type `no scmp subscriber force-single-sce` and press Enter.

How to Define 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 SCE platform to the SCMP peer device in seconds
 - Default = 5 seconds

Step 1 From the SCE(config)# prompt, type `scmp keepalive-interval interval` and press Enter.

How to Define the Reconnect Interval Parameter

The reconnect interval is the amount of time between attempts by the SCE platform to reconnect with an SCMP peer. The 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 SCE platform to reconnect with an SCMP peer, in seconds
 - Default = 30 seconds

Step 1 From the SCE(config)# prompt, type `scmp reconnect-interval interval` and press Enter.

How to Define the Loss-of-Sync Timeout Parameter

The loss of sync timeout interval is the amount of time between loss of connection between the 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

Step 1 From the SCE(config)# prompt, type `scmp loss-of-sync-timeout interval` and press Enter.

How to Add an SCMP Peer Device

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

-
- Step 1** From the SCE(config)# prompt, type `scmp name peer_device_name radius radius_hostname secret shared_secret [auth-port auth-portnumber acct-port acct-portnumber]` and press Enter.
-

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

-
- Step 1** From the SCE(config if)# prompt, type `subscriber anonymous-group name group-name IP-range range [template template] scmp name peer-device-name` and press Enter.
-

How to Remove an Anonymous Group from the SCMP Peer Device

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.

-
- Step 1** From the SCE(config if)# prompt, type `no subscriber anonymous-group name group-name` and press Enter.
-

How to Delete 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

Step 1 From the SCE(config if)# prompt, type `no subscriber scmp name peer-device-nameall` and press Enter.

How to Delete 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
```

How to Define 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
```

How to Configure 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

Step 1 From the SCE(config)# prompt, type `ip radius-client retry limit times[timeout timeout]` and press Enter.

How to Monitor the SCMP Environment

You can monitor the following components of the SCMP environment:

- SCMP
- RADIUS client

How to Monitor the SCMP

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

Step 1 From the SCE>prompt, type `show scmp` and press Enter.

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

Step 1 From the SCE>prompt, type `show scmp all` and press Enter.

How to display the configuration for a specified SCMP peer device

Step 1 From the SCE>prompt, type `show scmp name device-name` and press Enter.

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

Step 1 From the SCE>prompt, type `show scmp all counters` and press Enter.

How to display the statistics for a specified SCMP peer device

Step 1 From the SCE>prompt, type `show scmp name peer_device_name counters` and press Enter.

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
```

How to Monitor the RADIUS Client

Use the following command to monitor the SCMP RADIUS client. This command displays the general configuration of the RADIUS client.

Step 1 From the SCE>prompt, type `show ip radius-client` and press Enter.



APPENDIX **A**

Monitoring SCE Platform Utilization

This module explains how to monitor SCE platforms that are installed in real traffic.

- [SCE Platform Utilization Indicators, page A-2](#)
- [Service Loss, page A-3](#)

As with any network device, the SCE platform has its performance and capacity envelopes. As the network evolves, the utilization of the SCE platform can increase and these envelopes might be reached. It is, therefore, advisable to monitor SCE platform to be sure that utilization remains at a level that supports reliable and consistent service.

When the 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 SCE platform will automatically filter all UDP traffic in certain cases as a last resort).

Monitoring the SCE platform can be divided onto two main areas:

- Monitoring SCE platform utilization
- Monitoring service loss

SCE Platform Utilization Indicators

The 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
- Flows capacity
- Subscribers capacity

CPU Utilization

- **SNMP**

tpCpuUtilization , available for each Traffic Processor (*TpInfoEntry*) in the *PCubeSeMib.mib* .

- **CLI command**

```
show snmp MIB pcube-SE-MIB traffic-processor | include tpCpuUtilization
```

It is advisable to consider sizing of the solution when the CPU utilization exceeds 75% regularly at peak hours.

Flows Capacity

- **SNMP**

tpFlowsCapacityUtilization , available for each Traffic Processor (*TpInfoEntry*) in the *PCubeSeMib.mib*

- **CLI command**

```
show snmp MIB pcube-SE-MIB traffic-processor | include tpFlowsCapacityUtilization
```

It is advisable to consider sizing of the solution when the flows capacity utilization exceeds 90% regularly at peak hours.

Subscribers Capacity

- **SNMP**

subscribersInfoEntry , available in the *PCubeSeMib.mib*

- **CLI command**

```
show snmp MIB pcube-SE-MIB subscriber
```

The SCE 2000 platform supports up to 80K subscribers, while the SCE 1000 supports up to 40K 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 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 SCE platform:

- Temporary – This might occur when some network pattern which is short in its nature occurred and caused the SCE platform to exhaust some of its resources temporarily. An example could be a DDoS attack that the SCE platform could not detect and filter.

This is usually measured in seconds.

- Permanent – In cases where the 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 SCE platform.

Monitoring Service Loss

- SNMP

tpServiceLoss MIB available for each traffic processor (*TpInfoEntry*)

It is expected that the 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.



APPENDIX **B**

Proprietary MIB Reference

This appendix describes the *pcube* proprietary MIB supported by the SCE platform. A MIB (Management Information Base) is a database of objects that can be monitored by a network management system (NMS). The SCE platform supports both the standard MIB-II and a proprietary Service Control Enterprise MIB. This proprietary *pcube* MIB enables the external management system to perform configuration, performance, troubleshooting and alerting operations specific to the SCE platform, and therefore not provided by the standard MIB.



Note

Information and proprietary MIB files supported by the SCOS can be downloaded from:
<http://www.cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml>
under the Cisco Service Routing Products section.

- [pcube Enterprise MIB](#), page B-2
- [pcubeModules \(1.3.6.1.4.1.5655.2\)](#), page B-6
- [pcubeWorkgroup \(1.3.6.1.4.1.5655.4\)](#), page B-16
- [Supported Standards](#), page B-95

pcube Enterprise MIB

The **pcube** Enterprise MIB splits into four main groups:

- Products
- Modules
- Management
- Workgroup

The **pcube** enterprise tree structure is defined in a MIB file named **PCUBE-SMI.my**.



Note

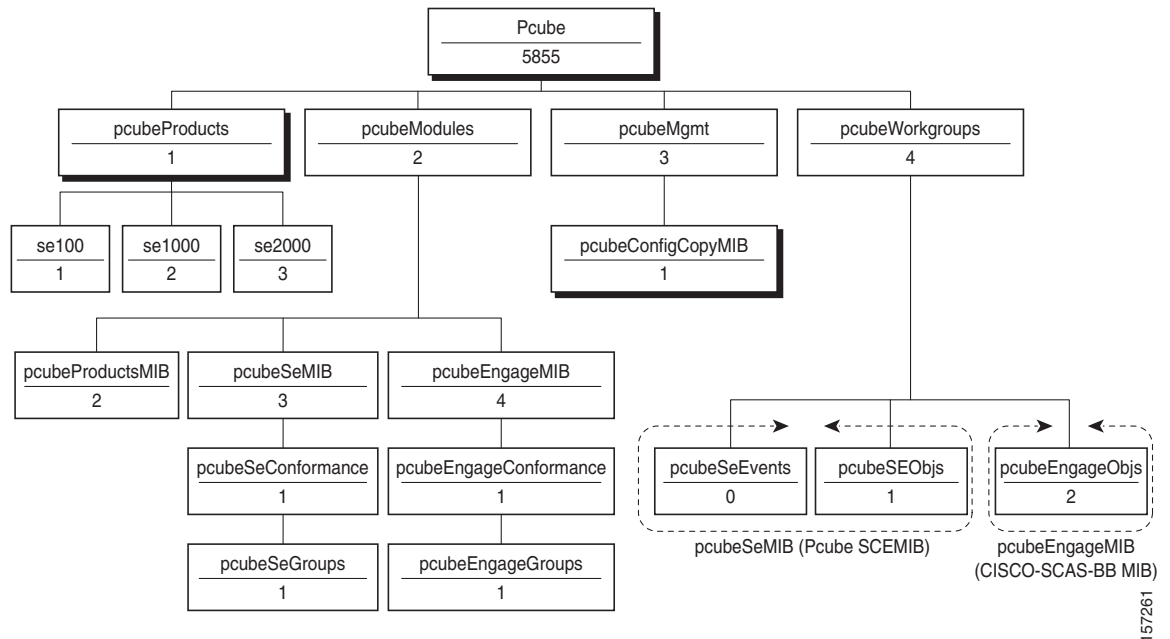
The following object identifier represents the **pcube** Enterprise MIB:
1.3.6.1.4.1.5655, or *iso.org.dod.internet.private.enterprise.pcube*.

The figure below illustrates the **pcube** Service Control Enterprise MIB structure.

Conventions used in the diagram:

- Dotted arrows surrounding a unit or units indicate that the component is described in the MIB file specified below the line.
- A shadowed box indicates that the component is described in its own MIB file.

Figure B-1 Cisco Service Control MIB Structure



The **pcubeProducts** subtree

The **pcubeProducts** subtree contains the OIDs of Cisco Service Control products. These OIDs are used only to describe the Cisco Service Control platforms, not as roots for other OIDs.

This subtree does not contain online data, just global definitions.

The **pcubeProductsMIB** is defined the MIB file: **PCUBE-PRODUCTS-MIB.my**

The pcubeModules subtree

The **pcubeModules** subtree provides a root object identifier under which MIB modules can be defined.

This subtree does not contain online data, just global definitions.

The pcubeMgmt subtree

The **pcubeMgmt** subtree is the root for **pcube** MIBs that are relevant to multiple products.

The pcubeWorkgroups subtree

The **pcubeWorkgroups** subtree contains the actual MIBs for Cisco Service Control devices and sub-devices. This is where the SCE platform, the dispatcher, and also the SCA BB application, hook their actual OIDs and notifications. This means that the branches in this subtree are defined in multiple MIB files.

The pcubeConfigCopyMib subtree

The **pcubeConfigCopy** MIB is a subset of the Cisco Config-Copy-MIB ported to the **pcube** enterprise subtree. It supports only local copying of running config to startup config.

The **pcubeConfigCopyMIB** is defined the MIB file: **PCUBE-CONFIG-COPY-MIB.my**

The config copy MIB is intended for use by all **pcube** products, and is therefore placed under the **pcubeMgmt** subtree.

The pcubeSE MIB

The **pcubeSE** MIB comprises two branches:

| | |
|----------------------|--|
| pcubeSeEvents | Contains the OIDs used for sending enterprise-specific notifications. |
| pcubeSEObjs | Contains the OIDs that belong to the SCE platform, divided into groups according to functionality. |

Application MIB Integration

The *pcubeSeMIB* described above contains information about the SCE platform. Information regarding the SML application running on the SCE platform is available through the following two interfaces:

- *application* and *subscriber* groups of the *pcubeSeMIB*
- *pcubeEngageMIB*

Application and Subscriber groups

These groups allow the operator to view information exposed by the application. In essence, these groups contain tables that allow access to viewables and tuneables defined in the installed SML application. This mechanism is general, and does not require specific adjustments of the SCOS platform for an application. Any installed application can define variables that would then automatically be accessible through the application or subscriber groups.

- [The Application Group, page B-4](#)
- [The Subscriber Group, page B-5](#)

The Application Group

The application group contains information about the application itself. It also contains two tables:

- **appPropertiesTable:**

Contains a list of properties of the application. These properties are SML viewables and tuneables that are defined in global scope. Each property has a name and a type (both are strings).

This table is updated when the application is installed, and has read-only access. It is therefore managed by the agent (the SCOS).
- **appPropertiesValueTable:**

Contains information about the application properties that appear in the *appPropertiesTable* , described above. For each such property, a value is available in one or more formats, as appropriate:

 - String
 - Integer
 - 64-bit Integer.
- This table is a read-write table, which uses the row-status mechanism. With this mechanism it is possible for a management station to add rows to an SNMP table in an pseudo-atomic way. The RowStatus type is defined in *SNMPv2-SMI* .

This table is cleared when an application is unloaded, or when the system resets. The manager may add rows to the table after the system starts, with properties taken from the *appPropertiesTable* , and then it is possible to poll these variables for their values. It is the manager's responsibility to read the value in the correct format (according to the "type" field in the *appPropertiesTable*).

Note that the values cannot be changed through SNMP. Tuneables and Viewables are the same in this respect, they can only be viewed.

Note as well that all the properties in the application group are global properties.

The Subscriber Group

The subscribers group contains global statistics about the subscribers. It also contains two tables, which work in a mechanism very similar to that of the application properties tables described above:

- subscribersPropertiesTable:

Contains a list of “subscriber properties” that the installed application has defined. These properties are SML viewables and tuneables that are defined in subscriber scope. Each property has a name and a type (both are strings).

This table is updated when the application is installed, and has read-only access. It is therefore it is managed by the agent (the SCOS).

- subscribersPropertiesValueTable:

Like the *appPropertiesValueTable*, this table also uses the RowStatus mechanism to allow addition and removal of rows by the manager. Once a row is created in the *subscriberValueTable*, it is possible to poll the value of the property for the requested subscriber. The values are given in three formats (string, integer, and 64-bit integer, it being the responsibility of the manager to read the appropriate field, according to the type of the property)

In order to retrieve the value of a subscriber property, the manager must supply the following:

- The name of the property (taken from the *subscriberPropertyTable*)
 - The name of the subscriber for which this value should be polled (the name of the subscriber should be known to the manager, it is not available through SNMP).
- Note that the values of the properties are read-only for viewables and tuneables alike.

This table is also used by the *pcubeEngageMIB*.

The Engage MIB (pcubeEngageMIB)

The application group of the *pcubeSeMIB* allows the manager to retrieve information about the application. However, due to its generality, it requires quite a bit of intelligence from the management station. The *pcubeEngageMIB* is an application-specific MIB that supplies information relevant only to the SCA BB (Engage) application.

All information that the application may want to publish is stored in viewables and tuneables. However, there is currently no automatic way for the application to declare that a certain OID should reflect a certain application variable. For this reason, the *Engage* application MIB is actually implemented by the SCOS.

The *Engage* subscriber tables rely on the SCOS subscriber value tables described above. Since there is no table that contains a list of the subscribers in the system, the only subscribers referred to in the *Engage* tables are the ones that appear in the SCOS *subscribersPropertiesValueTable*. The *spvIndex* is used to index the subscribers in the *pcubeEngageMIB* as well.

Using this Reference

This reference describes all existing MIB objects. For each object, the following information is presented:

| | |
|-------------|--|
| IDENTIFIERS | MIB descriptor name and OID |
| DESCRIPTION | Description of the object, including format and legal values, if applicable. |
| ACCESS | Access control associated with the object: <ul style="list-style-type: none"> • Read only (RO) • Read/Write (RW) |
| SYNTAX | The general format of the object. |

pcubeModules (1.3.6.1.4.1.5655.2)

pcubeModules provides a root object identifier from which module identity values may be assigned.

Modules:

- PCUBE-SE-MIB: pcubeSeMIB (1.3.6.1.4.1.5655.2.3)
- CISCO-SCAS-BB-MIB: pcubeEngageMIB (1.3.6.1.4.1.5655.2.4) (See the "SCA BB Proprietary MIB Reference" chapter in the *Cisco SCA BB Reference Guide* for a description of the CISCO-SCAS-BB-MIB.)
- [pcubeSeMIB \(1.3.6.1.4.1.5655.2.3\)](#), page B-6

pcubeSeMIB (1.3.6.1.4.1.5655.2.3)

Main SNMP MIB for the Cisco SCE products such as SCE 2000 and SCE 1000. This MIB provides configuration and runtime status for chassis, control modules, and line modules on the SCOS systems.

- [pcubeSeMIB Object Groups \(1.3.6.1.4.1.5655.2.3.1.1\)](#), page B-7
- [pcubeCompliances \(1.3.6.1.4.1.5655.2.3.1.2\)](#), page B-15

pcubeSeMIB Object Groups (1.3.6.1.4.1.5655.2.3.1.1)

Following is a list of the pcubeSeMIB object groups. Each object group consists of a number of related objects. All individual objects are listed and described in the pcubeWorkgroup section.

| | |
|--------------------------------|-------------------------------|
| pcubeSystemGroup | {1.3.6.1.4.1.5655.2.3.1.1.1} |
| pcubeChassisGroup | {1.3.6.1.4.1.5655.2.3.1.1.2} |
| pcubeModuleGroup | {1.3.6.1.4.1.5655.2.3.1.1.3} |
| pcubeLinkGroup | {1.3.6.1.4.1.5655.2.3.1.1.4} |
| pcubeDiskGroup | {1.3.6.1.4.1.5655.2.3.1.1.5} |
| pcubeRdrFormatterGroup | {1.3.6.1.4.1.5655.2.3.1.1.6} |
| pcubeLoggerGroup | {1.3.6.1.4.1.5655.2.3.1.1.7} |
| pcubeSubscribersGroup | {1.3.6.1.4.1.5655.2.3.1.1.8} |
| pcubeTrafficProcessorGroup | {1.3.6.1.4.1.5655.2.3.1.1.9} |
| pcubePortGroup | {1.3.6.1.4.1.5655.2.3.1.1.10} |
| pcubeTxQueuesGroup | {1.3.6.1.4.1.5655.2.3.1.1.11} |
| pcubeGlobalControllersGroup | {1.3.6.1.4.1.5655.2.3.1.1.12} |
| pcubeApplicationGroup | {1.3.6.1.4.1.5655.2.3.1.1.13} |
| pcubeTrafficCountersGroup | {1.3.6.1.4.1.5655.2.3.1.1.14} |
| pcubeAttackGroup | {1.3.6.1.4.1.5655.2.3.1.1.15} |
| pcubeVasTrafficForwardingGroup | {1.3.6.1.4.1.5655.2.3.1.1.16} |
| pcubeMplsVpnAutoLearnGroup | {1.3.6.1.4.1.5655.2.3.1.1.17} |
| pcubeTrapObjectsGroup | {1.3.6.1.4.1.5655.2.3.1.1.18} |

- [System Group: pcubeSystemGroup \(1.3.6.1.4.1.5655.2.3.1.1.1\)](#), page B-8
- [Chassis Group: pcubeChassisGroup \(1.3.6.1.4.1.5655.2.3.1.1.2\)](#), page B-8
- [Module Group: pcubeModuleGroup \(1.3.6.1.4.1.5655.2.3.1.1.3\)](#), page B-8
- [Link Group: pcubeLinkGroup \(1.3.6.1.4.1.5655.2.3.1.1.4\)](#), page B-9
- [Disk Group: pcubeDiskGroup \(1.3.6.1.4.1.5655.2.3.1.1.5\)](#), page B-9
- [RDR Formatter Group: pcubeRdrFormatterGroup \(1.3.6.1.4.1.5655.2.3.1.1.6\)](#), page B-9
- [Logger Group: pcubeLoggerGroup \(1.3.6.1.4.1.5655.2.3.1.1.7\)](#), page B-10
- [Subscribers Group: pcubeSubscribersGroup \(1.3.6.1.4.1.5655.2.3.1.1.8\)](#), page B-10
- [Traffic Processor Group: pcubeTrafficProcessorGroup \(1.3.6.1.4.1.5655.2.3.1.1.9\)](#), page B-11
- [Port Group: pcubePortGroup \(1.3.6.1.4.1.5655.2.3.1.1.10\)](#), page B-12
- [Transmit Queues Group: pcubeTxQueuesGroup \(1.3.6.1.4.1.5655.2.3.1.1.11\)](#), page B-13
- [Global Controllers Group: pcubeGlobalControllersGroup \(1.3.6.1.4.1.5655.2.3.1.1.12\)](#), page B-13
- [Application Group: pcubeApplicationGroup \(1.3.6.1.4.1.5655.2.3.1.1.13\)](#), page B-13
- [Traffic Counters Group: pcubeTrafficCountersGroup \(1.3.6.1.4.1.5655.2.3.1.1.14\)](#), page B-14
- [Attack Group: pcubeAttackGroup \(1.3.6.1.4.1.5655.2.3.1.1.15\)](#), page B-14

- VAS Group: `pcubeVasTrafficForwardingGroup` 1.3.6.1.4.1.5655.2.3.1.1.16), page B-14
- MPLS/VPN Group: `pcubeMplsVpnAutoLearnGroup` 1.3.6.1.4.1.5655.2.3.1.1.17), page B-14
- Traps Group: `pcubeTrapObjectsGroup` 1.3.6.1.4.1.5655.2.3.1.1.18), page B-15

System Group: `pcubeSystemGroup` (1.3.6.1.4.1.5655.2.3.1.1.1)

The System group provides data on the system-wide functionality of the SCE platform.

Objects:

1. `sysOperationalStatus` (1.3.6.1.4.1.5655.4.1.1.1), page B-30
2. `sysFailureRecovery` (1.3.6.1.4.1.5655.4.1.1.2), page B-30
3. `sysVersion` (1.3.6.1.4.1.5655.4.1.1.3), page B-31

Chassis Group: `pcubeChassisGroup` (1.3.6.1.4.1.5655.2.3.1.1.2)

The Chassis group defines and identifies the chassis, as well as environmental alarms related to the chassis.

Objects:

1. `pchassisSysType` (1.3.6.1.4.1.5655.4.1.2.1), page B-31
2. `pchassisPowerSupplyAlarm` (1.3.6.1.4.1.5655.4.1.2.2), page B-31
3. `pchassisFansAlarm` (1.3.6.1.4.1.5655.4.1.2.3), page B-31
4. `pchassisTempAlarm` (1.3.6.1.4.1.5655.4.1.2.4), page B-32
5. `pchassisVoltageAlarm` (1.3.6.1.4.1.5655.4.1.2.5), page B-32
6. `pchassisNumSlots` (1.3.6.1.4.1.5655.4.1.2.6), page B-32
7. `pchassisSlotConfig` (1.3.6.1.4.1.5655.4.1.2.7), page B-33
8. `pchassisPsuType` (1.3.6.1.4.1.5655.4.1.2.8), page B-33
9. `pchassisLineFeedAlarm` (1.3.6.1.4.1.5655.4.1.2.9), page B-33

Module Group: `pcubeModuleGroup` (1.3.6.1.4.1.5655.2.3.1.1.3)

The Module group identifies and defines the modules, or cards, in the SCE platform.

Objects:

1. `pmoduleIndex` (1.3.6.1.4.1.5655.4.1.3.1.1.1), page B-34
2. `pmoduleType` (1.3.6.1.4.1.5655.4.1.3.1.1.2), page B-35
3. `pmoduleNumTrafficProcessors` (1.3.6.1.4.1.5655.4.1.3.1.1.3), page B-35
4. `pmoduleSlotNum` (1.3.6.1.4.1.5655.4.1.3.1.1.4), page B-35
5. `pmoduleHwVersion` (1.3.6.1.4.1.5655.4.1.3.1.1.5), page B-36
6. `pmoduleNumPorts` (1.3.6.1.4.1.5655.4.1.3.1.1.6), page B-36
7. `pmoduleNumLinks` (1.3.6.1.4.1.5655.4.1.3.1.1.7), page B-36
8. `pmoduleConnectionMode` (1.3.6.1.4.1.5655.4.1.3.1.1.8), page B-36
9. `pmoduleSerialNumber` (1.3.6.1.4.1.5655.4.1.3.1.1.9), page B-37
10. `pmoduleUpStreamAttackFilteringTime` (1.3.6.1.4.1.5655.4.1.3.1.1.10), page B-37
11. `pmoduleUpStreamLastAttackFilteringTime` (1.3.6.1.4.1.5655.4.1.3.1.1.11), page B-37
12. `pmoduleDownStreamAttackFilteringTime` (1.3.6.1.4.1.5655.4.1.3.1.1.12), page B-37

13. [pmoduleDownStreamLastAttackFilteringTime \(1.3.6.1.4.1.5655.4.1.3.1.1.13\)](#), page B-38
14. [pmoduleAttackObjectsClearTime \(1.3.6.1.4.1.5655.4.1.3.1.1.14\)](#), page B-38
15. [pmoduleAdminStatus \(1.3.6.1.4.1.5655.4.1.3.1.1.15\)](#), page B-38
16. [pmoduleOperStatus \(1.3.6.1.4.1.5655.4.1.3.1.1.16\)](#), page B-38

Link Group: pcubeLinkGroup (1.3.6.1.4.1.5655.2.3.1.1.4)

The Link group defines and identifies the link. It provides information regarding the mode of operation of the link defined for each status of the platform.

Objects:

1. [linkModuleIndex \(1.3.6.1.4.1.5655.4.1.4.1.1.1\)](#), page B-39
2. [linkIndex \(1.3.6.1.4.1.5655.4.1.4.1.1.2\)](#), page B-40
3. [linkAdminModeOnActive \(1.3.6.1.4.1.5655.4.1.4.1.1.3\)](#), page B-40
4. [linkAdminModeOnFailure \(1.3.6.1.4.1.5655.4.1.4.1.1.4\)](#), page B-40
5. [linkOperMode \(1.3.6.1.4.1.5655.4.1.4.1.1.5\)](#), page B-41
6. [linkStatusReflectionEnable \(1.3.6.1.4.1.5655.4.1.4.1.1.6\)](#), page B-41
7. [linkSubscriberSidePortIndex \(1.3.6.1.4.1.5655.4.1.4.1.1.7\)](#), page B-41
8. [linkNetworkSidePortIndex \(1.3.6.1.4.1.5655.4.1.4.1.1.8\)](#), page B-42

Disk Group: pcubeDiskGroup (1.3.6.1.4.1.5655.2.3.1.1.5)

The Disk group provides data regarding the space utilization on the disk.

Objects:

1. [diskNumUsedBytes \(1.3.6.1.4.1.5655.4.1.5.1\)](#), page B-42
2. [diskNumFreeBytes \(1.3.6.1.4.1.5655.4.1.5.2\)](#), page B-42

RDR Formatter Group: pcubeRdrFormatterGroup 1.3.6.1.4.1.5655.2.3.1.1.6)

The RDR Formatter provides information regarding RDR Formatter destinations (Collection Managers), as well as RDR statistics.

Objects:

1. [rdrFormatterEnable \(1.3.6.1.4.1.5655.4.1.6.1\)](#), page B-42
2. [rdrFormatterDestIPAddr \(1.3.6.1.4.1.5655.4.1.6.2.1.1\)](#), page B-43
3. [rdrFormatterDestPort \(1.3.6.1.4.1.5655.4.1.6.2.1.2\)](#), page B-44
4. [rdrFormatterDestPriority \(1.3.6.1.4.1.5655.4.1.6.2.1.3\)](#), page B-44
5. [rdrFormatterDestStatus \(1.3.6.1.4.1.5655.4.1.6.2.1.4\)](#), page B-44
6. [rdrFormatterDestConnectionStatus \(1.3.6.1.4.1.5655.4.1.6.2.1.5\)](#), page B-45
7. [rdrFormatterDestNumReportsSent \(1.3.6.1.4.1.5655.4.1.6.2.1.6\)](#), page B-45
8. [rdrFormatterDestNumReportsDiscarded \(1.3.6.1.4.1.5655.4.1.6.2.1.7\)](#), page B-45
9. [rdrFormatterDestReportRate \(1.3.6.1.4.1.5655.4.1.6.2.1.8\)](#), page B-45
10. [rdrFormatterDestReportRatePeak \(1.3.6.1.4.1.5655.4.1.6.2.1.9\)](#), page B-46
11. [rdrFormatterDestReportRatePeakTime \(1.3.6.1.4.1.5655.4.1.6.2.1.10\)](#), page B-46
12. [rdrFormatterNumReportsSent \(1.3.6.1.4.1.5655.4.1.6.3\)](#), page B-46

13. [rdrFormatterNumReportsDiscarded \(1.3.6.1.4.1.5655.4.1.6.4\)](#), page B-46
14. [rdrFormatterClearCountersTime \(1.3.6.1.4.1.5655.4.1.6.5\)](#), page B-47
15. [rdrFormatterReportRate \(1.3.6.1.4.1.5655.4.1.6.6\)](#), page B-47
16. [rdrFormatterReportRatePeak \(1.3.6.1.4.1.5655.4.1.6.7\)](#), page B-47
17. [rdrFormatterReportRatePeakTime \(1.3.6.1.4.1.5655.4.1.6.8\)](#), page B-47
18. [rdrFormatterProtocol \(1.3.6.1.4.1.5655.4.1.6.9\)](#), page B-48
19. [rdrFormatterForwardingMode \(1.3.6.1.4.1.5655.4.1.6.10\)](#), page B-48
20. [rdrFormatterCategoryIndex \(1.3.6.1.4.1.5655.4.1.6.11.1.1\)](#), page B-49
21. [rdrFormatterCategoryName \(1.3.6.1.4.1.5655.4.1.6.11.1.2\)](#), page B-49
22. [rdrFormatterCategoryNumReportsSent \(1.3.6.1.4.1.5655.4.1.6.11.1.3\)](#), page B-49
23. [rdrFormatterCategoryNumReportsDiscarded \(1.3.6.1.4.1.5655.4.1.6.11.1.4\)](#), page B-50
24. [rdrFormatterCategoryReportRate \(1.3.6.1.4.1.5655.4.1.6.11.1.5\)](#), page B-50
25. [rdrFormatterCategoryReportRatePeak \(1.3.6.1.4.1.5655.4.1.6.11.1.6\)](#), page B-50
26. [rdrFormatterCategoryReportRatePeakTime \(1.3.6.1.4.1.5655.4.1.6.11.1.7\)](#), page B-50
27. [rdrFormatterCategoryNumReportsQueued \(1.3.6.1.4.1.5655.4.1.6.11.1.8\)](#), page B-51
28. [rdrFormatterCategoryDestPriority \(1.3.6.1.4.1.5655.4.1.6.12.1.1\)](#), page B-51
29. [rdrFormatterCategoryDestStatus \(1.3.6.1.4.1.5655.4.1.6.12.1.2\)](#), page B-52

Logger Group: pcubeLoggerGroup (1.3.6.1.4.1.5655.2.3.1.1.7)

The Logger group is responsible for logging the system synchronous and asynchronous events.

Objects:

1. [loggerUserLogEnable \(1.3.6.1.4.1.5655.4.1.7.1\)](#), page B-52
2. [loggerUserLogNumInfo \(1.3.6.1.4.1.5655.4.1.7.2\)](#), page B-52
3. [loggerUserLogNumWarning \(1.3.6.1.4.1.5655.4.1.7.3\)](#), page B-53
4. [loggerUserLogNumError \(1.3.6.1.4.1.5655.4.1.7.4\)](#), page B-53
5. [loggerUserLogNumFatal \(1.3.6.1.4.1.5655.4.1.7.5\)](#), page B-53
6. [loggerUserLogClearCountersTime \(1.3.6.1.4.1.5655.4.1.7.6\)](#), page B-53

Subscribers Group: pcubeSubscribersGroup 1.3.6.1.4.1.5655.2.3.1.1.8)

The Subscribers group provides statistics concerning the number of subscribers and subscriber mappings. It also provides data on the subscriber properties and the value of those properties for a specified subscriber.

Objects:

1. [subscribersNumIntroduced \(1.3.6.1.4.1.5655.4.1.8.1.1.1\)](#), page B-54
2. [subscribersNumFree \(1.3.6.1.4.1.5655.4.1.8.1.1.2\)](#), page B-55
3. [subscribersNumIpAddrMappings \(1.3.6.1.4.1.5655.4.1.8.1.1.3\)](#), page B-55
4. [subscribersNumIpAddrMappingsFree \(1.3.6.1.4.1.5655.4.1.8.1.1.4\)](#), page B-55
5. [subscribersNumIpRangeMappings \(1.3.6.1.4.1.5655.4.1.8.1.1.5\)](#), page B-55
6. [subscribersNumIpRangeMappingsFree \(1.3.6.1.4.1.5655.4.1.8.1.1.6\)](#), page B-56

7. subscribersNumVlanMappings (1.3.6.1.4.1.5655.4.1.8.1.1.7), page B-56
8. subscribersNumVlanMappingsFree (1.3.6.1.4.1.5655.4.1.8.1.1.8), page B-56
9. subscribersNumActive (1.3.6.1.4.1.5655.4.1.8.1.1.9), page B-56
10. subscribersNumActivePeak (1.3.6.1.4.1.5655.4.1.8.1.1.10), page B-57
11. subscribersNumActivePeakTime (1.3.6.1.4.1.5655.4.1.8.1.1.11), page B-57
12. subscribersNumUpdates (1.3.6.1.4.1.5655.4.1.8.1.1.12), page B-57
13. subscribersCountersClearTime (1.3.6.1.4.1.5655.4.1.8.1.1.13), page B-57
14. subscribersNumTpIpRangeMappings (1.3.6.1.4.1.5655.4.1.8.1.1.14), page B-58
15. subscribersNumTpIpRangeMappingsFree (1.3.6.1.4.1.5655.4.1.8.1.1.15), page B-58
16. subscribersNumAnonymous (1.3.6.1.4.1.5655.4.1.8.1.1.16), page B-58
17. subscribersNumWithSessions (1.3.6.1.4.1.5655.4.1.8.1.1.17), page B-58
18. spIndex (1.3.6.1.4.1.5655.4.1.8.2.1.1), page B-59
19. spName (1.3.6.1.4.1.5655.4.1.8.2.1.2), page B-59
20. spType (1.3.6.1.4.1.5655.4.1.8.2.1.3), page B-60
21. spvSubName (1.3.6.1.4.1.5655.4.1.8.3.1.2), page B-61
22. spvPropertyName (1.3.6.1.4.1.5655.4.1.8.3.1.3), page B-61
23. spvRowStatus (1.3.6.1.4.1.5655.4.1.8.3.1.4), page B-61
24. spvPropertyStringValue (1.3.6.1.4.1.5655.4.1.8.3.1.5), page B-62
25. spvPropertyUintValue (1.3.6.1.4.1.5655.4.1.8.3.1.6), page B-62
26. spvPropertyCounter64Value (1.3.6.1.4.1.5655.4.1.8.3.1.7), page B-62

Traffic Processor Group: pcubeTrafficProcessorGroup 1.3.6.1.4.1.5655.2.3.1.1.9)

The Traffic Processor group provides statistics regarding the traffic flow handled by each traffic processor.

Objects:

1. tpModuleIndex (1.3.6.1.4.1.5655.4.1.9.1.1.1), page B-63
2. tpIndex (1.3.6.1.4.1.5655.4.1.9.1.1.2), page B-63
3. tpTotalNumHandledPackets (1.3.6.1.4.1.5655.4.1.9.1.1.3), page B-64
4. tpTotalNumHandledFlows (1.3.6.1.4.1.5655.4.1.9.1.1.4), page B-64
5. tpNumActiveFlows (1.3.6.1.4.1.5655.4.1.9.1.1.5), page B-64
6. tpNumActiveFlowsPeak (1.3.6.1.4.1.5655.4.1.9.1.1.6), page B-64
7. tpNumActiveFlowsPeakTime (1.3.6.1.4.1.5655.4.1.9.1.1.7), page B-65
8. tpNumTcpActiveFlows (1.3.6.1.4.1.5655.4.1.9.1.1.8), page B-65
9. TpNumTcpActiveFlowsPeak (1.3.6.1.4.1.5655.4.1.9.1.1.9), page B-65
10. tpNumTcpActiveFlowsPeakTime (1.3.6.1.4.1.5655.4.1.9.1.1.10), page B-65
11. tpNumUdpActiveFlows (1.3.6.1.4.1.5655.4.1.9.1.1.11), page B-66
12. tpNumUdpActiveFlowsPeak (1.3.6.1.4.1.5655.4.1.9.1.1.12), page B-66
13. tpNumUdpActiveFlowsPeakTime (1.3.6.1.4.1.5655.4.1.9.1.1.13), page B-66

14. [tpNumNonTcpUdpActiveFlows \(1.3.6.1.4.1.5655.4.1.9.1.1.14\)](#), page B-66
15. [tpNumNonTcpUdpActiveFlowsPeak \(1.3.6.1.4.1.5655.4.1.9.1.1.15\)](#), page B-67
16. [tpNumNonTcpUdpActiveFlowsPeakTime \(1.3.6.1.4.1.5655.4.1.9.1.1.16\)](#), page B-67
17. [tpTotalNumBlockedPackets \(1.3.6.1.4.1.5655.4.1.9.1.1.17\)](#), page B-67
18. [tpTotalNumBlockedFlows \(1.3.6.1.4.1.5655.4.1.9.1.1.18\)](#), page B-67
19. [tpTotalNumDiscardedPacketsDueToBwLimit \(1.3.6.1.4.1.5655.4.1.9.1.1.19\)](#), page B-68
20. [tpTotalNumWredDiscardedPackets \(1.3.6.1.4.1.5655.4.1.9.1.1.20\)](#), page B-68
21. [tpTotalNumFragments \(1.3.6.1.4.1.5655.4.1.9.1.1.21\)](#), page B-68
22. [tpTotalNumNonIpPackets \(1.3.6.1.4.1.5655.4.1.9.1.1.22\)](#), page B-68
23. [tpTotalNumIpCrcErrPackets \(1.3.6.1.4.1.5655.4.1.9.1.1.23\)](#), page B-69
24. [tpTotalNumIpLengthErrPackets \(1.3.6.1.4.1.5655.4.1.9.1.1.24\)](#), page B-69
25. [tpTotalNumIpBroadcastPackets \(1.3.6.1.4.1.5655.4.1.9.1.1.25\)](#), page B-69
26. [tpTotalNumTtlErrPackets \(1.3.6.1.4.1.5655.4.1.9.1.1.26\)](#), page B-69
27. [tpTotalNumTcpUdpCrcErrPackets \(1.3.6.1.4.1.5655.4.1.9.1.1.27\)](#), page B-70
28. [tpClearCountersTime \(1.3.6.1.4.1.5655.4.1.9.1.1.28\)](#), page B-70
29. [tpHandledPacketsRate \(1.3.6.1.4.1.5655.4.1.9.1.1.29\)](#), page B-70
30. [tpHandledPacketsRatePeak \(1.3.6.1.4.1.5655.4.1.9.1.1.30\)](#), page B-70
31. [tpHandledPacketsRatePeakTime \(1.3.6.1.4.1.5655.4.1.9.1.1.31\)](#), page B-71
32. [tpHandledFlowsRate \(1.3.6.1.4.1.5655.4.1.9.1.1.32\)](#), page B-71
33. [tpHandledFlowsRatePeak \(1.3.6.1.4.1.5655.4.1.9.1.1.33\)](#), page B-71
34. [tpHandledFlowsRatePeakTime \(1.3.6.1.4.1.5655.4.1.9.1.1.34\)](#), page B-71
35. [tpCpuUtilization \(1.3.6.1.4.1.5655.4.1.9.1.1.35\)](#), page B-72
36. [tpCpuUtilizationPeak \(1.3.6.1.4.1.5655.4.1.9.1.1.36\)](#), page B-72
37. [tpCpuUtilizationPeakTime \(1.3.6.1.4.1.5655.4.1.9.1.1.37\)](#), page B-72
38. [tpFlowsCapacityUtilization \(1.3.6.1.4.1.5655.4.1.9.1.1.38\)](#), page B-72
39. [tpFlowsCapacityUtilizationPeak \(1.3.6.1.4.1.5655.4.1.9.1.1.39\)](#), page B-73
40. [tpFlowsCapacityUtilizationPeakTime \(1.3.6.1.4.1.5655.4.1.9.1.1.40\)](#), page B-73
41. [tpServiceLoss \(1.3.6.1.4.1.5655.4.1.9.1.1.41\)](#), page B-73

Port Group: pcubePortGroup (1.3.6.1.4.1.5655.2.3.1.1.10)

The Port group provides data regarding the port, such as its type and speed.

Objects:

1. [pportModuleIndex \(1.3.6.1.4.1.5655.4.1.10.1.1.1\)](#), page B-74
2. [pportIndex \(1.3.6.1.4.1.5655.4.1.10.1.1.2\)](#), page B-74
3. [pportType \(1.3.6.1.4.1.5655.4.1.10.1.1.3\)](#), page B-75
4. [pportNumTxQueues \(1.3.6.1.4.1.5655.4.1.10.1.1.4\)](#), page B-75
5. [pportIfIndex \(1.3.6.1.4.1.5655.4.1.10.1.1.5\)](#), page B-75
6. [pportAdminSpeed \(1.3.6.1.4.1.5655.4.1.10.1.1.6\)](#), page B-75

7. [pportAdminDuplex \(1.3.6.1.4.1.5655.4.1.10.1.1.7\)](#), page B-76
8. [pportOperDuplex \(1.3.6.1.4.1.5655.4.1.10.1.1.8\)](#), page B-76
9. [pportLinkIndex \(1.3.6.1.4.1.5655.4.1.10.1.1.9\)](#), page B-76
10. [pportOperStatus \(1.3.6.1.4.1.5655.4.1.10.1.1.10\)](#), page B-76

Transmit Queues Group: pcubeTxQueuesGroup 1.3.6.1.4.1.5655.2.3.1.1.11)

The Transmit Queues group provides data regarding the transmit queue counters.

Objects:

1. [txQueuesModuleIndex \(1.3.6.1.4.1.5655.4.1.11.1.1.1\)](#), page B-77
2. [txQueuesPortIndex \(1.3.6.1.4.1.5655.4.1.11.1.1.2\)](#), page B-77
3. [txQueuesQueueIndex \(1.3.6.1.4.1.5655.4.1.11.1.1.3\)](#), page B-78
4. [txQueuesDescription \(1.3.6.1.4.1.5655.4.1.11.1.1.4\)](#), page B-78
5. [txQueuesBandwidth \(1.3.6.1.4.1.5655.4.1.11.1.1.5\)](#), page B-78
6. [txQueuesUtilization \(1.3.6.1.4.1.5655.4.1.11.1.1.6\)](#), page B-78
7. [txQueuesUtilizationPeak \(1.3.6.1.4.1.5655.4.1.11.1.1.7\)](#), page B-79
8. [txQueuesUtilizationPeakTime \(1.3.6.1.4.1.5655.4.1.11.1.1.8\)](#), page B-79
9. [txQueuesClearCountersTime \(1.3.6.1.4.1.5655.4.1.11.1.1.9\)](#), page B-79
10. [txQueuesDroppedBytes \(1.3.6.1.4.1.5655.4.1.11.1.1.10\)](#), page B-79

Global Controllers Group: pcubeGlobalControllersGroup 1.3.6.1.4.1.5655.2.3.1.1.12)

The Global Controllers group provides data regarding the Global Controllers configuration and counters.

Objects:

1. [globalControllersModuleIndex \(1.3.6.1.4.1.5655.4.1.12.1.1.1\)](#), page B-80
2. [globalControllersPortIndex \(1.3.6.1.4.1.5655.4.1.12.1.1.2\)](#), page B-81
3. [globalControllersIndex \(1.3.6.1.4.1.5655.4.1.12.1.1.3\)](#), page B-81
4. [globalControllersDescription \(1.3.6.1.4.1.5655.4.1.12.1.1.4\)](#), page B-81
5. [globalControllersBandwidth \(1.3.6.1.4.1.5655.4.1.12.1.1.5\)](#), page B-81
6. [globalControllersUtilization \(1.3.6.1.4.1.5655.4.1.12.1.1.6\)](#), page B-82
7. [globalControllersUtilizationPeak \(1.3.6.1.4.1.5655.4.1.12.1.1.7\)](#), page B-82
8. [globalControllersUtilizationPeakTime \(1.3.6.1.4.1.5655.4.1.12.1.1.8\)](#), page B-82
9. [globalControllersClearCountersTime \(1.3.6.1.4.1.5655.4.1.12.1.1.9\)](#), page B-82
10. [globalControllersDroppedBytes \(1.3.6.1.4.1.5655.4.1.12.1.1.10\)](#), page B-83

Application Group: pcubeApplicationGroup 1.3.6.1.4.1.5655.2.3.1.1.13)

The Application group indicates which application is installed in the SCE platform, and what the properties of the application and values of those properties are.

Objects:

1. [appName \(1.3.6.1.4.1.5655.4.1.13.1.1.1\)](#), page B-83
2. [appDescription \(1.3.6.1.4.1.5655.4.1.13.1.1.2\)](#), page B-84
3. [appVersion \(1.3.6.1.4.1.5655.4.1.13.1.1.3\)](#), page B-84

4. [apIndex \(1.3.6.1.4.1.5655.4.1.13.2.1.1\)](#), page B-85
5. [apName \(1.3.6.1.4.1.5655.4.1.13.2.1.2\)](#), page B-85
6. [apType \(1.3.6.1.4.1.5655.4.1.13.2.1.3\)](#), page B-85
7. [apvPropertyName \(1.3.6.1.4.1.5655.4.1.13.3.1.2\)](#), page B-87
8. [apvRowStatus \(1.3.6.1.4.1.5655.4.1.13.3.1.3\)](#), page B-87
9. [apvPropertyStringValue \(1.3.6.1.4.1.5655.4.1.13.3.1.4\)](#), page B-87
10. [apvPropertyUintValue \(1.3.6.1.4.1.5655.4.1.13.3.1.5\)](#), page B-87
11. [apvPropertyCounter64Value \(1.3.6.1.4.1.5655.4.1.13.3.1.6\)](#), page B-88

Traffic Counters Group: pcubeTrafficCountersGroup 1.3.6.1.4.1.5655.2.3.1.1.14)

The Traffic Counters group provides information regarding the value of different the traffic counters.

Objects:

1. [trafficCounterIndex \(1.3.6.1.4.1.5655.4.1.14.1.1.1\)](#), page B-88
2. [trafficCounterValue \(1.3.6.1.4.1.5655.4.1.14.1.1.2\)](#), page B-89
3. [trafficCounterName \(1.3.6.1.4.1.5655.4.1.14.1.1.3\)](#), page B-89
4. [trafficCounterType \(1.3.6.1.4.1.5655.4.1.14.1.1.4\)](#), page B-89

Attack Group: pcubeAttackGroup (1.3.6.1.4.1.5655.2.3.1.1.15)

The Attack group provides information regarding detected attacks, aggregated by attack type.

Objects:

1. [attackTypeIndex \(1.3.6.1.4.1.5655.4.1.15.1.1.1\)](#), page B-90
2. [attackTypeName \(1.3.6.1.4.1.5655.4.1.15.1.1.2\)](#), page B-90
3. [attackTypeCurrentNumAttacks \(1.3.6.1.4.1.5655.4.1.15.1.1.3\)](#), page B-90
4. [attackTypeTotalNumAttacks \(1.3.6.1.4.1.5655.4.1.15.1.1.4\)](#), page B-91
5. [attackTypeTotalNumFlows \(1.3.6.1.4.1.5655.4.1.15.1.1.5\)](#), page B-91
6. [attackTypeTotalNumSeconds \(1.3.6.1.4.1.5655.4.1.15.1.1.6\)](#), page B-91

VAS Group: pcubeVasTrafficForwardingGroup 1.3.6.1.4.1.5655.2.3.1.1.16)

The VAS group provides information regarding the status of the VAS servers.

Objects:

1. [vasServerIndex \(1.3.6.1.4.1.5655.4.1.16.1.1.1\)](#), page B-92
2. [vasServerId \(1.3.6.1.4.1.5655.4.1.16.1.1.2\)](#), page B-92
3. [vasServerAdminStatus \(1.3.6.1.4.1.5655.4.1.16.1.1.3\)](#), page B-92
4. [vasServerOperStatus \(1.3.6.1.4.1.5655.4.1.16.1.1.4\)](#), page B-93

MPLS/VPN Group: pcubeMplsVpnAutoLearnGroup 1.3.6.1.4.1.5655.2.3.1.1.17)

The MPLS/VPN Group provides data regarding MPLS/VPN auto-learning.

Objects:

1. [mplsVpnMaxHWMappings \(1.3.6.1.4.1.5655.4.1.17.1.1.1\)](#), page B-93
2. [mplsVpnCurrentHWMappings \(1.3.6.1.4.1.5655.4.1.17.1.1.2\)](#), page B-94

Traps Group: pcubeTrapObjectsGroup 1.3.6.1.4.1.5655.2.3.1.1.18)

The traps group includes strings used for notifications.

Objects:

1. **pcubeSeEventGenericString1**
2. **pcubeSeEventGenericString2**
3. **pullRequestNumber**

pcubeCompliances (1.3.6.1.4.1.5655.2.3.1.2)

Module compliance is a compliance statement defined in this MIB module that defines which groups must be implemented.

- [pcubeCompliance module-compliances 1.3.6.1.4.1.5655.2.3.1.2.1](#), page B-15

pcubeCompliance module-compliances 1.3.6.1.4.1.5655.2.3.1.2.1)

A compliance statement defined in this MIB module, for SCE platform SNMP agents.

Module Name:

pcubeSeMIB

Mandatory Groups

1. **pcubeSystemGroup**
2. **pcubeChassisGroup**
3. **pcubeModuleGroup**
4. **pcubeLinkGroup**
5. **pcubeDiskGroup**
6. **pcubeRdrFormatterGroup**
7. **pcubeLoggerGroup**
8. **pcubeSubscribersGroup**
9. **pcubeTrafficProcessorGroup**
10. **pcubePortGroup**
11. **pcubeTxQueuesGroup**
12. **pcubeGlobalControllersGroup**
13. **pcubeApplicationGroup**
14. **pcubeTrafficCountersGroup**
15. **pcubeAttackGroup**
16. **pcubeVasTrafficForwardingGroup**
17. **pcubeMplsVpnAutoLearnGroup**
18. **pcubeTrapObjectsGroup**

pcubeWorkgroup (1.3.6.1.4.1.5655.4)

pcubeWorkgroup is the main subtree for objects and events of the Cisco SCE platform products.

- [Notification Types](#), page B-16
- [pcubeSe Objects](#), page B-24

Notification Types

Following is a list of the SCE platform notification types

| | |
|--|---------------------------|
| operationalStatusOperationalTrap | {1.3.6.1.4.1.5655.4.0.1} |
| operationalStatusWarningTrap | {1.3.6.1.4.1.5655.4.0.2} |
| operationalStatusFailureTrap | {1.3.6.1.4.1.5655.4.0.3} |
| systemResetTrap | {1.3.6.1.4.1.5655.4.0.4} |
| chassisTempAlarmOnTrap | {1.3.6.1.4.1.5655.4.0.5} |
| chassisTempAlarmOffTrap | {1.3.6.1.4.1.5655.4.0.6} |
| chassisVoltageAlarmOnTrap | {1.3.6.1.4.1.5655.4.0.7} |
| chassisFansAlarmOnTrap | {1.3.6.1.4.1.5655.4.0.8} |
| chassisPowerSupplyAlarmOnTrap | {1.3.6.1.4.1.5655.4.0.9} |
| rdrActiveConnectionTrap | {1.3.6.1.4.1.5655.4.0.10} |
| rdrNoActiveConnectionTrap | {1.3.6.1.4.1.5655.4.0.11} |
| rdrConnectionUpTrap | {1.3.6.1.4.1.5655.4.0.12} |
| rdrConnectionDownTrap | {1.3.6.1.4.1.5655.4.0.13} |
| loggerUserLogIsFullTrap | {1.3.6.1.4.1.5655.4.0.18} |
| sntpClockDriftWarnTrap | {1.3.6.1.4.1.5655.4.0.19} |
| linkModeBypassTrap | {1.3.6.1.4.1.5655.4.0.20} |
| linkModeForwardingTrap | {1.3.6.1.4.1.5655.4.0.21} |
| linkModeCutoffTrap | {1.3.6.1.4.1.5655.4.0.22} |
| moduleAttackFilterActivatedTrap | {1.3.6.1.4.1.5655.4.0.25} |
| moduleAttackFilterDeactivatedTrap | {1.3.6.1.4.1.5655.4.0.26} |
| moduleEmAgentGenericTrap | {1.3.6.1.4.1.5655.4.0.27} |
| linkModeSniffingTrap | {1.3.6.1.4.1.5655.4.0.28} |
| moduleRedundancyReadyTrap | {1.3.6.1.4.1.5655.4.0.29} |
| moduleRedundantConfigurationMismatchTrap | {1.3.6.1.4.1.5655.4.0.30} |
| moduleLostRedundancyTrap | {1.3.6.1.4.1.5655.4.0.31} |
| moduleSmConnectionDownTrap | {1.3.6.1.4.1.5655.4.0.32} |
| moduleSmConnectionUpTrap | {1.3.6.1.4.1.5655.4.0.33} |
| moduleOperStatusChangeTrap | {1.3.6.1.4.1.5655.4.0.34} |
| portOperStatusChangeTrap | {1.3.6.1.4.1.5655.4.0.35} |

| | |
|--|---------------------------|
| chassisLineFeedAlarmOnTrap | {1.3.6.1.4.1.5655.4.0.36} |
| rdrFormatterCategoryDiscardingReportsTrap | {1.3.6.1.4.1.5655.4.0.37} |
| rdrFormatterCategoryStoppedDiscardingReportsTrap | {1.3.6.1.4.1.5655.4.0.38} |
| sessionStartedTrap | {1.3.6.1.4.1.5655.4.0.39} |
| sessionEndedTrap | {1.3.6.1.4.1.5655.4.0.40} |
| sessionDeniedAccessTrap | {1.3.6.1.4.1.5655.4.0.41} |
| sessionBadLoginTrap | {1.3.6.1.4.1.5655.4.0.42} |
| illegalSubscriberMappingTrap | {1.3.6.1.4.1.5655.4.0.43} |
| loggerLineAttackLogFullTrap | {1.3.6.1.4.1.5655.4.0.44} |
| vasServerOperationStatusChangeTrap | {1.3.6.1.4.1.5655.4.0.45} |
| pullRequestNumber | {1.3.6.1.4.1.5655.4.0.46} |
| pullRequestRetryFailedTrap | {1.3.6.1.4.1.5655.4.0.47} |
| mplsVpnTotalHWMappingsThresholdExceededTrap | {1.3.6.1.4.1.5655.4.0.48} |

- [operationalStatusOperationalTrap \(1.3.6.1.4.1.5655.4.0.1\)](#), page B-18
- [operationalStatusWarningTrap \(1.3.6.1.4.1.5655.4.0.2\)](#), page B-18
- [operationalStatusFailureTrap \(1.3.6.1.4.1.5655.4.0.3\)](#), page B-18
- [systemResetTrap \(1.3.6.1.4.1.5655.4.0.4\)](#), page B-18
- [chassisTempAlarmOnTrap \(1.3.6.1.4.1.5655.4.0.5\)](#), page B-18
- [chassisTempAlarmOffTrap \(1.3.6.1.4.1.5655.4.0.6\)](#), page B-19
- [chassisVoltageAlarmOnTrap \(1.3.6.1.4.1.5655.4.0.7\)](#), page B-19
- [chassisFansAlarmOnTrap \(1.3.6.1.4.1.5655.4.0.8\)](#), page B-19
- [chassisPowerSupplyAlarmOnTrap \(1.3.6.1.4.1.5655.4.0.9\)](#), page B-19
- [rdrActiveConnectionTrap \(1.3.6.1.4.1.5655.4.0.10\)](#), page B-19
- [rdrNoActiveConnectionTrap \(1.3.6.1.4.1.5655.4.0.11\)](#), page B-19
- [rdrConnectionUpTrap \(1.3.6.1.4.1.5655.4.0.12\)](#), page B-19
- [rdrConnectionDownTrap \(1.3.6.1.4.1.5655.4.0.13\)](#), page B-19
- [loggerUserLogIsFullTrap \(1.3.6.1.4.1.5655.4.0.18\)](#), page B-19
- [sntpClockDriftWarnTrap \(1.3.6.1.4.1.5655.4.0.19\)](#), page B-19
- [linkModeBypassTrap \(1.3.6.1.4.1.5655.4.0.20\)](#), page B-20
- [linkModeForwardingTrap \(1.3.6.1.4.1.5655.4.0.21\)](#), page B-20
- [linkModeCutoffTrap \(1.3.6.1.4.1.5655.4.0.22\)](#), page B-20
- [moduleAttackFilterActivatedTrap \(1.3.6.1.4.1.5655.4.0.25\)](#), page B-20
- [moduleAttackFilterDeactivatedTrap \(1.3.6.1.4.1.5655.4.0.26\)](#), page B-21
- [moduleEmAgentGenericTrap \(1.3.6.1.4.1.5655.4.0.27\)](#), page B-21
- [linkModeSniffingTrap \(1.3.6.1.4.1.5655.4.0.28\)](#), page B-21

- [moduleRedundancyReadyTrap \(1.3.6.1.4.1.5655.4.0.29\)](#), page B-21
- [moduleRedundantConfigurationMismatchTrap \(1.3.6.1.4.1.5655.4.0.30\)](#), page B-21
- [moduleLostRedundancyTrap \(1.3.6.1.4.1.5655.4.0.31\)](#), page B-21
- [moduleSmConnectionDownTrap \(1.3.6.1.4.1.5655.4.0.32\)](#), page B-22
- [moduleSmConnectionUpTrap \(1.3.6.1.4.1.5655.4.0.33\)](#), page B-22
- [moduleOperStatusChangeTrap \(1.3.6.1.4.1.5655.4.0.34\)](#), page B-22
- [portOperStatusChangeTrap \(1.3.6.1.4.1.5655.4.0.35\)](#), page B-22
- [chassisLineFeedAlarmOnTrap \(1.3.6.1.4.1.5655.4.0.36\)](#), page B-22
- [rdrFormatterCategoryDiscardingReportsTrap \(1.3.6.1.4.1.5655.4.0.37\)](#), page B-22
- [rdrFormatterCategoryStoppedDiscardingReportsTrap \(1.3.6.1.4.1.5655.4.0.38\)](#), page B-22
- [sessionStartedTrap \(1.3.6.1.4.1.5655.4.0.39\)](#), page B-22
- [sessionEndedTrap \(1.3.6.1.4.1.5655.4.0.40\)](#), page B-22
- [sessionDeniedAccessTrap \(1.3.6.1.4.1.5655.4.0.41\)](#), page B-22
- [sessionBadLoginTrap \(1.3.6.1.4.1.5655.4.0.42\)](#), page B-23
- [illegalSubscriberMappingTrap \(1.3.6.1.4.1.5655.4.0.43\)](#), page B-23
- [loggerLineAttackLogFullTrap \(1.3.6.1.4.1.5655.4.0.44\)](#), page B-23
- [vasServerOperationalStatusChangeTrap \(1.3.6.1.4.1.5655.4.0.45\)](#), page B-23
- [pullRequestNumber \(1.3.6.1.4.1.5655.4.0.46\)](#), page B-23
- [pullRequestRetryFailedTrap \(1.3.6.1.4.1.5655.4.0.47\)](#), page B-23
- [mplsVpnTotalHWMappingsThresholdExceededTrap \(1.3.6.1.4.1.5655.4.0.48\)](#), page B-23

operationalStatusOperationalTrap (1.3.6.1.4.1.5655.4.0.1)

The system operational state of the SCE platform has changed to *Operational*(3).

operationalStatusWarningTrap (1.3.6.1.4.1.5655.4.0.2)

The system operational state of the SCE platform has changed to *Warning*(4).

operationalStatusFailureTrap (1.3.6.1.4.1.5655.4.0.3)

The system operational state of the SCE platform has changed to *Failure*(5).

systemResetTrap (1.3.6.1.4.1.5655.4.0.4)

The agent entity is about to reset itself either per user request or due to a fatal event.

chassisTempAlarmOnTrap (1.3.6.1.4.1.5655.4.0.5)

The **chassisTempAlarm** object in this MIB has transitioned to the *On*(3) state, indicating that the temperature is too high.

chassisTempAlarmOffTrap (1.3.6.1.4.1.5655.4.0.6)

The **chassisTempAlarm** object in this MIB has transitioned to the *Off(2)* state, indicating that the temperature level is back to normal.

chassisVoltageAlarmOnTrap (1.3.6.1.4.1.5655.4.0.7)

The **chassisVoltageAlarm** object in this MIB has transitioned to the *On(3)* state, indicating that the voltage level is out of safe bounds.

chassisFansAlarmOnTrap (1.3.6.1.4.1.5655.4.0.8)

The **chassisFansAlarm** object in this MIB has transitioned to the *On(3)* state, indicating fan malfunction.

chassisPowerSupplyAlarmOnTrap (1.3.6.1.4.1.5655.4.0.9)

The **chassisPowerSupplyAlarm** object in this MIB has transitioned to the *On(3)* state, indicating power supply malfunction.

rdrActiveConnectionTrap (1.3.6.1.4.1.5655.4.0.10)

One of the RDR-formatter connections has become the active connection.

rdrNoActiveConnectionTrap (1.3.6.1.4.1.5655.4.0.11)

There is no active connection between the RDR-formatter and any Collection Manager.

rdrConnectionUpTrap (1.3.6.1.4.1.5655.4.0.12)

The **rdrFormatterDestConnectionStatus** object in this MIB has transitioned to *Up(2)*, indicating that one of the RDR-formatter connections was established.

rdrConnectionDownTrap (1.3.6.1.4.1.5655.4.0.13)

The **rdrFormatterDestConnectionStatus** object in this MIB has transitioned to *Down(3)*, indicating that one of the RDR-formatter connections was disconnected.

loggerUserLogIsFullTrap (1.3.6.1.4.1.5655.4.0.18)

The User log file is full. The agent entity then rolls to the next file.

sntpClockDriftWarnTrap (1.3.6.1.4.1.5655.4.0.19)

The SNTP agent has not received an SNTP time update for a long period, which may result in a time drift of the agent entity's clock.

linkModeBypassTrap (1.3.6.1.4.1.5655.4.0.20)

The link mode has changed to bypass.

linkModeForwardingTrap (1.3.6.1.4.1.5655.4.0.21)

The link mode has changed to forwarding.

linkModeCutoffTrap (1.3.6.1.4.1.5655.4.0.22)

The link mode has changed to cutoff.

moduleAttackFilterActivatedTrap (1.3.6.1.4.1.5655.4.0.25)

The attack filter module has detected an attack and activated a filter. The type of attack-filter that was activated is returned in pcubeSeEventGenericString1.

Following are several examples of pcubeSeEventGenericString1 for various scenarios:

- **Attack detected automatically** (the number of open flows or ddos-suspected flows has exceeded the maximum configured for the attack detector):
 - **Source of the attack is detected** (at the subscriber side, IP address = 10.1.4.134, attacking the network side using UDP, number of open flows = 10000, configured action is 'report'):
- Attack detected: Attack from IP address 10.1.4.134, from subscriber side, protocol UDP. 10000 concurrent open flows detected, 57 concurrent Ddos-suspected flows detected. Action is: Report.
 - **Target of the attack is detected** (at the network side, IP address = 10.1.4.135, being attacked from the subscriber side using ICMP, number of ddos-suspected flows = 500, configured action is 'block'):
- Attack detected: Attack on IP address 10.1.4.135, from subscriber side, protocol ICMP. 745 concurrent open flows detected, 500 concurrent Ddos-suspected flows detected. Action is: Block.
- **Forced filtering** using the 'force-filter' command:
 - Action is 'block', attack-direction is attack-source, side is subscriber, IP address = 10.1.1.1, and protocol is TCP:
- Attack filter: Forced block of flows from IP address 10.1.1.1, from subscriber side, protocol TCP. Attack forced using a force-filter command.
 - When the action is 'report', attack-direction is attack-destination, side is subscriber, IP address = 10.1.1.1, and protocol is Other:
- Attack filter: Forced report to IP address 10.1.1.1, from network side, protocol Other. Attack forced using a force-filter command.

moduleAttackFilterDeactivatedTrap (1.3.6.1.4.1.5655.4.0.26)

The attack filter module has removed a filter that was previously activated.

- Attack filter type — in pcubeSeEventGenericString1 (refer to corresponding moduleAttackFilterActivatedTrap)
- Reason for deactivating the filter — in pcubeSeEventGenericString2

Following are several examples of pcubeSeEventGenericString1 for various scenarios:

- **Attack end detected automatically** (the number of open flows or ddos-suspected flows drops below the minimum value configured for the attack detector):
End-of-attack detected — Attack on IP address 10.1.4.135, from subscriber side, protocol UDP. Action is: Report. Duration 20 seconds, attack comprised of 11736 flows.
End-of-attack detected — Attack from IP address 10.1.4.134, from subscriber side, protocol ICMP. Action is: Block. Duration 10 seconds, attack comprised of 2093 flows.
- **Attack end forced** by a 'dont-filter', or a previous 'force-filter' command is removed:
Attack filter — Forced to end block of flows from IP address 10.1.1.1, from subscriber side, protocol TCP. Attack end forced using a 'no force-filter' or a 'dont-filter' command. Duration 6 seconds, 1 flows blocked.
Attack filter — Forced to end report to IP address 10.1.1.1, from network side, protocol Other. Attack end forced using a 'no force-filter' or a 'dont-filter' command. Duration 13 seconds, attack comprised of 1 flows.

moduleEmAgentGenericTrap (1.3.6.1.4.1.5655.4.0.27)

A generic trap used by the Cisco management agent.

- Trap name — in pcubeSeEventGenericString1 (refer to corresponding moduleAttackFilterActivatedTrap)
- Relevant parameter — in pcubeSeEventGenericString2

linkModeSniffingTrap (1.3.6.1.4.1.5655.4.0.28)

The agent entity has detected that the **linkOperMode** object in this MIB has changed to *sniffing(5)*.

moduleRedundancyReadyTrap (1.3.6.1.4.1.5655.4.0.29)

The module was able to connect and synch with a redundant entity, and is now ready to handle fail-over if needed.

moduleRedundantConfigurationMismatchTrap (1.3.6.1.4.1.5655.4.0.30)

The module was not able to synch with a redundant entity, due to an incompatibility in essential configuration parameters between the module and the redundant entity.

moduleLostRedundancyTrap (1.3.6.1.4.1.5655.4.0.31)

The module has lost the ability to perform the fail-over procedure.

moduleSmConnectionDownTrap (1.3.6.1.4.1.5655.4.0.32)

The virtual connection to the SM (Subscriber Manager) is broken.

moduleSmConnectionUpTrap (1.3.6.1.4.1.5655.4.0.33)

The virtual connection to the SM is up and working.

moduleOperStatusChangeTrap (1.3.6.1.4.1.5655.4.0.34)

The value of **moduleOperStatus** has changed.

portOperStatusChangeTrap (1.3.6.1.4.1.5655.4.0.35)

The value of the **portOperStatus** object of the **portIndex** has changed, indicating that the link was either forced down or the force down was released.

chassisLineFeedAlarmOnTrap (1.3.6.1.4.1.5655.4.0.36)

The agent entity has detected that the **chassisLineFeed** object in this MIB has changed to the *on(3)* state.

rdrFormatterCategoryDiscardingReportsTrap (1.3.6.1.4.1.5655.4.0.37)

The agent entity has detected that reports sent to this category are being discarded.

The **rdrFormatterCategoryNumReportsDiscarded** object in this MIB counts the number of discarded reports.

rdrFormatterCategoryStoppedDiscardingReportsTrap (1.3.6.1.4.1.5655.4.0.38)

The agent entity has detected that reports sent to this category are no longer being discarded.

The **rdrFormatterCategoryNumReportsDiscarded** object in this MIB counts the number of discarded reports.

sessionStartedTrap (1.3.6.1.4.1.5655.4.0.39)

The agent entity has accepted a new session. The **pcubeSeEventGenericString1** contains the session type (telnet/SSH) and client IP address.

sessionEndedTrap (1.3.6.1.4.1.5655.4.0.40)

The agent entity has detected the end of a session. The **pcubeSeEventGenericString1** contains the session type (telnet/SSH) and client IP address.

sessionDeniedAccessTrap (1.3.6.1.4.1.5655.4.0.41)

The agent entity has refused a session from unauthorized source. The **pcubeSeEventGenericString1** contains the session type (telnet/SSH) and client IP address.

sessionBadLoginTrap (1.3.6.1.4.1.5655.4.0.42)

The agent entity has detected attempt to login with a wrong password. The `pcubeSeEventGenericString1` contains the session type (telnet/SSH) and client IP address.

illegalSubscriberMappingTrap (1.3.6.1.4.1.5655.4.0.43)

The agent entity has detected that an external entity has attempted to create an illegal or inconsistent subscriber mapping.

`pcubeSeEventGenericString1` contains a message describing the problem.

loggerLineAttackLogFullTrap (1.3.6.1.4.1.5655.4.0.44)

The agent entity has detected that the attack log is full and a new log file is opened.

vasServerOperationalStatusChangeTrap (1.3.6.1.4.1.5655.4.0.45)

The agent entity has detected a change in the operational status of a VAS server.

pullRequestNumber (1.3.6.1.4.1.5655.4.0.46)

The number of pull requests currently issued for the anonymous subscriber identified in the `pullRequestRetryFailedTrap`.

Always returns a value of 0.

pullRequestRetryFailedTrap (1.3.6.1.4.1.5655.4.0.47)

An unknown subscriber could not be identified after a certain number of pull requests, and is suspected to be an intruder.

`pcubeSeEventGenericString1` contains subscriber ID.

mplsVpnTotalHWMappingsThresholdExceededTrap (1.3.6.1.4.1.5655.4.0.48)

The value of `mplsVpnCurrentHWMappings` exceeds the allowed threshold.

pcubeSe Objects

The pcubeSe objects provide configuration and runtime status for the SCE platform.

- [sysOperationalStatus \(1.3.6.1.4.1.5655.4.1.1.1\)](#), page B-30
- [sysFailureRecovery \(1.3.6.1.4.1.5655.4.1.1.2\)](#), page B-30
- [sysVersion \(1.3.6.1.4.1.5655.4.1.1.3\)](#), page B-31
- [pchassisSysType \(1.3.6.1.4.1.5655.4.1.2.1\)](#), page B-31
- [pchassisPowerSupplyAlarm \(1.3.6.1.4.1.5655.4.1.2.2\)](#), page B-31
- [pchassisFansAlarm \(1.3.6.1.4.1.5655.4.1.2.3\)](#), page B-31
- [pchassisTempAlarm \(1.3.6.1.4.1.5655.4.1.2.4\)](#), page B-32
- [pchassisVoltageAlarm \(1.3.6.1.4.1.5655.4.1.2.5\)](#), page B-32
- [pchassisNumSlots \(1.3.6.1.4.1.5655.4.1.2.6\)](#), page B-32
- [pchassisSlotConfig \(1.3.6.1.4.1.5655.4.1.2.7\)](#), page B-33
- [pchassisPsuType \(1.3.6.1.4.1.5655.4.1.2.8\)](#), page B-33
- [pchassisLineFeedAlarm \(1.3.6.1.4.1.5655.4.1.2.9\)](#), page B-33
- [pmoduleTable \(1.3.6.1.4.1.5655.4.1.3.1\)](#), page B-34
- [pmoduleEntry \(1.3.6.1.4.1.5655.4.1.3.1.1\)](#), page B-34
- [pmoduleIndex \(1.3.6.1.4.1.5655.4.1.3.1.1.1\)](#), page B-34
- [pmoduleType \(1.3.6.1.4.1.5655.4.1.3.1.1.2\)](#), page B-35
- [pmoduleNumTrafficProcessors \(1.3.6.1.4.1.5655.4.1.3.1.1.3\)](#), page B-35
- [pmoduleSlotNum \(1.3.6.1.4.1.5655.4.1.3.1.1.4\)](#), page B-35
- [pmoduleHwVersion \(1.3.6.1.4.1.5655.4.1.3.1.1.5\)](#), page B-36
- [pmoduleNumPorts \(1.3.6.1.4.1.5655.4.1.3.1.1.6\)](#), page B-36
- [pmoduleNumLinks \(1.3.6.1.4.1.5655.4.1.3.1.1.7\)](#), page B-36
- [pmoduleConnectionMode \(1.3.6.1.4.1.5655.4.1.3.1.1.8\)](#), page B-36
- [pmoduleSerialNumber \(1.3.6.1.4.1.5655.4.1.3.1.1.9\)](#), page B-37
- [pmoduleUpStreamAttackFilteringTime \(1.3.6.1.4.1.5655.4.1.3.1.1.10\)](#), page B-37
- [pmoduleUpStreamLastAttackFilteringTime \(1.3.6.1.4.1.5655.4.1.3.1.1.11\)](#), page B-37
- [pmoduleDownStreamAttackFilteringTime \(1.3.6.1.4.1.5655.4.1.3.1.1.12\)](#), page B-37
- [pmoduleDownStreamLastAttackFilteringTime \(1.3.6.1.4.1.5655.4.1.3.1.1.13\)](#), page B-38
- [pmoduleAttackObjectsClearTime \(1.3.6.1.4.1.5655.4.1.3.1.1.14\)](#), page B-38
- [pmoduleAdminStatus \(1.3.6.1.4.1.5655.4.1.3.1.1.15\)](#), page B-38
- [pmoduleOperStatus \(1.3.6.1.4.1.5655.4.1.3.1.1.16\)](#), page B-38
- [linkTable \(1.3.6.1.4.1.5655.4.1.4.1\)](#), page B-39
- [linkEntry \(1.3.6.1.4.1.5655.4.1.4.1.1\)](#), page B-39
- [linkModuleIndex \(1.3.6.1.4.1.5655.4.1.4.1.1.1\)](#), page B-39
- [linkIndex \(1.3.6.1.4.1.5655.4.1.4.1.1.2\)](#), page B-40
- [linkAdminModeOnActive \(1.3.6.1.4.1.5655.4.1.4.1.1.3\)](#), page B-40

- [linkAdminModeOnFailure \(1.3.6.1.4.1.5655.4.1.4.1.1.4\)](#), page B-40
- [linkOperMode \(1.3.6.1.4.1.5655.4.1.4.1.1.5\)](#), page B-41
- [linkStatusReflectionEnable \(1.3.6.1.4.1.5655.4.1.4.1.1.6\)](#), page B-41
- [linkSubscriberSidePortIndex \(1.3.6.1.4.1.5655.4.1.4.1.1.7\)](#), page B-41
- [linkNetworkSidePortIndex \(1.3.6.1.4.1.5655.4.1.4.1.1.8\)](#), page B-42
- [diskNumUsedBytes \(1.3.6.1.4.1.5655.4.1.5.1\)](#), page B-42
- [diskNumFreeBytes \(1.3.6.1.4.1.5655.4.1.5.2\)](#), page B-42
- [rdrFormatterEnable \(1.3.6.1.4.1.5655.4.1.6.1\)](#), page B-42
- [rdrFormatterDestTable \(1.3.6.1.4.1.5655.4.1.6.2\)](#), page B-43
- [rdrFormatterDestEntry \(1.3.6.1.4.1.5655.4.1.6.2.1\)](#), page B-43
- [rdrFormatterDestIPAddr \(1.3.6.1.4.1.5655.4.1.6.2.1.1\)](#), page B-43
- [rdrFormatterDestPort \(1.3.6.1.4.1.5655.4.1.6.2.1.2\)](#), page B-44
- [rdrFormatterDestPriority \(1.3.6.1.4.1.5655.4.1.6.2.1.3\)](#), page B-44
- [rdrFormatterDestStatus \(1.3.6.1.4.1.5655.4.1.6.2.1.4\)](#), page B-44
- [rdrFormatterDestConnectionStatus \(1.3.6.1.4.1.5655.4.1.6.2.1.5\)](#), page B-45
- [rdrFormatterDestNumReportsSent \(1.3.6.1.4.1.5655.4.1.6.2.1.6\)](#), page B-45
- [rdrFormatterDestNumReportsDiscarded \(1.3.6.1.4.1.5655.4.1.6.2.1.7\)](#), page B-45
- [rdrFormatterDestReportRate \(1.3.6.1.4.1.5655.4.1.6.2.1.8\)](#), page B-45
- [rdrFormatterDestReportRatePeak \(1.3.6.1.4.1.5655.4.1.6.2.1.9\)](#), page B-46
- [rdrFormatterDestReportRatePeakTime \(1.3.6.1.4.1.5655.4.1.6.2.1.10\)](#), page B-46
- [rdrFormatterNumReportsSent \(1.3.6.1.4.1.5655.4.1.6.3\)](#), page B-46
- [rdrFormatterNumReportsDiscarded \(1.3.6.1.4.1.5655.4.1.6.4\)](#), page B-46
- [rdrFormatterClearCountersTime \(1.3.6.1.4.1.5655.4.1.6.5\)](#), page B-47
- [rdrFormatterReportRate \(1.3.6.1.4.1.5655.4.1.6.6\)](#), page B-47
- [rdrFormatterReportRatePeak \(1.3.6.1.4.1.5655.4.1.6.7\)](#), page B-47
- [rdrFormatterReportRatePeakTime \(1.3.6.1.4.1.5655.4.1.6.8\)](#), page B-47
- [rdrFormatterProtocol \(1.3.6.1.4.1.5655.4.1.6.9\)](#), page B-48
- [rdrFormatterForwardingMode \(1.3.6.1.4.1.5655.4.1.6.10\)](#), page B-48
- [rdrFormatterCategoryTable \(1.3.6.1.4.1.5655.4.1.6.11\)](#), page B-48
- [rdrFormatterCategoryEntry \(1.3.6.1.4.1.5655.4.1.6.11.1\)](#), page B-49
- [rdrFormatterCategoryIndex \(1.3.6.1.4.1.5655.4.1.6.11.1.1\)](#), page B-49
- [rdrFormatterCategoryName \(1.3.6.1.4.1.5655.4.1.6.11.1.2\)](#), page B-49
- [rdrFormatterCategoryNumReportsSent \(1.3.6.1.4.1.5655.4.1.6.11.1.3\)](#), page B-49
- [rdrFormatterCategoryNumReportsDiscarded \(1.3.6.1.4.1.5655.4.1.6.11.1.4\)](#), page B-50
- [rdrFormatterCategoryReportRate \(1.3.6.1.4.1.5655.4.1.6.11.1.5\)](#), page B-50
- [rdrFormatterCategoryReportRatePeak \(1.3.6.1.4.1.5655.4.1.6.11.1.6\)](#), page B-50
- [rdrFormatterCategoryReportRatePeakTime \(1.3.6.1.4.1.5655.4.1.6.11.1.7\)](#), page B-50
- [rdrFormatterCategoryNumReportsQueued \(1.3.6.1.4.1.5655.4.1.6.11.1.8\)](#), page B-51

- rdrFormatterCategoryDestTable (1.3.6.1.4.1.5655.4.1.6.12), page B-51
- rdrFormatterCategoryDestEntry (1.3.6.1.4.1.5655.4.1.6.12.1), page B-51
- rdrFormatterCategoryDestPriority (1.3.6.1.4.1.5655.4.1.6.12.1.1), page B-51
- rdrFormatterCategoryDestStatus (1.3.6.1.4.1.5655.4.1.6.12.1.2), page B-52
- loggerUserLogEnable (1.3.6.1.4.1.5655.4.1.7.1), page B-52
- loggerUserLogNumInfo (1.3.6.1.4.1.5655.4.1.7.2), page B-52
- loggerUserLogNumWarning (1.3.6.1.4.1.5655.4.1.7.3), page B-53
- loggerUserLogNumError (1.3.6.1.4.1.5655.4.1.7.4), page B-53
- loggerUserLogNumFatal (1.3.6.1.4.1.5655.4.1.7.5), page B-53
- loggerUserLogClearCountersTime (1.3.6.1.4.1.5655.4.1.7.6), page B-53
- subscribersInfoTable (1.3.6.1.4.1.5655.4.1.8.1), page B-54
- subscribersInfoEntry (1.3.6.1.4.1.5655.4.1.8.1.1), page B-54
- subscribersNumIntroduced (1.3.6.1.4.1.5655.4.1.8.1.1.1), page B-54
- subscribersNumFree (1.3.6.1.4.1.5655.4.1.8.1.1.2), page B-55
- subscribersNumIpAddrMappings (1.3.6.1.4.1.5655.4.1.8.1.1.3), page B-55
- subscribersNumIpAddrMappingsFree (1.3.6.1.4.1.5655.4.1.8.1.1.4), page B-55
- subscribersNumIpRangeMappings (1.3.6.1.4.1.5655.4.1.8.1.1.5), page B-55
- subscribersNumIpRangeMappingsFree (1.3.6.1.4.1.5655.4.1.8.1.1.6), page B-56
- subscribersNumVlanMappings (1.3.6.1.4.1.5655.4.1.8.1.1.7), page B-56
- subscribersNumVlanMappingsFree (1.3.6.1.4.1.5655.4.1.8.1.1.8), page B-56
- subscribersNumActive (1.3.6.1.4.1.5655.4.1.8.1.1.9), page B-56
- subscribersNumActivePeak (1.3.6.1.4.1.5655.4.1.8.1.1.10), page B-57
- subscribersNumActivePeakTime (1.3.6.1.4.1.5655.4.1.8.1.1.11), page B-57
- subscribersNumUpdates (1.3.6.1.4.1.5655.4.1.8.1.1.12), page B-57
- subscribersCountersClearTime (1.3.6.1.4.1.5655.4.1.8.1.1.13), page B-57
- subscribersNumTpIpRangeMappings (1.3.6.1.4.1.5655.4.1.8.1.1.14), page B-58
- subscribersNumTpIpRangeMappingsFree (1.3.6.1.4.1.5655.4.1.8.1.1.15), page B-58
- subscribersNumAnonymous (1.3.6.1.4.1.5655.4.1.8.1.1.16), page B-58
- subscribersNumWithSessions (1.3.6.1.4.1.5655.4.1.8.1.1.17), page B-58
- subscribersPropertiesTable (1.3.6.1.4.1.5655.4.1.8.2), page B-59
- subscribersPropertiesEntry (1.3.6.1.4.1.5655.4.1.8.2.1), page B-59
- spIndex (1.3.6.1.4.1.5655.4.1.8.2.1.1), page B-59
- spName (1.3.6.1.4.1.5655.4.1.8.2.1.2), page B-59
- spType (1.3.6.1.4.1.5655.4.1.8.2.1.3), page B-60
- subscriberPropertiesValuesTable (1.3.6.1.4.1.5655.4.1.8.3), page B-60
- subscriberPropertiesValueEntry (1.3.6.1.4.1.5655.4.1.8.3.1), page B-60
- spvIndex (1.3.6.1.4.1.5655.4.1.8.3.1.1), page B-61
- spvSubName (1.3.6.1.4.1.5655.4.1.8.3.1.2), page B-61

- spvPropertyName (1.3.6.1.4.1.5655.4.1.8.3.1.3), page B-61
- spvRowStatus (1.3.6.1.4.1.5655.4.1.8.3.1.4), page B-61
- spvPropertyStringValue (1.3.6.1.4.1.5655.4.1.8.3.1.5), page B-62
- spvPropertyUintValue (1.3.6.1.4.1.5655.4.1.8.3.1.6), page B-62
- spvPropertyCounter64Value (1.3.6.1.4.1.5655.4.1.8.3.1.7), page B-62
- tpInfoTable (1.3.6.1.4.1.5655.4.1.9.1), page B-62
- tpInfoEntry (1.3.6.1.4.1.5655.4.1.9.1.1), page B-63
- tpModuleIndex (1.3.6.1.4.1.5655.4.1.9.1.1.1), page B-63
- tpIndex (1.3.6.1.4.1.5655.4.1.9.1.1.2), page B-63
- tpTotalNumHandledPackets (1.3.6.1.4.1.5655.4.1.9.1.1.3), page B-64
- tpTotalNumHandledFlows (1.3.6.1.4.1.5655.4.1.9.1.1.4), page B-64
- tpNumActiveFlows (1.3.6.1.4.1.5655.4.1.9.1.1.5), page B-64
- tpNumActiveFlowsPeak (1.3.6.1.4.1.5655.4.1.9.1.1.6), page B-64
- tpNumActiveFlowsPeakTime (1.3.6.1.4.1.5655.4.1.9.1.1.7), page B-65
- tpNumTcpActiveFlows (1.3.6.1.4.1.5655.4.1.9.1.1.8), page B-65
- TpNumTcpActiveFlowsPeak (1.3.6.1.4.1.5655.4.1.9.1.1.9), page B-65
- tpNumTcpActiveFlowsPeakTime (1.3.6.1.4.1.5655.4.1.9.1.1.10), page B-65
- tpNumUdpActiveFlows (1.3.6.1.4.1.5655.4.1.9.1.1.11), page B-66
- tpNumUdpActiveFlowsPeak (1.3.6.1.4.1.5655.4.1.9.1.1.12), page B-66
- tpNumUdpActiveFlowsPeakTime (1.3.6.1.4.1.5655.4.1.9.1.1.13), page B-66
- tpNumNonTcpUdpActiveFlows (1.3.6.1.4.1.5655.4.1.9.1.1.14), page B-66
- tpNumNonTcpUdpActiveFlowsPeak (1.3.6.1.4.1.5655.4.1.9.1.1.15), page B-67
- tpNumNonTcpUdpActiveFlowsPeakTime (1.3.6.1.4.1.5655.4.1.9.1.1.16), page B-67
- tpTotalNumBlockedPackets (1.3.6.1.4.1.5655.4.1.9.1.1.17), page B-67
- tpTotalNumBlockedFlows (1.3.6.1.4.1.5655.4.1.9.1.1.18), page B-67
- tpTotalNumDiscardedPacketsDueToBwLimit (1.3.6.1.4.1.5655.4.1.9.1.1.19), page B-68
- tpTotalNumWredDiscardedPackets (1.3.6.1.4.1.5655.4.1.9.1.1.20), page B-68
- tpTotalNumFragments (1.3.6.1.4.1.5655.4.1.9.1.1.21), page B-68
- tpTotalNumNonIpPackets (1.3.6.1.4.1.5655.4.1.9.1.1.22), page B-68
- tpTotalNumIpCrcErrPackets (1.3.6.1.4.1.5655.4.1.9.1.1.23), page B-69
- tpTotalNumIpLengthErrPackets (1.3.6.1.4.1.5655.4.1.9.1.1.24), page B-69
- tpTotalNumIpBroadcastPackets (1.3.6.1.4.1.5655.4.1.9.1.1.25), page B-69
- tpTotalNumTtlErrPackets (1.3.6.1.4.1.5655.4.1.9.1.1.26), page B-69
- tpTotalNumTcpUdpCrcErrPackets (1.3.6.1.4.1.5655.4.1.9.1.1.27), page B-70
- tpClearCountersTime (1.3.6.1.4.1.5655.4.1.9.1.1.28), page B-70
- tpHandledPacketsRate (1.3.6.1.4.1.5655.4.1.9.1.1.29), page B-70
- tpHandledPacketsRatePeak (1.3.6.1.4.1.5655.4.1.9.1.1.30), page B-70
- tpHandledPacketsRatePeakTime (1.3.6.1.4.1.5655.4.1.9.1.1.31), page B-71

- [tpHandledFlowsRate \(1.3.6.1.4.1.5655.4.1.9.1.1.32\)](#), page B-71
- [tpHandledFlowsRatePeak \(1.3.6.1.4.1.5655.4.1.9.1.1.33\)](#), page B-71
- [tpHandledFlowsRatePeakTime \(1.3.6.1.4.1.5655.4.1.9.1.1.34\)](#), page B-71
- [tpCpuUtilization \(1.3.6.1.4.1.5655.4.1.9.1.1.35\)](#), page B-72
- [tpCpuUtilizationPeak \(1.3.6.1.4.1.5655.4.1.9.1.1.36\)](#), page B-72
- [tpCpuUtilizationPeakTime \(1.3.6.1.4.1.5655.4.1.9.1.1.37\)](#), page B-72
- [tpFlowsCapacityUtilization \(1.3.6.1.4.1.5655.4.1.9.1.1.38\)](#), page B-72
- [tpFlowsCapacityUtilizationPeak \(1.3.6.1.4.1.5655.4.1.9.1.1.39\)](#), page B-73
- [tpFlowsCapacityUtilizationPeakTime \(1.3.6.1.4.1.5655.4.1.9.1.1.40\)](#), page B-73
- [tpServiceLoss \(1.3.6.1.4.1.5655.4.1.9.1.1.41\)](#), page B-73
- [pportTable \(1.3.6.1.4.1.5655.4.1.10.1\)](#), page B-73
- [pportEntry \(1.3.6.1.4.1.5655.4.1.10.1.1\)](#), page B-74
- [pportModuleIndex \(1.3.6.1.4.1.5655.4.1.10.1.1.1\)](#), page B-74
- [pportIndex \(1.3.6.1.4.1.5655.4.1.10.1.1.2\)](#), page B-74
- [pportType \(1.3.6.1.4.1.5655.4.1.10.1.1.3\)](#), page B-75
- [pportNumTxQueues \(1.3.6.1.4.1.5655.4.1.10.1.1.4\)](#), page B-75
- [pportIfIndex \(1.3.6.1.4.1.5655.4.1.10.1.1.5\)](#), page B-75
- [pportAdminSpeed \(1.3.6.1.4.1.5655.4.1.10.1.1.6\)](#), page B-75
- [pportAdminDuplex \(1.3.6.1.4.1.5655.4.1.10.1.1.7\)](#), page B-76
- [pportOperDuplex \(1.3.6.1.4.1.5655.4.1.10.1.1.8\)](#), page B-76
- [pportLinkIndex \(1.3.6.1.4.1.5655.4.1.10.1.1.9\)](#), page B-76
- [pportOperStatus \(1.3.6.1.4.1.5655.4.1.10.1.1.10\)](#), page B-76
- [txQueuesTable \(1.3.6.1.4.1.5655.4.1.11.1\)](#), page B-77
- [txQueuesEntry \(1.3.6.1.4.1.5655.4.1.11.1.1\)](#), page B-77
- [txQueuesModuleIndex \(1.3.6.1.4.1.5655.4.1.11.1.1.1\)](#), page B-77
- [txQueuesPortIndex \(1.3.6.1.4.1.5655.4.1.11.1.1.2\)](#), page B-77
- [txQueuesQueueIndex \(1.3.6.1.4.1.5655.4.1.11.1.1.3\)](#), page B-78
- [txQueuesDescription \(1.3.6.1.4.1.5655.4.1.11.1.1.4\)](#), page B-78
- [txQueuesBandwidth \(1.3.6.1.4.1.5655.4.1.11.1.1.5\)](#), page B-78
- [txQueuesUtilization \(1.3.6.1.4.1.5655.4.1.11.1.1.6\)](#), page B-78
- [txQueuesUtilizationPeak \(1.3.6.1.4.1.5655.4.1.11.1.1.7\)](#), page B-79
- [txQueuesUtilizationPeakTime \(1.3.6.1.4.1.5655.4.1.11.1.1.8\)](#), page B-79
- [txQueuesClearCountersTime \(1.3.6.1.4.1.5655.4.1.11.1.1.9\)](#), page B-79
- [txQueuesDroppedBytes \(1.3.6.1.4.1.5655.4.1.11.1.1.10\)](#), page B-79
- [globalControllersTable \(1.3.6.1.4.1.5655.4.1.12.1\)](#), page B-80
- [globalControllersEntry \(1.3.6.1.4.1.5655.4.1.12.1.1\)](#), page B-80
- [globalControllersModuleIndex \(1.3.6.1.4.1.5655.4.1.12.1.1.1\)](#), page B-80
- [globalControllersPortIndex \(1.3.6.1.4.1.5655.4.1.12.1.1.2\)](#), page B-81

- [globalControllersIndex \(1.3.6.1.4.1.5655.4.1.12.1.1.3\)](#), page B-81
- [globalControllersDescription \(1.3.6.1.4.1.5655.4.1.12.1.1.4\)](#), page B-81
- [globalControllersBandwidth \(1.3.6.1.4.1.5655.4.1.12.1.1.5\)](#), page B-81
- [globalControllersUtilization \(1.3.6.1.4.1.5655.4.1.12.1.1.6\)](#), page B-82
- [globalControllersUtilizationPeak \(1.3.6.1.4.1.5655.4.1.12.1.1.7\)](#), page B-82
- [globalControllersUtilizationPeakTime \(1.3.6.1.4.1.5655.4.1.12.1.1.8\)](#), page B-82
- [globalControllersClearCountersTime \(1.3.6.1.4.1.5655.4.1.12.1.1.9\)](#), page B-82
- [globalControllersDroppedBytes \(1.3.6.1.4.1.5655.4.1.12.1.1.10\)](#), page B-83
- [appInfoTable \(1.3.6.1.4.1.5655.4.1.13.1\)](#), page B-83
- [appInfoEntry \(1.3.6.1.4.1.5655.4.1.13.1.1\)](#), page B-83
- [appName \(1.3.6.1.4.1.5655.4.1.13.1.1.1\)](#), page B-83
- [appDescription \(1.3.6.1.4.1.5655.4.1.13.1.1.2\)](#), page B-84
- [appVersion \(1.3.6.1.4.1.5655.4.1.13.1.1.3\)](#), page B-84
- [appPropertiesTable \(1.3.6.1.4.1.5655.4.1.13.2\)](#), page B-84
- [appPropertiesEntry \(1.3.6.1.4.1.5655.4.1.13.2.1\)](#), page B-84
- [apIndex \(1.3.6.1.4.1.5655.4.1.13.2.1.1\)](#), page B-85
- [apName \(1.3.6.1.4.1.5655.4.1.13.2.1.2\)](#), page B-85
- [apType \(1.3.6.1.4.1.5655.4.1.13.2.1.3\)](#), page B-85
- [appPropertiesValuesTable \(1.3.6.1.4.1.5655.4.1.13.3\)](#), page B-86
- [appPropertiesValueEntry \(1.3.6.1.4.1.5655.4.1.13.3.1\)](#), page B-86
- [apvIndex \(1.3.6.1.4.1.5655.4.1.13.3.1.1\)](#), page B-86
- [apvPropertyName \(1.3.6.1.4.1.5655.4.1.13.3.1.2\)](#), page B-87
- [apvRowStatus \(1.3.6.1.4.1.5655.4.1.13.3.1.3\)](#), page B-87
- [apvPropertyStringValue \(1.3.6.1.4.1.5655.4.1.13.3.1.4\)](#), page B-87
- [apvPropertyUintValue \(1.3.6.1.4.1.5655.4.1.13.3.1.5\)](#), page B-87
- [apvPropertyCounter64Value \(1.3.6.1.4.1.5655.4.1.13.3.1.6\)](#), page B-88
- [trafficCountersTable \(1.3.6.1.4.1.5655.4.1.14.1\)](#), page B-88
- [trafficCountersEntry \(1.3.6.1.4.1.5655.4.1.14.1.1\)](#), page B-88
- [trafficCounterIndex \(1.3.6.1.4.1.5655.4.1.14.1.1.1\)](#), page B-88
- [trafficCounterValue \(1.3.6.1.4.1.5655.4.1.14.1.1.2\)](#), page B-89
- [trafficCounterName \(1.3.6.1.4.1.5655.4.1.14.1.1.3\)](#), page B-89
- [trafficCounterType \(1.3.6.1.4.1.5655.4.1.14.1.1.4\)](#), page B-89
- [attackTypeTable \(1.3.6.1.4.1.5655.4.1.15.1\)](#), page B-89
- [attackTypeEntry \(1.3.6.1.4.1.5655.4.1.15.1.1\)](#), page B-90
- [attackTypeIndex \(1.3.6.1.4.1.5655.4.1.15.1.1.1\)](#), page B-90
- [attackTypeName \(1.3.6.1.4.1.5655.4.1.15.1.1.2\)](#), page B-90
- [attackTypeCurrentNumAttacks \(1.3.6.1.4.1.5655.4.1.15.1.1.3\)](#), page B-90
- [attackTypeTotalNumAttacks \(1.3.6.1.4.1.5655.4.1.15.1.1.4\)](#), page B-91

- [attackTypeTotalNumFlows \(1.3.6.1.4.1.5655.4.1.15.1.1.5\)](#), page B-91
- [attackTypeTotalNumSeconds \(1.3.6.1.4.1.5655.4.1.15.1.1.6\)](#), page B-91
- [vasServersTable \(1.3.6.1.4.1.5655.4.1.16.1\)](#), page B-91
- [vasServerEntry \(1.3.6.1.4.1.5655.4.1.16.1.1\)](#), page B-92
- [vasServerIndex \(1.3.6.1.4.1.5655.4.1.16.1.1.1\)](#), page B-92
- [vasServerId \(1.3.6.1.4.1.5655.4.1.16.1.1.2\)](#), page B-92
- [vasServerAdminStatus \(1.3.6.1.4.1.5655.4.1.16.1.1.3\)](#), page B-92
- [vasServerOperStatus \(1.3.6.1.4.1.5655.4.1.16.1.1.4\)](#), page B-93
- [mplsVpnSoftwareCountersTable \(1.3.6.1.4.1.5655.4.1.17.1\)](#), page B-93
- [mplsVpnSoftwareCountersEntry \(1.3.6.1.4.1.5655.4.1.17.1.1\)](#), page B-93
- [mplsVpnMaxHWMappings \(1.3.6.1.4.1.5655.4.1.17.1.1.1\)](#), page B-93
- [mplsVpnCurrentHWMappings \(1.3.6.1.4.1.5655.4.1.17.1.1.2\)](#), page B-94

sysOperationalStatus (1.3.6.1.4.1.5655.4.1.1.1)

Indicates the operational status of the system.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

```

INTEGER { 1 (other)- none of the following 2 (boot)- the system is in boot process
3 (operational)- the system is operational 4 (warning)- the system is in Warning status
5 (failure)- the system is in Failure status
}

```

sysFailureRecovery (1.3.6.1.4.1.5655.4.1.1.2)

Indicates the behavior of the system after abnormal boot.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

```

INTEGER { 1 (other)- none of the following 2 (operational)- the system should enter
Operational mode after abnormal boot 3 (non-operational)- the system should enter Failure
mode after abnormal boot
}

```

sysVersion (1.3.6.1.4.1.5655.4.1.1.3)

The system version.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

`DisplayString`

pchassisSysType (1.3.6.1.4.1.5655.4.1.2.1)

The chassis system type.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

`INTEGER` {**1** (*other*)– none of the following**2** (*SE1000*)– SCE 1000 platform**3** (*SE100*)– SCE 100 platform**4** (*SE2000*)– SCE 2000 platform
}

pchassisPowerSupplyAlarm (1.3.6.1.4.1.5655.4.1.2.2)

Indicates whether the power supply to the chassis is normal. If the alarm is ‘on’, it means that one or more of the power supplies is not functional

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

`INTEGER` {**1** (*other*)– none of the following**2** (*off*)– the power supply to the chassis is normal**3** (*on*)– the power supply to the chassis is not normal, and probably one or more of the power supplies is not functional.
}

pchassisFansAlarm (1.3.6.1.4.1.5655.4.1.2.3)

Indicates whether all the fans on the chassis are functional.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

`INTEGER` {**1** (*other*)– none of the following**2** (*off*)– all fans are functional**3** (*on*)– one or more fans is not functional.
}

pchassisTempAlarm (1.3.6.1.4.1.5655.4.1.2.4)

Indicates the chassis temperature alarm status.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

```

INTEGER {1 (other)- none of the following2 (off)- temperature is within acceptable
range3 (on)- temperature is too high.
}

```

pchassisVoltageAlarm (1.3.6.1.4.1.5655.4.1.2.5)

Indicates the chassis internal voltage alarm status. If the alarm is 'on', it indicates that the voltage level of one or more unit in the chassis is not in the normal range.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

```

INTEGER {1 (other)- none of the following2 (off)- voltage level is within normal range3
(on)- voltage level is out of the acceptable bounds.
}

```

pchassisNumSlots (1.3.6.1.4.1.5655.4.1.2.6)

Indicates the number of slots in the chassis available for plug-in modules, including both currently occupied and empty slots.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

```

INTEGER (0..255)

```


pchassisSlotConfig (1.3.6.1.4.1.5655.4.1.2.7)

An indication of which slots in the chassis are occupied.

This is an integer value with bits set to indicate configured modules. It is expressed as the function:

Sum of $f(x)$ as x goes from 1 to the number of slots, where:

- no module inserted — $f(x) = 0$
- module inserted — $f(x) = \exp(2, x-1)$

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

INTEGER (0..65535)

pchassisPsuType (1.3.6.1.4.1.5655.4.1.2.8)

Indicates the type of the power supplies.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

INTEGER {**1** (other)— none of the following**2** (AC)— AC power supply**3** (DC)— DC power supply
}

pchassisLineFeedAlarm (1.3.6.1.4.1.5655.4.1.2.9)

Indicates whether the line feed to the chassis is connected and whether it is supplying power to the power supply unit.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

INTEGER {**1** (other)— none of the following**2** (OFF)— The line feed to the chassis is connected and has power**3** (ON)— The line feed to the chassis is not normal. One or both of the line feeds may not be connected properly or have no power.
}

pmoduleTable (1.3.6.1.4.1.5655.4.1.3.1)

A list of module entries containing information defining the modules in the chassis.

The number of entries is the number of modules in the chassis.

| | |
|--------|----------------|
| Access | not-accessible |
|--------|----------------|

Syntax

Sequence of pmoduleEntry

pmoduleEntry (1.3.6.1.4.1.5655.4.1.3.1.1)

Entry containing a number of parameters defining the physical characteristics of one module in the chassis.

| | |
|--------|----------------|
| Access | not-accessible |
|--------|----------------|

Index

{pmoduleIndex}

Syntax**SEQUENCE**

*{pmoduleIndexpmoduleTypepmoduleNumTrafficProcessorspmoduleSlotNumpmoduleHwVersionpmoduleNumPorts
pmoduleNumLinkspmoduleConnectionModepmoduleSerialNumberpmoduleUpStreamAttackFilteringTime
pmoduleUpStreamLastAttackFilteringTimepmoduleDownStreamAttackFilteringTimepmoduleDown
StreamLastAttackFilteringTimepmoduleAttackObjectsClearTimepmoduleAdminStatuspmoduleOperSta
tus}*

pmoduleIndex (1.3.6.1.4.1.5655.4.1.3.1.1.1)

An ID number identifying the module. A unique value for each module within the chassis.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

INTEGER (1..255)

pmoduleType (1.3.6.1.4.1.5655.4.1.3.1.1.2)

The type of module.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

```

INTEGER { 1 (other)- none of the following 2 (gbe2Module)- 2 port Gigabit Ethernet line
interface, 2 Fast Ethernet 10/100 management interfaces 3 (fe2Module)- 2 port Fast
Ethernet line interface, 1 Fast Ethernet 10/100 management interface 4 (gbe4Module)- 4
port Gigabit Ethernet line interface, 2 Fast Ethernet 10/100 management interfaces 5
(fe4Module)- 4 port Fast Ethernet line interface, 2 Fast Ethernet 10/100 management
interfaces 6 (oc12-4Module)- 4 port OC12 line interface, 2 Fast Ethernet 10/100 management
interfaces 7 (fe8Module)- 8 port Fast Ethernet line interface, 2 Fast Ethernet 10/100
management interfaces
}

```

pmoduleNumTrafficProcessors (1.3.6.1.4.1.5655.4.1.3.1.1.3)

The number of traffic processors supported by the module.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

```

INTEGER (0..255)

```

pmoduleSlotNum (1.3.6.1.4.1.5655.4.1.3.1.1.4)

The number of the slot in the chassis in which the module is installed.

Valid entries are from 1 to the value of **pchassisNumSlot**s.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

```

INTEGER (1..255)

```

pmoduleHwVersion (1.3.6.1.4.1.5655.4.1.3.1.1.5)

The hardware version of the module.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

`DisplayString`

pmoduleNumPorts (1.3.6.1.4.1.5655.4.1.3.1.1.6)

The number of ports supported by the module.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

`INTEGER (0..255)`

pmoduleNumLinks (1.3.6.1.4.1.5655.4.1.3.1.1.7)

The number of links carrying inband traffic that are supported by the module. The link is uniquely defined by the two ports that are at its endpoints.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

`INTEGER (0..255)`

pmoduleConnectionMode (1.3.6.1.4.1.5655.4.1.3.1.1.8)

Indicates the connection mode of the module.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

`INTEGER {1 (other)- none of the following2 (inline)- SCE is both receiving and transmitting traffic on the line ports.3 (receive-only)- SCE can only receive packets from the line ports. This mode is suitable for external splitting topology.4 (inline-cascade)- SCE is both receiving and transmitting traffic on the line ports and the cascade ports.5 (receive-only-cascade)- SCE can only receive packets from the line and the cascade ports. This mode is suitable for external splitting topology }`

pmoduleSerialNumber (1.3.6.1.4.1.5655.4.1.3.1.1.9)

The serial number of the module.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

`DisplayString`

pmoduleUpStreamAttackFilteringTime (1.3.6.1.4.1.5655.4.1.3.1.1.10)

The accumulated time (in hundredths of a second) during which attack up-stream traffic was filtered.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

`TimeTicks`

pmoduleUpStreamLastAttackFilteringTime (1.3.6.1.4.1.5655.4.1.3.1.1.11)

The time (in hundredths of a second) since the previous attack filtered in the up-stream traffic.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

`TimeTicks`

pmoduleDownStreamAttackFilteringTime (1.3.6.1.4.1.5655.4.1.3.1.1.12)

The accumulated time (in hundredths of a second) during which attack down-stream traffic was filtered.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

`TimeTicks`

pmoduleDownStreamLastAttackFilteringTime (1.3.6.1.4.1.5655.4.1.3.1.1.13)

The time (in hundredths of a second) since the previous attack filtered in the down-stream traffic.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

TimeTicks

pmoduleAttackObjectsClearTime (1.3.6.1.4.1.5655.4.1.3.1.1.14)

The time (in hundredths of a second) since the attack objects were cleared. Writing a 0 to this object causes the counters to be cleared.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

TimeTicks

pmoduleAdminStatus (1.3.6.1.4.1.5655.4.1.3.1.1.15)

Indicates whether the module is configured to handle traffic on startup or reboot (active), to be the hot standby.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

```
INTEGER {1 (other)- none of the following2 (primary)- Handle traffic on startup.3
(secondary)- Fail-over module on startup.
}
```

pmoduleOperStatus (1.3.6.1.4.1.5655.4.1.3.1.1.16)

Indicates whether the module is currently handling (active), or is on standby.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

```
INTEGER {1 (other)- none of the following2 (active)- Currently is handling traffic.3
(standby)- Currently is the fail-over module.
}
```

linkTable (1.3.6.1.4.1.5655.4.1.4.1)

A list of link entries containing information regarding the configuration and status of the links that pass through the SCE platform and carry in-band traffic.

The number of entries is determined by the number of modules in the chassis and the number of links on each module.

| | |
|--------|----------------|
| Access | not-accessible |
|--------|----------------|

Syntax

Sequence of linkEntry

linkEntry (1.3.6.1.4.1.5655.4.1.4.1.1)

Entry containing information about the Link.

| | |
|--------|----------------|
| Access | not-accessible |
|--------|----------------|

Index

{linkModuleIndex, linkIndex}

Syntax

SEQUENCE

{linkModuleIndexlinkIndexlinkAdminModeOnActivelinkAdminModeOnFailurelinkOperModelinkStatusReflectionEnablelinkSubscriberSidePortIndexlinkNetworkSidePortIndex}

linkModuleIndex (1.3.6.1.4.1.5655.4.1.4.1.1.1)

An index value (**pmoduleIndex**) that uniquely identifies the module where this link is located.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

INTEGER (1..255)

linkIndex (1.3.6.1.4.1.5655.4.1.4.1.1.2)

An index value that uniquely identifies the link within the specified module.

Valid entries are 1 to the value of **pmoduleNumLinks** for this module.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

INTEGER (1..255)

linkAdminModeOnActive (1.3.6.1.4.1.5655.4.1.4.1.1.3)

The desired mode of the link when the operating status of the module is active and it is not in boot or failure.

Possible values (*LinkModeType*):

- *Bypass*— the traffic is forwarded from one port to the other using an internal splitter.
- *Forwarding*— the traffic is forwarded by the internal hardware and software modules of the SCE platform.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

LinkModeType

linkAdminModeOnFailure (1.3.6.1.4.1.5655.4.1.4.1.1.4)

The desired mode of the link when the system status is failure.

Possible values (*LinkModeType*):

- *Bypass*— the traffic is forwarded from one port to the other using an internal splitter.
- *Cutoff*— all traffic is dropped by the SCE.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

LinkModeType

linkOperMode (1.3.6.1.4.1.5655.4.1.4.1.1.5)

The current operational mode of the link.

Possible values (*LinkModeType*):

- *Bypass*— the traffic is forwarded from one port to the other using an internal splitter with no processing taking place.
- *Forwarding*— the traffic is forwarded by the internal hardware and software modules of the SCE.
- *Sniffing*— the traffic is forwarded in the same manner as in Bypass mode, however it passes through and is analysed by the internal software and hardware modules of the SCE platform.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

LinkModeType

linkStatusReflectionEnable (1.3.6.1.4.1.5655.4.1.4.1.1.6)

Indicates whether failure of the physical link on one interface should trigger the failure of the link on the other interface on the module.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

INTEGER {1 (enabled)2 (disabled)}

linkSubscriberSidePortIndex (1.3.6.1.4.1.5655.4.1.4.1.1.7)

An index value that uniquely identifies this link with the related port that is connected to the subscriber side.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

INTEGER (0..255)

linkNetworkSidePortIndex (1.3.6.1.4.1.5655.4.1.4.1.1.8)

An index value that uniquely identifies this link with the related port that is connected to the network side.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

INTEGER (0..255)

diskNumUsedBytes (1.3.6.1.4.1.5655.4.1.5.1)

The number of used bytes on the disk.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Unsigned32 (0...4294967295)

diskNumFreeBytes (1.3.6.1.4.1.5655.4.1.5.2)

The number of free bytes on the disk.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Unsigned32 (0...4294967295)

rdrFormatterEnable (1.3.6.1.4.1.5655.4.1.6.1)

Indicates whether the RDR-formatter is enabled or disabled.

When the RDR-formatter is enabled, it sends the reports it gets from the traffic processors to the Collection Manager as defined in the rdrFormatterDestTable.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

INTEGER {1 (enabled) 2 (disabled)}

rdrFormatterDestTable (1.3.6.1.4.1.5655.4.1.6.2)

This table lists the addresses of Collection Managers.

If the RDR-formatter is enabled, the destination with the highest priority to which a TCP connection can be established is designated as the active connection, and would receive the reports generated by the traffic processors.

The table may contain a maximum of three entries.

| | |
|--------|----------------|
| Access | not-accessible |
|--------|----------------|

Syntax

Sequence of rdrFormatterDestEntry

rdrFormatterDestEntry (1.3.6.1.4.1.5655.4.1.6.2.1)

Entry defining one RDR destination.

| | |
|--------|----------------|
| Access | not-accessible |
|--------|----------------|

Index

{rdrFormatterDestIPAddr, rdrFormatterDestPort}

Syntax

SEQUENCE

{rdrFormatterDestIPAddr, rdrFormatterDestPort, rdrFormatterDestPriority, rdrFormatterDestStatus, rdrFormatterDestConnectionStatus, rdrFormatterDestNumReportsSent, rdrFormatterDestNumReportsDiscarded, rdrFormatterDestReportRate, rdrFormatterDestReportRatePeak, rdrFormatterDestReportRatePeakTime}

rdrFormatterDestIPAddr (1.3.6.1.4.1.5655.4.1.6.2.1.1)

The IP address of a Collection Manager.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

IP Address

rdrFormatterDestPort (1.3.6.1.4.1.5655.4.1.6.2.1.2)

The TCP port on which the Collection Manager listens and the to which the RDR-Formatter should connect.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

INTEGER (1...65535)

rdrFormatterDestPriority (1.3.6.1.4.1.5655.4.1.6.2.1.3)

The priority given to the Collection Manager. The active Collection Manager is the Collection Manager with the highest priority whose TCP connection is up.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

INTEGER (1...100)

rdrFormatterDestStatus (1.3.6.1.4.1.5655.4.1.6.2.1.4)

Indicates whether this destination is the active one.

In redundancy and simple-load-balancing modes there can be only one 'active' destination, which is the one to which the reports are sent. In multicast mode all destinations receive the active mode.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

INTEGER {**1** (*other*)– none of the following**2** (*active*)– this destination is where the reports are sent**3** (*standby*)– this destination is a backup
}

rdrFormatterDestConnectionStatus (1.3.6.1.4.1.5655.4.1.6.2.1.5)

The status of TCP connection to this destination.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

```

INTEGER {1 (other)- none of the following2 (up)- the TCP connection to this destination
is up3 (down)- the TCP connection to this destination is down
}

```

rdrFormatterDestNumReportsSent (1.3.6.1.4.1.5655.4.1.6.2.1.6)

The number of reports sent by the RDR-formatter to this destination.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

```

Unsigned32 (0...4294967295)

```

rdrFormatterDestNumReportsDiscarded (1.3.6.1.4.1.5655.4.1.6.2.1.7)

The number of reports dropped by the RDR-formatter at this destination.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

```

Unsigned32 (0...4294967295)

```

rdrFormatterDestReportRate (1.3.6.1.4.1.5655.4.1.6.2.1.8)

The current rate (in reports per second) of sending reports to this destination.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

```

Unsigned32 (0...4294967295)

```

rdrFormatterDestReportRatePeak (1.3.6.1.4.1.5655.4.1.6.2.1.9)

The maximum rate of sending reports to this destination.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Unsigned32 (0...4294967295)

rdrFormatterDestReportRatePeakTime (1.3.6.1.4.1.5655.4.1.6.2.1.10)

The time (in hundredths of a second) since the **rdrFormatterDestReportRatePeak** value occurred.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

TimeTicks

rdrFormatterNumReportsSent (1.3.6.1.4.1.5655.4.1.6.3)

The number of reports sent by the RDR-formatter.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Unsigned32 (0...4294967295)

rdrFormatterNumReportsDiscarded (1.3.6.1.4.1.5655.4.1.6.4)

The number of reports dropped by the RDR-formatter.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Unsigned32 (0...4294967295)

rdrFormatterClearCountersTime (1.3.6.1.4.1.5655.4.1.6.5)

The time (in hundredths of a second) since the RDR-formatter counters were last cleared. Writing a 0 to this object causes the RDR-formatter counters to be cleared.

| | |
|--------|----|
| Access | RW |
|--------|----|

Syntax

TimeTicks

rdrFormatterReportRate (1.3.6.1.4.1.5655.4.1.6.6)

The current rate (in reports per second) of sending reports to all destinations.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Unsigned32 (0...4294967295)

rdrFormatterReportRatePeak (1.3.6.1.4.1.5655.4.1.6.7)

The maximum rate of sending reports to all destinations.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Unsigned32 (0...4294967295)

rdrFormatterReportRatePeakTime (1.3.6.1.4.1.5655.4.1.6.8)

The time (in hundredths of a second) since the **rdrFormatterReportRatePeak** value occurred.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

TimeTicks

rdRFormatterProtocol (1.3.6.1.4.1.5655.4.1.6.9)

The RDR protocol currently in use.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

```

INTEGER { 1 (other)- none of the following 2 (RDRv1)- RDR protocol version 13 (RDRv2)-
RDR protocol version 2
}

```

rdRFormatterForwardingMode (1.3.6.1.4.1.5655.4.1.6.10)

The manner in which the RDR formatter sends the reports to the destinations.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

```

INTEGER { 1 (other)- none of the following 2 (redundancy)- all RDRs are sent to the
primary (active) destination, and all other destinations are in standby 3
(simpleLoadBalancing)- each successive RDR is sent to a different destination, one
destination after the other, in a round robin manner 4 (multicast)- all RDRs are sent to
all destinations
}

```

rdRFormatterCategoryTable (1.3.6.1.4.1.5655.4.1.6.11)

This table describes the different categories of RDRs and supplies some statistical information about the RDRs sent to these categories

| | |
|--------|----------------|
| Access | not-accessible |
|--------|----------------|

Syntax

Sequence of rdRFormatterCategoryEntry

rdrFormatterCategoryEntry (1.3.6.1.4.1.5655.4.1.6.11.1)

Entry containing information about the RDR formatter categories.

| | |
|--------|----------------|
| Access | not-accessible |
|--------|----------------|

Index

{rdrFormatterCategoryIndex}

Syntax**SEQUENCE**

{rdrFormatterCategoryIndexrdrFormatterCategoryNamerdrFormatterCategoryNumReportsSentrdrFormatterCategoryNumReportsDiscardedrdrFormatterCategoryReportRaterdrFormatterCategoryReportRatePeakrdrFormatterCategoryReportRatePeakTimerdrFormatterCategoryNumReportsQueued}

rdrFormatterCategoryIndex (1.3.6.1.4.1.5655.4.1.6.11.1.1)

The RDR formatter category number.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

INTEGER (1..4)

rdrFormatterCategoryName (1.3.6.1.4.1.5655.4.1.6.11.1.2)

The name of the category.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

DisplayString

rdrFormatterCategoryNumReportsSent (1.3.6.1.4.1.5655.4.1.6.11.1.3)

The number of reports sent by the RDR-formatter to this category.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Unsigned32 (0...4294967295)

rdrFormatterCategoryNumReportsDiscarded (1.3.6.1.4.1.5655.4.1.6.11.1.4)

The number of reports dropped by the RDR formatter for this category.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Unsigned32 (0...4294967295)

rdrFormatterCategoryReportRate (1.3.6.1.4.1.5655.4.1.6.11.1.5)

The rate of the reports (in reports per second) currently sent to this category.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Unsigned32 (0...4294967295)

rdrFormatterCategoryReportRatePeak (1.3.6.1.4.1.5655.4.1.6.11.1.6)

The maximum report rate sent to this category.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Unsigned32 (0...4294967295)

rdrFormatterCategoryReportRatePeakTime (1.3.6.1.4.1.5655.4.1.6.11.1.7)

The time (in hundredths of a second) since the **rdrFormatterCategoryReportRatePeak** value occurred.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

TimeTicks

rdrFormatterCategoryNumReportsQueued (1.3.6.1.4.1.5655.4.1.6.11.1.8)

The number of pending reports in this category.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Unsigned32 (0...4294967295)

rdrFormatterCategoryDestTable (1.3.6.1.4.1.5655.4.1.6.12)

This table describes the partition of the RDR destinations between the different categories and the priority and status of each destination in each category

| | |
|--------|----------------|
| Access | not-accessible |
|--------|----------------|

Syntax

Sequence of rdrFormatterCategoryDestEntry

rdrFormatterCategoryDestEntry (1.3.6.1.4.1.5655.4.1.6.12.1)

A destination table entry.

| | |
|--------|----------------|
| Access | not-accessible |
|--------|----------------|

Index

{rdrFormatterCategoryIndex, rdrFormatterDestIPAddr, rdrFormatterDestPort}

Syntax

SEQUENCE {rdrFormatterCategoryDestPriorityrdrFormatterCategoryDestStatus}

rdrFormatterCategoryDestPriority (1.3.6.1.4.1.5655.4.1.6.12.1.1)

The priority assigned to the Collection Manager for this category.

The active Collection Manager is the Collection Manager with the highest priority and a TCP connection that is up.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

INTEGER (1...100)

rdrFormatterCategoryDestStatus (1.3.6.1.4.1.5655.4.1.6.12.1.2)

Indicates whether the destination is currently active or standby.

In redundancy and in simple Load Balancing **rdrFormatterForwardingMode** there can be only one active destination, which is where the reports are currently being sent. In multicast mode, all destinations will be assigned the active(2) status.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

```
INTEGER {1 (other)- none of the following 2 (active)- this is the destination to which
reports are currently being sent 3 (standby)- this destination is a backup
}
```

loggerUserLogEnable (1.3.6.1.4.1.5655.4.1.7.1)

Indicates whether the logging of user information is enabled or disabled.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

```
INTEGER {1 (enabled) 2 (disabled)}
```

loggerUserLogNumInfo (1.3.6.1.4.1.5655.4.1.7.2)

The number of Info messages logged into the user log file since last reboot or last time the counter was cleared.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

```
Unsigned32 (0...4294967295)
```

loggerUserLogNumWarning (1.3.6.1.4.1.5655.4.1.7.3)

The number of **Warning** messages logged into the user log file since last reboot or last time the counter was cleared.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Unsigned32 (0...4294967295)

loggerUserLogNumError (1.3.6.1.4.1.5655.4.1.7.4)

The number of **Error** messages logged into the user log file since last reboot or last time the counter was cleared.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Unsigned32 (0...4294967295)

loggerUserLogNumFatal (1.3.6.1.4.1.5655.4.1.7.5)

The number of **Fatal** messages logged into the user log file since last reboot or last time the counter was cleared.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Unsigned32 (0...4294967295)

loggerUserLogClearCountersTime (1.3.6.1.4.1.5655.4.1.7.6)

The time (in hundredths of a second) since user log counters were last cleared.

Writing a 0 to this object causes the user log counters to be cleared.

| | |
|--------|----|
| Access | RW |
|--------|----|

Syntax

TimeTicks

subscribersInfoTable (1.3.6.1.4.1.5655.4.1.8.1)

Data regarding subscriber management operations performed.

| | |
|--------|----------------|
| Access | not-accessible |
|--------|----------------|

Syntax

Sequence of subscribersInfoEntry

subscribersInfoEntry (1.3.6.1.4.1.5655.4.1.8.1.1)

Entry describing the subscriber management operations performed on a certain module.

| | |
|--------|----------------|
| Access | not-accessible |
|--------|----------------|

Index

{pmoduleIndex}

Syntax

SEQUENCE {*subscribersNumIntroduced subscribersNumFree subscribersNumIpAddrMappings subscribersNumIpAddrMappingsFree subscribersNumIpRangeMappings subscribersNumIpRangeMappingsFree subscribersNumVlanMappings subscribersNumVlanMappingsFree subscribersNumActive subscribersNumActivePeak subscribersNumActivePeakTime subscribersNumUpdatessubscribersCountersClearTimesubscriberssubscribersNumTpIpRangeMapping subsubscribersNumTpIpRangeMappingsFreeCountersClearTimesubscribersNumAnonymoussubscribersNum WithSessions*}

subscribersNumIntroduced (1.3.6.1.4.1.5655.4.1.8.1.1.1)

The current number of subscribers introduced to the SCE. These subscribers may or may not have IP address or VLAN mappings. Subscribers who do not have mappings of any kind cannot be associated with traffic, and will be served by the SCE according to the default settings.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Unsigned32 (0...4294967295)

subscribersNumFree (1.3.6.1.4.1.5655.4.1.8.1.1.2)

The number of subscribers that may be introduced in addition to the currently introduced subscribers.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Unsigned32 (0...4294967295)

subscribersNumIpAddrMappings (1.3.6.1.4.1.5655.4.1.8.1.1.3)

The current number of IP address to subscriber mappings.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Unsigned32 (0...4294967295)

subscribersNumIpAddrMappingsFree (1.3.6.1.4.1.5655.4.1.8.1.1.4)

The number of free IP address to subscriber mappings that are available for defining new mappings.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Unsigned32 (0...4294967295)

subscribersNumIpRangeMappings (1.3.6.1.4.1.5655.4.1.8.1.1.5)

The current number of IP-range to subscriber mappings.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Unsigned32 (0...4294967295)

subscribersNumIpRangeMappingsFree (1.3.6.1.4.1.5655.4.1.8.1.1.6)

The number of free IP range to subscriber mappings that are available for defining new mappings.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Unsigned32 (0...4294967295)

subscribersNumVlanMappings (1.3.6.1.4.1.5655.4.1.8.1.1.7)

The current number of VLAN to subscriber mappings.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Unsigned32 (0...4294967295)

subscribersNumVlanMappingsFree (1.3.6.1.4.1.5655.4.1.8.1.1.8)

The number of free VLAN to subscriber mappings that are available for defining new mappings.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Unsigned32 (0...4294967295)

subscribersNumActive (1.3.6.1.4.1.5655.4.1.8.1.1.9)

The current number of active subscribers. These subscribers necessarily have IP address or VLAN mappings that define the traffic to be served according to the subscriber service agreement.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Unsigned32 (0...4294967295)

subscribersNumActivePeak (1.3.6.1.4.1.5655.4.1.8.1.1.10)

The peak value of **subscribersNumActive** since the last time it was cleared or the system started.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Unsigned32 (0...4294967295)

subscribersNumActivePeakTime (1.3.6.1.4.1.5655.4.1.8.1.1.11)

The time (in hundredths of a second) since the **subscribersNumActivePeak** value occurred.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

TimeTicks

subscribersNumUpdates (1.3.6.1.4.1.5655.4.1.8.1.1.12)

The accumulated number of subscribers database updates received by the SCE.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Unsigned32 (0...4294967295)

subscribersCountersClearTime (1.3.6.1.4.1.5655.4.1.8.1.1.13)

The time (in hundredths of a second) since the subscribers counters were cleared.

Writing a 0 to this object causes the counters to be cleared.

| | |
|--------|----|
| Access | RW |
|--------|----|

Syntax

TimeTicks

subscribersNumTplRangeMappings (1.3.6.1.4.1.5655.4.1.8.1.1.14)

The current number of IP range to Traffic Processor mappings.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Unsigned32 (0...4294967295)

subscribersNumTplRangeMappingsFree (1.3.6.1.4.1.5655.4.1.8.1.1.15)

The current number of IP range to Traffic Processor mappings that are available for defining new mappings.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Unsigned32 (0...4294967295)

subscribersNumAnonymous (1.3.6.1.4.1.5655.4.1.8.1.1.16)

The current number of anonymous subscribers.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Unsigned32 (0...4294967295)

subscribersNumWithSessions (1.3.6.1.4.1.5655.4.1.8.1.1.17)

The current number of subscribers with open sessions.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Unsigned32 (0...4294967295)

subscribersPropertiesTable (1.3.6.1.4.1.5655.4.1.8.2)

List of all subscriber properties. This table is updated each time an application is loaded on the SCE platform.

| | |
|--------|----------------|
| Access | not-accessible |
|--------|----------------|

Syntax

Sequence of subscribersPropertiesEntry

subscribersPropertiesEntry (1.3.6.1.4.1.5655.4.1.8.2.1)

Entry describing subscriber properties of the application relevant for a certain module.

| | |
|--------|----------------|
| Access | not-accessible |
|--------|----------------|

Index

{pmoduleIndex, spIndex}

Syntax

SEQUENCE {spIndexspNamespType}

spIndex (1.3.6.1.4.1.5655.4.1.8.2.1.1)

An index value that uniquely identifies the subscriber property.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

INTEGER (1..255)

spName (1.3.6.1.4.1.5655.4.1.8.2.1.2)

Name of the subscriber property.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

DisplayString

spType (1.3.6.1.4.1.5655.4.1.8.2.1.3)

Property type in respect to: variable type (integer, boolean, string etc), number of elements (scalar or array), and restrictions, if any.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

`DisplayString`

subscriberPropertiesValuesTable (1.3.6.1.4.1.5655.4.1.8.3)

The subscriber properties value table is used to provide values for the subscriber properties for a specific subscriber introduced into the SCE platform.

An entry must be created by setting the entry `spvRowStatus` object with `CreateAndGo` (4) before setting the name of the subscriber and the property requested. The property requested must be one of the properties from the `subscribersPropertiesTable`. To remove an entry set the `spvRowStatus` object with `Destroy` (6).

To poll the subscriber property, either of these objects should be polled:

- `spvPropertyStringValue`
- `spvPropertyUnitValue`

The table is cleared when the application is unloaded.

| | |
|--------|----------------|
| Access | not-accessible |
|--------|----------------|

Syntax

Sequence of subscribersPropertiesValueEntry

subscriberPropertiesValueEntry (1.3.6.1.4.1.5655.4.1.8.3.1)

Entry providing information on the value of one of the specified subscriber properties.

| | |
|--------|----------------|
| Access | not-accessible |
|--------|----------------|

Index

{pmoduleIndex, spvIndex}

Syntax

SEQUENCE

{SpvIndexspvSubNamespvPropertyNamespvRowStatuspvPropertyStringValuepvPropertyUnitValuepvPropertyCounter64Value}

spvIndex (1.3.6.1.4.1.5655.4.1.8.3.1.1)

An index value that uniquely identifies the entry.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

INTEGER (1.. 1024)

spvSubName (1.3.6.1.4.1.5655.4.1.8.3.1.2)

A name that uniquely identifies the subscriber.

| | |
|--------|----|
| Access | RC |
|--------|----|

Syntax

DisplayString (Size 1...40)

spvPropertyName (1.3.6.1.4.1.5655.4.1.8.3.1.3)

A name that uniquely identifies the subscriber property.

Array-type properties may be accessed one element at a time in C-like format. (For example: x[1], or y[1][2]).

| | |
|--------|----|
| Access | RC |
|--------|----|

Syntax

DisplayString (Size 1...128)

spvRowStatus (1.3.6.1.4.1.5655.4.1.8.3.1.4)

Controls creation of a table entry. Only setting CreateAndGo (4) and Destroy (6) will change the status of the entry.

| | |
|--------|----|
| Access | RC |
|--------|----|

Syntax

RowStatus

spvPropertyStringValue (1.3.6.1.4.1.5655.4.1.8.3.1.5)

The value of the subscriber property in display string format.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

DisplayString (SIZE 0...128)

spvPropertyUintValue (1.3.6.1.4.1.5655.4.1.8.3.1.6)

The value of the subscriber property in Uint format.

If the property cannot be cast to Uint format, getting this object returns zero.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Unsigned32 (0...4294967295)

spvPropertyCounter64Value (1.3.6.1.4.1.5655.4.1.8.3.1.7)

The value of the subscriber property in Counter64 format.

If the property cannot be cast to Counter64 format, getting this object returns zero.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Counter64

tpInfoTable (1.3.6.1.4.1.5655.4.1.9.1)

The Traffic Processor Info table consists of data regarding traffic handled by the traffic processors, classified by packets and flows.

| | |
|--------|----------------|
| Access | not-accessible |
|--------|----------------|

Syntax

Sequence of tpInfoEntry

tpInfoEntry (1.3.6.1.4.1.5655.4.1.9.1.1)

Entry containing information from the traffic processors.

| | |
|--------|----------------|
| Access | not-accessible |
|--------|----------------|

Index

{tpModuleIndex, tpIndex}

Syntax

SEQUENCE {*tpModuleIndex**tpIndex**tpTotalNumHandledPacketstpTotalNumHandledFlows*
tpNumActiveFlows *tpNumActiveFlowsPeak* *tpNumActiveFlowsPeakTime* *tpNumTcpActiveFlows*
tpNumTcpActiveFlowsPeak *tpNumTcpActiveFlowsPeakTime* *tpNumUdpActiveFlows*
tpNumUdpActiveFlowsPeak *tpNumUdpActiveFlowsPeakTime*
tpNumNonTcpUdpActiveFlowstpNumNonTcpUdpActiveFlowsPeak *tpNumNonTcpUdpActiveFlowsPeakTime*
tpTotalNumBlockedPacketstpTotalNumBlockedFlowstpTotalNumDiscardedPacketsDueToBwLimittpTota
lNumWredDiscardedPacketstpTotalNumFragmentstpTotalNumNonIpPacketstpTotalNumIpCrcErrPackets
tpTotalNumIpLengthErrPacketstpTotalNumIpBroadcastPacketstpTotalNumTtlErrPacketstpTotalNumT
*cpUdpCrcErrPacketstpClearCountersTimetpHandledPacketsRate**tpHandledPacketsRatePeak**tpHandled*
*PacketsRatePeakTime**tpHandledFlowsRate**tpHandledFlowsRatePeak**tpHandledFlowsRatePeakTime**tpCpu*
*Utilization**tpCpuUtilizationPeak**tpCpuUtilizationPeakTime**tpFlowsCapacityUtilization**tpFlowsCa*
*capacityUtilizationPeak**tpFlowsCapacityUtilizationPeakTime**tpServiceLoss*}

tpModuleIndex (1.3.6.1.4.1.5655.4.1.9.1.1.1)

An index value (**pmoduleIndex**) that uniquely identifies the module in which this traffic processor is located.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

INTEGER (1...255)

tpIndex (1.3.6.1.4.1.5655.4.1.9.1.1.2)

An index value that uniquely identifies the traffic processor within the specified module. The value is determined by the location of the traffic processor on the module.

Valid entries are 1 to the value of **pmoduleNumTrafficProcessors** for the specified module.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

INTEGER (1...255)

tpTotalNumHandledPackets (1.3.6.1.4.1.5655.4.1.9.1.1.3)

The accumulated number of packets handled by this traffic processor since last reboot or last time this counter was cleared.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Unsigned32 (0...4294967295)

tpTotalNumHandledFlows (1.3.6.1.4.1.5655.4.1.9.1.1.4)

The accumulated number of flows handled by this traffic processor since last reboot or last time this counter was cleared.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Unsigned32 (0...4294967295)

tpNumActiveFlows (1.3.6.1.4.1.5655.4.1.9.1.1.5)

The number of flows currently being handled by this traffic processor.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Unsigned32 (0...4294967295)

tpNumActiveFlowsPeak (1.3.6.1.4.1.5655.4.1.9.1.1.6)

The peak value of **tpNumActiveFlows** since the last time it was cleared or the system started.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Unsigned32 (0...4294967295)

tpNumActiveFlowsPeakTime (1.3.6.1.4.1.5655.4.1.9.1.1.7)

The time (in hundredths of a second) since the **tpNumActiveFlowsPeak** value occurred.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

TimeTicks

tpNumTcpActiveFlows (1.3.6.1.4.1.5655.4.1.9.1.1.8)

The number of TCP flows currently being handled by this traffic processor.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Unsigned32 (0...4294967295)

TpNumTcpActiveFlowsPeak (1.3.6.1.4.1.5655.4.1.9.1.1.9)

The peak value of **tpNumTcpActiveFlows** since the last time it was cleared or the system started.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Unsigned32 (0...4294967295)

tpNumTcpActiveFlowsPeakTime (1.3.6.1.4.1.5655.4.1.9.1.1.10)

The time (in hundredths of a second) since the **tpNumTcpActiveFlowsPeak** value occurred.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

TimeTicks

tpNumUdpActiveFlows (1.3.6.1.4.1.5655.4.1.9.1.1.11)

The number of UDP flows currently being handled by the traffic processor.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Unsigned32 (0...4294967295)

tpNumUdpActiveFlowsPeak (1.3.6.1.4.1.5655.4.1.9.1.1.12)

The peak value of **tpNumUdpActiveFlows** since the last time it was cleared or the system started.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Unsigned32 (0...4294967295)

tpNumUdpActiveFlowsPeakTime (1.3.6.1.4.1.5655.4.1.9.1.1.13)

The time (in hundredths of a second) since the **tpNumUdpActiveFlowsPeak** value occurred.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

TimeTicks

tpNumNonTcpUdpActiveFlows (1.3.6.1.4.1.5655.4.1.9.1.1.14)

The number of non TCP/UDP flows currently being handled by the traffic processor.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Unsigned32 (0...4294967295)

tpNumNonTcpUdpActiveFlowsPeak (1.3.6.1.4.1.5655.4.1.9.1.1.15)

The peak value of **tpNumNonTcpUdpActiveFlows** since the last time it was cleared or the system started.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Unsigned32 (0...4294967295)

tpNumNonTcpUdpActiveFlowsPeakTime (1.3.6.1.4.1.5655.4.1.9.1.1.16)

The time (in hundredths of a second) since the **tpNumNonTcpUdpActiveFlowsPeak** value occurred.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

TimeTicks

tpTotalNumBlockedPackets (1.3.6.1.4.1.5655.4.1.9.1.1.17)

The accumulated number of packets discarded by the traffic processor according to application blocking rules.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Unsigned32 (0...4294967295)

tpTotalNumBlockedFlows (1.3.6.1.4.1.5655.4.1.9.1.1.18)

The accumulated number of flows discarded by the traffic processor according to application blocking rules.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Unsigned32 (0...4294967295)

tpTotalNumDiscardedPacketsDueToBwLimit (1.3.6.1.4.1.5655.4.1.9.1.1.19)

The accumulated number of packets discarded by the traffic processor due to subscriber bandwidth limitations.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Unsigned32 (0...4294967295)

tpTotalNumWredDiscardedPackets (1.3.6.1.4.1.5655.4.1.9.1.1.20)

The accumulated number of packets discarded by the traffic processor due to congestion in the queues.

The value in this counter is absolute only when accelerate-packet-drops mode is disabled. When accelerate-packet-drops mode is enabled, it provides only a relative value indicating the trend of the number of packet drops, with a factor of approximately 1:6. (See [How to Count Dropped Packets, page 6-20](#)).

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Unsigned32 (0...4294967295)

tpTotalNumFragments (1.3.6.1.4.1.5655.4.1.9.1.1.21)

The accumulated number of fragmented packets handled by the traffic processor.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Unsigned32 (0...4294967295)

tpTotalNumNonIpPackets (1.3.6.1.4.1.5655.4.1.9.1.1.22)

The accumulated number of non IP packets handled by the traffic processor.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Unsigned32 (0...4294967295)

tpTotalNumIpCrcErrPackets (1.3.6.1.4.1.5655.4.1.9.1.1.23)

The accumulated number of packets with IP CRC error handled by the traffic processor.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Unsigned32 (0...4294967295)

tpTotalNumIpLengthErrPackets (1.3.6.1.4.1.5655.4.1.9.1.1.24)

The accumulated number of packets with IP length error handled by the traffic processor.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Unsigned32 (0...4294967295)

tpTotalNumIpBroadcastPackets (1.3.6.1.4.1.5655.4.1.9.1.1.25)

The accumulated number of IP broadcast packets handled by the traffic processor.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Unsigned32 (0...4294967295)

tpTotalNumTtlErrPackets (1.3.6.1.4.1.5655.4.1.9.1.1.26)

The accumulated number of packets with TTL error handled by the traffic processor.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Unsigned32 (0...4294967295)

tpTotalNumTcpUdpCrcErrPackets (1.3.6.1.4.1.5655.4.1.9.1.1.27)

The accumulated number of TCP/UDP packets with CRC error handled by the traffic processor.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Unsigned32 (0...4294967295)

tpClearCountersTime (1.3.6.1.4.1.5655.4.1.9.1.1.28)

The time (in hundredths of a second) since the traffic processor statistics counters were last cleared. Writing a 0 to this object causes the RDR-formatter counters to be cleared.

| | |
|--------|----|
| Access | RW |
|--------|----|

Syntax

TimeTicks

tpHandledPacketsRate (1.3.6.1.4.1.5655.4.1.9.1.1.29)

The rate in packets per second of the packets handled by this traffic processor.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Unsigned32 (0... 4294967295)

tpHandledPacketsRatePeak (1.3.6.1.4.1.5655.4.1.9.1.1.30)

The peak value of **tpHandledPacketsRate** since the last time it was cleared or the system started.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Unsigned32 (0...4294967295)

tpHandledPacketsRatePeakTime (1.3.6.1.4.1.5655.4.1.9.1.1.31)

The time (in hundredths of a second) since the **tpHandledPacketsRatePeak** value occurred.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

TimeTicks

tpHandledFlowsRate (1.3.6.1.4.1.5655.4.1.9.1.1.32)

The rate in flows start per second of the flows handled by this traffic processor.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Unsigned32 (0...4294967295)

tpHandledFlowsRatePeak (1.3.6.1.4.1.5655.4.1.9.1.1.33)

The peak value of **tpHandledFlowsRate** since the last time it was cleared or the system started.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Unsigned32 (0...4294967295)

tpHandledFlowsRatePeakTime (1.3.6.1.4.1.5655.4.1.9.1.1.34)

The time (in hundredths of a second) since the **tpHandledFlowsRatePeak** value occurred.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

TimeTicks

tpCpuUtilization (1.3.6.1.4.1.5655.4.1.9.1.1.35)

The current percentage of CPU utilization.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

INTEGER (1..100)

tpCpuUtilizationPeak (1.3.6.1.4.1.5655.4.1.9.1.1.36)

The peak value of **tpCpuUtilization** since the last time it was cleared or the system started.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

INTEGER (1..100)

tpCpuUtilizationPeakTime (1.3.6.1.4.1.5655.4.1.9.1.1.37)

The time (in hundredths of a second) since the **pCpuUtilizationPeak** value occurred.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

TimeTicks

tpFlowsCapacityUtilization (1.3.6.1.4.1.5655.4.1.9.1.1.38)

The percentage of flows capacity utilization.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

INTEGER (1..100)

tpFlowsCapacityUtilizationPeak (1.3.6.1.4.1.5655.4.1.9.1.1.39)

The peak value of tpFlowsCapacityUtilization since the last time it was cleared or the system started.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

INTEGER (1..100)

tpFlowsCapacityUtilizationPeakTime (1.3.6.1.4.1.5655.4.1.9.1.1.40)

The time (in hundredths of a second) since the **tpFlowsCapacityUtilizationPeak** value occurred.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

TimeTicks

tpServiceLoss (1.3.6.1.4.1.5655.4.1.9.1.1.41)

The relative amount of service loss in this traffic processor, in units of 0.001%, since last reboot or last time this counter was cleared.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

INTEGER (1..100000)

pportTable (1.3.6.1.4.1.5655.4.1.10.1)

A list of port entries.

The number of entries is determined by the number of modules in the chassis and the number of ports on each module.

| | |
|--------|----------------|
| Access | not-accessible |
|--------|----------------|

Syntax

Sequence of pportEntry

pportEntry (1.3.6.1.4.1.5655.4.1.10.1.1)

Entry containing information for a specified port on a module.

| | |
|--------|----------------|
| Access | not-accessible |
|--------|----------------|

Index

{pportModuleIndex, pportIndex}

Syntax**SEQUENCE**

{pportModuleIndex pportIndex pportType pportNumTxQueues pportIfIndex pportAdminSpeed pportAdminDuplex pportOperDuplex pportLinkIndex pportOperStatus}

pportModuleIndex (1.3.6.1.4.1.5655.4.1.10.1.1.1)

An index value (**pmoduleIndex**) that uniquely identifies the module where the port is located.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

INTEGER (1..255)

pportIndex (1.3.6.1.4.1.5655.4.1.10.1.1.2)

An index value that uniquely identifies the port within the specified module. The value is determined by the location of the port on the module.

Valid entries are 1 to the value of **pmoduleNumPorts** for this module.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

INTEGER (1..255)

pportType (1.3.6.1.4.1.5655.4.1.10.1.1.3)

The type of physical layer medium dependent interface on the port.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

```

INTEGER { 1 (other)- none of the following 11 (e100BaseTX)- UTP Fast Ethernet (Cat 5) 28
(e1000BaseSX)- Short Wave fiber Giga Ethernet
}

```

pportNumTxQueues (1.3.6.1.4.1.5655.4.1.10.1.1.4)

The number of transmit queues supported by this port.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

```

INTEGER (1..255)

```

pportIfIndex (1.3.6.1.4.1.5655.4.1.10.1.1.5)

The value of the instance of the ifIndex object, defined in MIB-II, for this port.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

```

INTEGER (1..255)

```

pportAdminSpeed (1.3.6.1.4.1.5655.4.1.10.1.1.6)

The desired speed of the port. The current operational speed of the port can be determined from ifSpeed.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

```

INTEGER { 1 (autoNegotiation)- 10000000 (s10000000)- 10 Mbps 100000000 (s100000000)- 100
Mbps 1000000000 (s1000000000)- 1 Gbps
}

```

pportAdminDuplex (1.3.6.1.4.1.5655.4.1.10.1.1.7)

The desired duplex of the port.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

```
INTEGER {1 (half)2 (full)4 (auto)}
```

pportOperDuplex (1.3.6.1.4.1.5655.4.1.10.1.1.8)

Indicates whether the port is operating in half-duplex or full-duplex.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

```
INTEGER {1 (half)2 (full)}
```

pportLinkIndex (1.3.6.1.4.1.5655.4.1.10.1.1.9)

The **linkIndex** of the link to which this port belongs.

Value of 0 indicates that this port is not associated with any link.

Value of -1 indicates that this port is associated with multiple links.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

```
INTEGER (-1..255)
```

pportOperStatus (1.3.6.1.4.1.5655.4.1.10.1.1.10)

The status of the port. If the port is down, the reason is indicated.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

```
INTEGER {1 (other)- none of the following2 (up)- the port is up3
(reflectionForcingDown)- the port is currently forced down due to the link reflection
mechanism4 (redundancyForcingDown)- the port is currently forced down due to redundancy
reasons5 (otherDown)- the port is down due to other reasons
}
```

txQueuesTable (1.3.6.1.4.1.5655.4.1.11.1)

A list of information for each SCE platform transmit queue.

| | |
|--------|----------------|
| Access | not-accessible |
|--------|----------------|

Syntax

Sequence of txQueuesEntry

txQueuesEntry (1.3.6.1.4.1.5655.4.1.11.1.1)

Entry containing information for a specified SCE transmit queue.

| | |
|--------|----------------|
| Access | not-accessible |
|--------|----------------|

Index

{txQueuesModuleIndex, txQueuesPortIndex, txQueuesQueueIndex}

Syntax**SEQUENCE**

{txQueuesModuleIndextxQueuesPortIndextxQueuesQueueIndextxQueuesDescriptiontxQueuesBandwidthtxQueuesUtilizationtxQueuesUtilizationPeaktxQueuesUtilizationPeakTimetxQueuesClearCounterstime}

txQueuesModuleIndex (1.3.6.1.4.1.5655.4.1.11.1.1.1)

An index value (**pmoduleIndex**) that uniquely identifies the module where the queue is located.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

INTEGER (1..255)

txQueuesPortIndex (1.3.6.1.4.1.5655.4.1.11.1.1.2)

An index value that uniquely identifies the port on which the queue is located.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

INTEGER (1..255)

txQueuesQueueIndex (1.3.6.1.4.1.5655.4.1.11.1.1.3)

An index value that uniquely identifies the queue within the specified port. The value is determined by the location of the queue on the port.

Valid entries are 1 to the value of **pportNumTxQueues** for the specified port.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

INTEGER (1..255)

txQueuesDescription (1.3.6.1.4.1.5655.4.1.11.1.1.4)

Description of the transmit queue.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

DisplayString

txQueuesBandwidth (1.3.6.1.4.1.5655.4.1.11.1.1.5)

The bandwidth in kbps configured for this queue.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

INTEGER (1...1000000)

txQueuesUtilization (1.3.6.1.4.1.5655.4.1.11.1.1.6)

The percentage of bandwidth utilization relative to the to the configured rate.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

INTEGER (0...100)

txQueuesUtilizationPeak (1.3.6.1.4.1.5655.4.1.11.1.1.7)

The peak value of **txQueuesUtilization** since the last time it was cleared or the system started.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

INTEGER (0...100)

txQueuesUtilizationPeakTime (1.3.6.1.4.1.5655.4.1.11.1.1.8)

The time (in hundredths of a second) since the **txQueuesUtilizationPeak** value occurred.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

TimeTicks

txQueuesClearCountersTime (1.3.6.1.4.1.5655.4.1.11.1.1.9)

The time (in hundredths of a second) since the transmit queues statistics counters were last cleared.

Writing a 0 to this object causes the transmit queues counters to be cleared.

| | |
|--------|----|
| Access | RW |
|--------|----|

Syntax

TimeTicks

txQueuesDroppedBytes (1.3.6.1.4.1.5655.4.1.11.1.1.10)

Number of dropped bytes. Valid only if the system is configured to count dropped bytes per TX queue.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Counter64

globalControllersTable (1.3.6.1.4.1.5655.4.1.12.1)

A list of information for each global controller.

| | |
|--------|----------------|
| Access | not-accessible |
|--------|----------------|

Syntax

Sequence of globalControllersEntry

globalControllersEntry (1.3.6.1.4.1.5655.4.1.12.1.1)

Entry containing information for a specified global controller.

| | |
|--------|----------------|
| Access | not-accessible |
|--------|----------------|

Index

{globalControllersModuleIndex, globalControllersPortIndex, globalControllersIndex}

Syntax**SEQUENCE**

{globalControllersModuleIndexglobalControllersPortIndexglobalControllersIndexglobalControllersDescriptionglobalControllersBandwidthglobalControllersUtilizationglobalControllersUtilizationPeakglobalControllersUtilizationPeakTimeglobalControllersClearCounterstimeglobalControllersDroppedBytes}

globalControllersModuleIndex (1.3.6.1.4.1.5655.4.1.12.1.1.1)

An index value (**pmoduleIndex**) that uniquely identifies the module where the Global Controller is located.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

INTEGER (1..255)

globalControllersPortIndex (1.3.6.1.4.1.5655.4.1.12.1.1.2)

An index value that uniquely identifies the port on which the Global Controller is located.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

INTEGER (1..255)

globalControllersIndex (1.3.6.1.4.1.5655.4.1.12.1.1.3)

An index value that uniquely identifies this Global Controller within the specified port.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

INTEGER (1..255)

globalControllersDescription (1.3.6.1.4.1.5655.4.1.12.1.1.4)

Description of the Global Controller.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

DisplayString

globalControllersBandwidth (1.3.6.1.4.1.5655.4.1.12.1.1.5)

The bandwidth in kbps configured for this Global Controller.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

INTEGER (1...1000000)

globalControllersUtilization (1.3.6.1.4.1.5655.4.1.12.1.1.6)

The percentage of bandwidth utilization relative to the configured rate (**globalControllersBandwidth**).

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

INTEGER (0...100)

globalControllersUtilizationPeak (1.3.6.1.4.1.5655.4.1.12.1.1.7)

The peak value of **bwLimitersUtilization** since the last time it was cleared or the system started.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

INTEGER (0...100)

globalControllersUtilizationPeakTime (1.3.6.1.4.1.5655.4.1.12.1.1.8)

The time (in hundredths of a second) since the **globalControllersUtilizationPeak** value occurred.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

TimeTicks

globalControllersClearCountersTime (1.3.6.1.4.1.5655.4.1.12.1.1.9)

The time (in hundredths of a second) since the Global Controller statistics counters were last cleared. Writing a 0 to this object causes the Global Controller counters to be cleared.

| | |
|--------|----|
| Access | RW |
|--------|----|

Syntax

TimeTicks

globalControllersDroppedBytes (1.3.6.1.4.1.5655.4.1.12.1.1.10)

Number of dropped bytes. Valid only if the system is configured to count dropped bytes per global controller.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Counter64

appInfoTable (1.3.6.1.4.1.5655.4.1.13.1)

Information identifying the application that is currently installed in the SCE platform.

| | |
|--------|----------------|
| Access | not-accessible |
|--------|----------------|

Syntax

Sequence of appInfoEntry

appInfoEntry (1.3.6.1.4.1.5655.4.1.13.1.1)

Entry containing identifying information for the application that is currently installed in the SCE platform.

| | |
|--------|----------------|
| Access | not-accessible |
|--------|----------------|

Index

{moduleIndex}

Syntax

SEQUENCE {appNameappDescriptionappVersion}

appName (1.3.6.1.4.1.5655.4.1.13.1.1.1)

Name of the application currently installed in the SCE platform. This object returns an empty string if no application is currently installed.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

DisplayString

appDescription (1.3.6.1.4.1.5655.4.1.13.1.1.2)

Description of the application currently installed in the SCE platform.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

`DisplayString`

appVersion (1.3.6.1.4.1.5655.4.1.13.1.1.3)

Version information for the application currently installed in the SCE platform.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

`DisplayString`

appPropertiesTable (1.3.6.1.4.1.5655.4.1.13.2)

List of all properties available for the application. The table is cleared when the application is unloaded.

| | |
|--------|----------------|
| Access | not-accessible |
|--------|----------------|

Syntax

Sequence of appPropertiesEntry

appPropertiesEntry (1.3.6.1.4.1.5655.4.1.13.2.1)

Entry describing one of the properties available for the application.

| | |
|--------|----------------|
| Access | not-accessible |
|--------|----------------|

Index

{moduleIndex, apIndex}

Syntax

SEQUENCE {apIndexapNameapType}

apIndex (1.3.6.1.4.1.5655.4.1.13.2.1.1)

An index value that uniquely identifies the property.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

INTEGER (1..255)

apName (1.3.6.1.4.1.5655.4.1.13.2.1.2)

Name of the property.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

DisplayString

apType (1.3.6.1.4.1.5655.4.1.13.2.1.3)

Property type in respect to: variable type (integer, boolean, string etc), number of elements (scalar or array), and restrictions, if any.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

DisplayString

appPropertiesValuesTable (1.3.6.1.4.1.5655.4.1.13.3)

The applications properties value table is used to provide specific values for the applications properties.

An entry must be created by setting the entry `apvRowStatus` object with `CreateAndGo (4)` before setting the name of the property requested. The property requested must be one of the properties from the `appPropertiesTable`. To remove an entry set the `apvRowStatus` object with `Destroy (6)`.

To poll the application property, any of these objects should be polled:

- `apvPropertyValue`
- `apvPropertyUnitValue`
- `apvPropertyCounter64` object.

The table is cleared when the application is unloaded.

| | |
|--------|----------------|
| Access | not-accessible |
|--------|----------------|

Syntax

Sequence of `appPropertiesValueEntry`

appPropertiesValueEntry (1.3.6.1.4.1.5655.4.1.13.3.1)

Entry providing information on the value of one of the specified application properties.

| | |
|--------|----------------|
| Access | not-accessible |
|--------|----------------|

Index

{`moduleIndex`, `apvIndex`}

Syntax

SEQUENCE

{`apvIndex``apvPropertyName``apvRowStatus``apvPropertyStringValue``apvPropertyUnitValue``apvPropertyCounter64Value`}

apvIndex (1.3.6.1.4.1.5655.4.1.13.3.1.1)

An index value that uniquely identifies the property.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

INTEGER (1.. 1024)

apvPropertyName (1.3.6.1.4.1.5655.4.1.13.3.1.2)

A name that uniquely identifies the application property.

Array-type properties may be accessed one element at a time in C-like format. (For example: x[1], or y[1][2]).

| | |
|--------|----|
| Access | RC |
|--------|----|

Syntax

DisplayString

apvRowStatus (1.3.6.1.4.1.5655.4.1.13.3.1.3)

Controls creation of a table entry.

| | |
|--------|----|
| Access | RC |
|--------|----|

Syntax

RowStatus

apvPropertyStringValue (1.3.6.1.4.1.5655.4.1.13.3.1.4)

The value of the application property in display string format.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

DisplayString (SIZE 0...128)

apvPropertyUintValue (1.3.6.1.4.1.5655.4.1.13.3.1.5)

The value of the application property in Uint format.

If the property cannot be cast to Uint format, getting this object returns zero.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Unsigned32 (0...4294967295)

apvPropertyCounter64Value (1.3.6.1.4.1.5655.4.1.13.3.1.6)

The value of the application property in Counter64 format.

If the property cannot be cast to Counter64 format, getting this object returns zero.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Counter64

trafficCountersTable (1.3.6.1.4.1.5655.4.1.14.1)

A list of information for each traffic counter.

| | |
|--------|----------------|
| Access | not-accessible |
|--------|----------------|

Syntax

Sequence of trafficCountersEntry

trafficCountersEntry (1.3.6.1.4.1.5655.4.1.14.1.1)

Entry containing information for a specified traffic counter.

| | |
|--------|----------------|
| Access | not-accessible |
|--------|----------------|

Index

{trafficCounterIndex}

Syntax

SEQUENCE *{trafficCounterIndex trafficCounterValue trafficCounterName trafficCounterType}*

trafficCounterIndex (1.3.6.1.4.1.5655.4.1.14.1.1.1)

An index value that uniquely identifies the counter.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

INTEGER *(1..255)*

trafficCounterValue (1.3.6.1.4.1.5655.4.1.14.1.1.2)

The 64 bit counter value.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

`Counter64`

trafficCounterName (1.3.6.1.4.1.5655.4.1.14.1.1.3)

The name of the counter.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

`DisplayString`

trafficCounterType (1.3.6.1.4.1.5655.4.1.14.1.1.4)

Defines whether the traffic counters counts by packets (3) or by bytes (2).

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

```

INTEGER {1 (other)- none of the following2 (bytes)- counts by bytes3 (packets)- counts
by packets
}

```

attackTypeTable (1.3.6.1.4.1.5655.4.1.15.1)

A list of information for defined attack types.

| | |
|--------|----------------|
| Access | not-accessible |
|--------|----------------|

Syntax

Sequence of AttackTypeEntry

attackTypeEntry (1.3.6.1.4.1.5655.4.1.15.1.1)

Entry containing information for a specified attack type.

| | |
|--------|----------------|
| Access | not-accessible |
|--------|----------------|

Index

{moduleIndex, attackTypeIndex}

Syntax**SEQUENCE**

{attackTypeIndexattackTypeNameattackTypeCurrentNumAttacksattackTypeTotalNumAttacksattackTypeTotalNumFlowsattackTypeTotalNumSeconds}

attackTypeIndex (1.3.6.1.4.1.5655.4.1.15.1.1.1)

An index value that uniquely identifies the attack type.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

INTEGER (1..255)

attackTypeName (1.3.6.1.4.1.5655.4.1.15.1.1.2)

The name of the attack type.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

DisplayString

attackTypeCurrentNumAttacks (1.3.6.1.4.1.5655.4.1.15.1.1.3)

The number of attacks currently detected of this type.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Unsigned32 (0...4294967295)

attackTypeTotalNumAttacks (1.3.6.1.4.1.5655.4.1.15.1.1.4)

The total number of attacks of this type detected since last clear.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Unsigned32 (0...4294967295)

attackTypeTotalNumFlows (1.3.6.1.4.1.5655.4.1.15.1.1.5)

The total number of flows in attacks of this type detected since last clear.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Counter64

attackTypeTotalNumSeconds (1.3.6.1.4.1.5655.4.1.15.1.1.6)

The total duration (in seconds) of attacks of this type detected since last clear.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Unsigned32 (0...4294967295)

vasServersTable (1.3.6.1.4.1.5655.4.1.16.1)

A list of information for the VAS servers.

| | |
|--------|----------------|
| Access | not-accessible |
|--------|----------------|

Syntax

Sequence of VasServerEntry

vasServerEntry (1.3.6.1.4.1.5655.4.1.16.1.1)

Entry containing information for a specified VAS server.

| | |
|--------|----------------|
| Access | not-accessible |
|--------|----------------|

Index

{vasServerIndex}

Syntax

SEQUENCE { *vasServerIndex* *vasServerId* *vasServerAdminStatus* *vasServerOperStatus* }

vasServerIndex (1.3.6.1.4.1.5655.4.1.16.1.1.1)

An index value that uniquely identifies the VAS server.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Unsigned32 (0...4294967295)

vasServerId (1.3.6.1.4.1.5655.4.1.16.1.1.2)

The VAS server ID number in the system.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

Unsigned32 (0...4294967295)

vasServerAdminStatus (1.3.6.1.4.1.5655.4.1.16.1.1.3)

The administrative status of the VAS server.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

INTEGER { **1** (*other*) **2** (*up*) **3** (*down*) }

vasServerOperStatus (1.3.6.1.4.1.5655.4.1.16.1.1.4)

The operational status of the VAS server.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

INTEGER {1 (other)2 (up)3 (down)}

mplsVpnSoftwareCountersTable (1.3.6.1.4.1.5655.4.1.17.1)

A list of information on various system software counters related to MPLS/VPN auto-learning.

| | |
|--------|----------------|
| Access | not-accessible |
|--------|----------------|

Syntax

Sequence of mplsVpnSoftwareCountersEntry

mplsVpnSoftwareCountersEntry (1.3.6.1.4.1.5655.4.1.17.1.1)

Entry containing information regarding MPLS/VPN auto-learning.

| | |
|--------|----------------|
| Access | not-accessible |
|--------|----------------|

Syntax

SEQUENCE {mplsVpnMaxHWMappingsmplsVpnCurrentHWMappings}

mplsVpnMaxHWMappings (1.3.6.1.4.1.5655.4.1.17.1.1.1)

The maximum number of hardware mappings permitted.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

INTEGER (1..1000000)

mplsVpnCurrentHWMappings (1.3.6.1.4.1.5655.4.1.17.1.1.2)

The current number of hardware mappings in the system.

| | |
|--------|----|
| Access | RO |
|--------|----|

Syntax

INTEGER (1..1000000)

Supported Standards

The SCE platform supports the SNMP related standards listed in the following table.

Table B-1 Supported SNMP Standards

| Document Name | Description |
|--|--|
| RFC 1155: Structure and Identification of Management Information for TCP/IP-based Internets | K. McCloghrie and M. T. Rose, (May 1990). Contains MIB object definitions. (Obsoletes RFC 1065) |
| RFC 1157: A Simple Network Management Protocol | J. D. Case, M. Fedor, M. L. Schoffstall, and C. Davin, (May 1990). Defines SNMP. (Obsoletes RFC 1098) |
| RFC 1212: Concise MIB Definitions | K. McCloghrie (March 1991). Defines a format for producing MIB modules |
| RFC 1213: Management Information Base Network Management of TCP/IP based internets: MIB-II | K. McCloghrie and M. T. Rose, eds., (March 1991). Defines MIB-II. (Obsoletes RFC 1158) |
| RFC 1215: Convention for Defining Traps for Use with the SNMP | M. T. Rose, ed. (March 1991). |
| RFC 1901: Introduction to Community-based SNMPv2 | SNMPv2 WG, J. Case, K. McCloghrie, M. T. Rose, S. Waldbusser, (January 1996). Defines "Community-based SNMPv2." (Experimental. Obsoletes RFC 1441) |
| RFC 1905: Protocol Operations for Version 2 of the Simple Network Management Protocol (SNMPv2) | Obsoletes: 1448 (January 1996) |
| RFC 1906: Transport Mappings for Version 2 of the Simple Network Management Protocol (SNMPv2) | Obsoletes: 1449 (January 1996) |
| RFC 2233: The Interfaces Group MIB using SMIV2. Obsoleted by RFC-2863: The Interfaces Group MIB | Extensions for the ifTable. |

