



**7450 ETHERNET SERVICE SWITCH
7750 SERVICE ROUTER
7950 EXTENSIBLE ROUTING SYSTEM**

**SYSTEM MANAGEMENT GUIDE
RELEASE 15.0.R1**

3HE 11979 AAAA TQZZA 01

Issue: 01

March 2017

Nokia is a registered trademark of Nokia Corporation. Other products and company names mentioned herein may be trademarks or tradenames of their respective owners.

The information presented is subject to change without notice. No responsibility is assumed for inaccuracies contained herein.

© 2017 Nokia.

Contains proprietary/trade secret information which is the property of Nokia and must not be made available to, or copied or used by anyone outside Nokia without its written authorization. Not to be used or disclosed except in accordance with applicable agreements.

Table of Contents

1	Getting Started	13
1.1	About This Guide.....	13
1.2	Router Configuration Process	14
2	Security	15
2.1	Authentication, Authorization, and Accounting	15
2.1.1	Authentication.....	16
2.1.1.1	Local Authentication	17
2.1.1.2	RADIUS Authentication	18
2.1.1.3	TACACS+ Authentication	22
2.1.1.4	LDAP Authentication	23
2.1.2	Authorization.....	30
2.1.2.1	Local Authorization	30
2.1.2.2	RADIUS Authorization	30
2.1.2.3	TACACS+ Authorization.....	31
2.1.3	Accounting.....	33
2.1.3.1	RADIUS Accounting	34
2.1.3.2	TACACS+ Accounting	34
2.2	Security Controls	34
2.2.1	When a Server Does Not Respond	35
2.2.2	Access Request Flow	36
2.3	Centralized CPU Protection.....	37
2.3.1	CPU Protection Extensions for ETH-CFM.....	40
2.3.2	ETH-CFM Ingress Squelching	43
2.4	Distributed CPU Protection (DCP).....	45
2.4.1	Applicability of Distributed CPU Protection.....	48
2.4.2	Log Events, Statistics, Status and SNMP support.....	49
2.4.3	DCP Policer Resource Management.....	49
2.4.4	Operational Guidelines and Tips	50
2.5	Classification-Based Priority for Extracted Protocol Traffic	52
2.6	Vendor-Specific Attributes (VSAs).....	53
2.7	Other Security Features	54
2.7.1	Secure Shell (SSH)	54
2.7.2	SSH PKI Authentication.....	56
2.7.2.1	Key Generation.....	56
2.7.3	Per Peer CPM Queuing.....	57
2.7.4	CPM Filters and Traffic Management.....	57
2.7.5	TTL Security for BGP and LDP	58
2.7.6	Exponential Login Backoff.....	59
2.7.7	User Lockout	60
2.7.8	CLI Login Scripts	60
2.7.9	802.1x Network Access Control	61
2.7.10	TCP Enhanced Authentication Option.....	61
2.7.10.1	Packet Formats	62
2.7.10.2	Keychain.....	63

2.7.11	gRPC Authentication	65
2.8	Configuration Notes	67
2.8.1	General	67
2.9	Configuring Security with CLI	69
2.10	Setting Up Security Attributes	70
2.10.1	Configuring Authentication	70
2.10.2	Configuring Authorization	71
2.10.3	Configuring Accounting	71
2.11	Security Configurations	72
2.12	Configuration Tasks	72
2.13	Security Configuration Procedures	74
2.13.1	Configuring Management Access Filters	74
2.13.2	Configuring IP CPM Filters Policy	75
2.13.3	Configuring MAC CPM Filters	76
2.13.4	Configuring IPv6 CPM Filters	77
2.13.5	Configuring CPM Queues	77
2.13.6	IPSec Certificates Parameters	79
2.13.7	Configuring Profiles	80
2.13.7.1	Parameters	80
2.13.7.2	Wildcards	82
2.13.7.3	CLI Session Resource Management	83
2.13.8	Configuring Users	85
2.13.9	Configuring Keychains	85
2.13.10	Copying and Overwriting Users and Profiles	86
2.13.10.1	User	86
2.13.10.2	Profile	88
2.14	RADIUS Configurations	90
2.14.1	Configuring RADIUS Authentication	90
2.14.2	Configuring RADIUS Authorization	91
2.14.3	Configuring RADIUS Accounting	91
2.15	Configuring 802.1x RADIUS Policies	92
2.16	TACACS+ Configurations	93
2.16.1	Enabling TACACS+ Authentication	93
2.16.2	Configuring TACACS+ Authorization	93
2.16.3	Configuring TACACS+ Accounting	94
2.16.4	Enabling SSH	95
2.17	LDAP Configurations	95
2.17.1	Configuring LDAP Authentication	95
2.17.2	Configuring Redundant Servers	97
2.17.3	Enabling SSH	98
2.18	Configuring Login Controls	98
2.19	Security Configuration Command Reference	99
2.19.1	Command Hierarchies	99
2.19.1.1	Security Commands	99
2.19.1.2	Login Control Commands	115
2.19.2	Command Descriptions	117
2.19.2.1	General Security Commands	117
2.19.2.2	LLDP Commands	122
2.19.2.3	Login, Telnet, SSH and FTP Commands	125

2.19.2.4	Management Access Filter Commands.....	134
2.19.2.5	Password Commands	151
2.19.2.6	Public Key Infrastructure (PKI) Commands.....	161
2.19.2.7	Profile Management Commands	177
2.19.2.8	User Management Commands.....	181
2.19.2.9	CLI Session Management Commands	191
2.19.2.10	RADIUS Client Commands	193
2.19.2.11	TACACS+ Client Commands	198
2.19.2.12	LDAP Client Commands.....	203
2.19.2.13	Generic 802.1x COMMANDS	208
2.19.2.14	Keychain Authentication	212
2.19.2.15	CLI Script Commands	218
2.19.2.16	CPM Filter Commands	220
2.19.2.17	CPM Queue Commands	238
2.19.2.18	TTL Security Commands.....	239
2.19.2.19	gRPC Commands.....	241
2.19.2.20	CPU Protection Commands	242
2.19.2.21	Distributed CPU Protection Commands	252
2.19.2.22	Extracted Protocol Traffic Priority Commands.....	260
2.20	Security Show, Clear, Debug, Tools, and Admin Command Reference	261
2.20.1	Command Hierarchies.....	261
2.20.1.1	Show Commands	261
2.20.1.2	Clear Commands.....	262
2.20.1.3	Debug Commands.....	263
2.20.1.4	Tools Commands	263
2.20.1.5	Admin Commands.....	263
2.20.2	Command Descriptions	264
2.20.2.1	Show Commands	264
2.20.2.2	Clear Commands.....	310
2.20.2.3	Debug Commands.....	313
2.20.2.4	Tools Commands	314
2.20.2.5	Admin Commands.....	316
3	SNMP	317
3.1	In This Chapter	317
3.2	SNMP Overview	318
3.2.1	SNMP Architecture	318
3.2.2	Management Information Base	319
3.2.3	SNMP Protocol Operations	319
3.2.4	SNMP Versions	319
3.2.5	Management Information Access Control	320
3.2.6	User-Based Security Model Community Strings	320
3.2.7	Views.....	321
3.2.8	Access Groups	321
3.2.9	Users	322
3.2.10	Per-VPRN Logs and SNMP Access.....	322
3.2.11	Per-SNMP Community Source IP Address Validation	323
3.3	Which SNMP Version to Use?.....	323

3.4	Configuration Notes.....	324
3.4.1	General.....	324
3.5	Configuring SNMP with CLI.....	327
3.6	SNMP Configuration Overview.....	327
3.6.1	Configuring SNMPv1 and SNMPv2c.....	327
3.6.2	Configuring SNMPv3.....	328
3.7	Basic SNMP Security Configuration.....	328
3.8	Configuring SNMP Components.....	329
3.8.1	Configuring a Community String.....	330
3.8.2	Configuring View Options.....	330
3.8.3	Configuring Access Options.....	331
3.8.4	Configuring USM Community Options.....	332
3.8.5	Configuring Other SNMP Parameters.....	333
3.9	SNMP Configuration Command Reference.....	335
3.9.1	Command Hierarchies.....	335
3.9.1.1	SNMP System Commands.....	335
3.9.1.2	SNMP Security Commands.....	335
3.9.2	Command Descriptions.....	336
3.9.2.1	SNMP System Commands.....	336
3.9.2.2	SNMP Security Commands.....	339
3.10	SNMP Show Command Reference.....	349
3.10.1	Command Hierarchies.....	349
3.10.1.1	Show Commands.....	349
3.10.2	Command Descriptions.....	349
3.10.2.1	Show Commands.....	349
4	NETCONF.....	371
4.1	In This Chapter.....	371
4.2	NETCONF Overview.....	371
4.3	NETCONF in SR OS.....	373
4.3.1	YANG Data Models.....	373
4.3.2	Transport and Sessions.....	375
4.3.3	Datstores and URLs.....	376
4.3.4	NETCONF Operations and Capabilities.....	378
4.3.4.1	<get>.....	379
4.3.4.2	<get-config>.....	381
4.3.4.3	<edit-config>.....	381
4.3.4.4	<copy-config> and <delete-config>.....	381
4.3.4.5	<lock>.....	382
4.3.4.6	<unlock>.....	383
4.3.4.7	<commit>.....	384
4.3.4.8	<discard-changes>.....	384
4.3.4.9	<validate>.....	384
4.3.5	Data Model, Datstore and Operation Combinations.....	385
4.3.6	General NETCONF Behavior.....	385
4.3.6.1	System-Provisioned Configuration (SPC) Objects.....	395
4.4	Establishing a NETCONF Session.....	398
4.5	XML Content Layer.....	400
4.5.1	<get> with XML Content Layer.....	400

4.5.2	<edit-config> with XML Content Layer	402
4.5.3	<get-config> with XML Content Layer	411
4.6	XML Content Layer Examples	418
4.7	CLI Content Layer	423
4.8	CLI Content Layer Examples	424
4.9	NETCONF Configuration Command Reference	429
4.9.1	Command Hierarchies	430
4.9.1.1	NETCONF System Commands	430
4.9.1.2	NETCONF Security Commands	430
4.9.2	Configuration Commands	431
4.9.2.1	NETCONF System Commands	431
4.9.2.2	NETCONF Security Commands	432
4.10	NETCONF Show Command Reference	435
4.10.1	Command Hierarchies	435
4.10.1.1	Show Commands	435
4.10.2	Command Descriptions	435
4.10.2.1	Show Commands	435
4.11	NETCONF Admin Command Reference	439
4.11.1	Command Hierarchies	439
4.11.1.1	Admin Commands	439
4.11.2	Command Descriptions	440
4.11.2.1	Admin Commands	440
5	Event and Accounting Logs	441
5.1	In This Chapter	441
5.2	Logging Overview	442
5.3	Log Destinations	443
5.3.1	Console	444
5.3.2	Session	444
5.3.3	Memory Logs	444
5.3.4	Log Files	445
5.3.5	SNMP Trap Group	447
5.3.6	Syslog	447
5.4	Event Logs	449
5.4.1	Event Sources	450
5.4.2	Event Control	451
5.4.3	Log Manager and Event Logs	452
5.4.4	Event Filter Policies	453
5.4.5	Event Log Entries	454
5.4.6	Simple Logger Event Throttling	456
5.4.7	Default System Log	456
5.4.8	Event Handling System	457
5.5	Customizing Syslog Messages Using Python	466
5.5.1	Python Engine for Syslog	466
5.5.1.1	Python Syslog APIs	467
5.5.1.2	Timestamp Format Manipulation	470
5.5.2	Python Processing Efficiency	472
5.5.3	Python Backpressure	472
5.5.4	Event Selection for Python Processing	473

5.5.5	Modifying a Log File	475
5.5.6	Deleting a Log File.....	475
5.5.7	Modifying a File ID.....	476
5.5.8	Modifying a Syslog ID.....	477
5.5.9	Modifying an SNMP Trap Group	478
5.5.10	Deleting an SNMP Trap Group.....	478
5.5.11	Modifying a Log Filter	479
5.5.12	Modifying Event Control Parameters.....	480
5.5.13	Returning to the Default Event Control Configuration	481
5.6	Accounting Logs.....	482
5.6.1	Accounting Records	482
5.6.2	Accounting Files	506
5.6.3	Design Considerations	506
5.6.4	Reporting and Time-Based Accounting.....	507
5.6.5	Overhead Reduction in Accounting: Custom Record	507
5.6.5.1	User Configurable Records	507
5.6.5.2	Changed Statistics Only	508
5.6.5.3	Configurable Accounting Records.....	508
5.6.5.4	Significant Change Only Reporting	509
5.6.6	Immediate Completion of Records	509
5.6.6.1	Record Completion for XML Accounting	509
5.6.7	AA Accounting per Forwarding Class.....	510
5.7	Configuration Notes.....	510
5.8	Configuring Logging with CLI	511
5.9	Log Configuration Overview	511
5.9.1	Log Types.....	511
5.10	Basic Event Log Configuration	512
5.11	Common Configuration Tasks	512
5.11.1	Configuring an Event Log	513
5.11.2	Configuring a File ID.....	513
5.11.3	Configuring an Accounting Policy.....	514
5.11.4	Configuring Event Control	515
5.11.5	Configuring a Log Filter	515
5.11.6	Configuring an SNMP Trap Group	516
5.11.6.1	Setting the Replay Parameter	517
5.11.6.2	Shutdown In-Band Port	518
5.11.6.3	No Shutdown Port	520
5.11.7	Configuring a Syslog Target.....	521
5.11.7.1	Configuring an Accounting Custom Record	522
5.12	Log Configuration Command Reference.....	525
5.12.1	Command Hierarchies.....	525
5.12.1.1	Log Configuration Commands	525
5.12.1.2	Accounting Policy Commands.....	526
5.12.1.3	Custom Record Commands	526
5.12.1.4	File ID Commands.....	529
5.12.1.5	Event Filter Commands.....	529
5.12.1.6	Event Handling System (EHS) Commands	530
5.12.1.7	Event Trigger Commands.....	530
5.12.1.8	Log ID Commands.....	531

5.12.1.9	SNMP Trap Group Commands	531
5.12.1.10	Syslog Commands	532
5.12.2	Command Descriptions	533
5.12.2.1	Generic Commands.....	533
5.12.2.2	Log Configuration Commands	535
5.12.2.3	File ID Commands.....	539
5.12.2.4	Log Filter Commands	543
5.12.2.5	Log Filter Entry Commands.....	544
5.12.2.6	Log Filter Entry Match Commands	545
5.12.2.7	Event Handling System (EHS) Commands	550
5.12.2.8	Event Trigger Commands.....	552
5.12.2.9	Syslog Commands	554
5.12.2.10	SNMP Trap Groups.....	560
5.12.2.11	Accounting Policy Commands	570
5.13	Log Command Reference	595
5.13.1	Command Hierarchies.....	595
5.13.1.1	Show Commands	596
5.13.1.2	Clear Command	596
5.13.2	Command Descriptions	597
5.13.2.1	Show Commands	597
5.13.2.2	Clear Commands.....	627
6	sFlow	629
6.1	In This Chapter	629
6.2	sFlow Overview	629
6.3	sFlow Features.....	630
6.3.1	sFlow Counter Polling Architecture	630
6.3.2	sFlow Support on Logical Ethernet Ports	631
6.3.3	sFlow SAP Counter Map.....	632
6.3.4	sFlow Record Formats	632
6.4	sFlow Command Reference	637
6.4.1	Command Hierarchies.....	637
6.4.1.1	System Commands	637
6.4.1.2	Show Commands	637
6.5	sFlow Configuration Command Descriptions	639
6.5.1	Command Descriptions	639
6.5.1.1	System Commands	639
6.6	sFlow Show Command Descriptions.....	643
6.6.1	Command Descriptions	643
6.6.1.1	Show Commands	643
7	Telemetry	647
7.1	In This Chapter	647
7.2	Telemetry Overview.....	647
7.3	About Telemetry	648
7.3.1	gRPC in Telemetry	648
7.3.2	Operations Layer.....	651
7.3.3	Schema Paths	654
7.4	Telemetry Examples.....	655

7.5	gRPC Command Reference.....	663
7.5.1	Command Hierarchies.....	663
7.5.1.1	System Commands.....	663
7.5.1.2	QoS Commands.....	663
7.6	Telemetry Configuration Command Descriptions.....	665
7.6.1	Command Descriptions.....	665
7.6.1.1	System Commands.....	665
7.6.1.2	QoS Commands.....	666
7.7	gRPC Show, Admin Command Reference.....	667
7.7.1	Command Hierarchies.....	667
7.7.1.1	Show Commands.....	667
7.7.1.2	Admin Commands.....	667
7.7.2	Command Descriptions.....	667
7.7.2.1	Show Commands.....	668
7.7.2.2	Admin Commands.....	669
8	TLS.....	671
8.1	In This Chapter.....	671
8.2	TLS Overview.....	671
8.3	TLS Server Interaction with Applications.....	672
8.3.1	TLS Application Support.....	672
8.4	TLS Handshake.....	673
8.5	TLS Client Certificate.....	674
8.6	TLS Symmetric Key Rollover.....	675
8.7	Supported TLS Ciphers.....	675
8.8	SR OS Certificate Management.....	676
8.8.1	Certificate Profile.....	676
8.8.2	TLS Server Authentication of the Client Certificate CN Field.....	677
8.8.3	CN Regexp Format.....	677
8.9	Operational Guidelines.....	677
8.9.1	Server Authentication Behavior.....	677
8.9.2	Client TLS Profile and Trust Anchor Behavior and Scale.....	678
8.10	LDAP Redundancy and TLS.....	679
8.11	Basic TLS Configuration.....	681
8.12	Common Configuration Tasks.....	682
8.12.1	Configuring a Server TLS Profile.....	682
8.12.2	Configuring a Client TLS Profile.....	682
8.12.3	Configuring a TLS Client or TLS Server Certificate.....	683
8.12.4	Configuring a TLS Trust Anchor.....	683
8.13	TLS Command Reference.....	685
8.13.1	Command Hierarchies.....	685
8.13.1.1	Security TLS Commands.....	685
8.13.1.2	LDAP TLS Profile Commands.....	686
8.13.1.3	Admin Commands.....	686
8.13.2	Command Descriptions.....	687
8.13.2.1	Security TLS Commands.....	687
8.13.2.2	LDAP TLS Profile Commands.....	694
8.13.2.3	Admin Commands.....	695
8.14	TLS Show Command Reference.....	697

8.14.1	Command Hierarchies	697
8.14.1.1	Show Commands	697
8.14.2	Command Descriptions	698
8.14.2.1	Show Commands	698
9	Facility Alarms	701
9.1	In This Chapter	701
9.2	Facility Alarms Overview	701
9.3	Facility Alarms vs. Log Events	702
9.4	Facility Alarm Severities and Alarm LED Behavior.....	703
9.5	Facility Alarm Hierarchy.....	704
9.6	Facility Alarm List	705
9.7	Configuring Logging with CLI	717
9.8	Basic Facility Alarm Configuration.....	717
9.9	Common Configuration Tasks.....	717
9.9.1	Configuring the Maximum Number of Alarms To Clear.....	717
9.10	Facility Alarms Configuration Command Reference.....	719
9.10.1	Command Hierarchies.....	719
9.10.1.1	Facility Alarm Configuration Commands	719
9.10.2	Command Descriptions	719
9.10.2.1	Generic Commands.....	719
9.11	Facility Alarms Show Command Reference	721
9.11.1	Command Hierarchies.....	721
9.11.1.1	Show Commands	721
9.11.2	Command Descriptions	721
9.11.2.1	Show Commands	721
10	Standards and Protocol Support	723

1 Getting Started

1.1 About This Guide

This guide describes system concepts and provides configuration explanations and examples to configure SR-OS boot option file (BOF), file system and system management functions.

This guide is organized into functional chapters and provides concepts and descriptions of the implementation flow, as well as Command Line Interface (CLI) syntax and command usage.

The topics and commands described in this document apply to the:

- 7450 ESS
- 7750 SR
- 7950 XRS

[Table 1](#) lists the available chassis types for each SR OS router.

Table 1 Supported SR OS Router Chassis Types

7450 ESS	7750 SR	7950 XRS
<ul style="list-style-type: none"> • 7450 ESS-7/12 running in standard mode (not mixed-mode) 	<ul style="list-style-type: none"> • 7450 ESS-7/12 running in mixed-mode (not standard mode) • 7750 SR-a4/a8 • 7750 SR-c4/c12 • 7750 SR-1e/2e/3e • 7750 SR-7/12 • 7750 SR-12e 	<ul style="list-style-type: none"> • 7950 XRS-16c • 7950 XRS-20/40

For a list of unsupported features by platform and chassis, refer to the *SR OS R15.0.Rx Software Release Notes*, part number 3HE 12060 000x TQZZA.

Command outputs shown in this guide are examples only; actual displays may differ depending on supported functionality and user configuration.



Note: This guide generically covers Release 15.0 content and may contain some content that will be released in later maintenance loads. Please refer to the *SR OS R15.0.Rx Software Release Notes*, part number 3HE 12060 000x TQZZA, for information on features supported in each load of the Release 15.0 software.

1.2 Router Configuration Process

[Table 2](#) lists the tasks necessary to configure system security and access functions and logging features on the 7450 ESS, 7750 SR, and 7950 XRS platforms. Each chapter in this book is presented in an overall logical configuration flow. Each section describes a software area and provides CLI syntax and command usage to configure parameters for a functional area.

Table 2 Configuration Process

Area	Task	Chapter	Supported platforms
System security	Configure system security parameters, such as authentication, authorization, and accounting.	Getting Started	All
Network management	Configure SNMP elements.	SNMP	All
Secure network management	Configure NETCONF elements.	NETCONF	All
Operational functions	Configure event and accounting logs.	Event and Accounting Logs	All
Counter management	Configure sFlow elements.	sFlow	7750 SR and 7950 XRS
Reference	List of IEEE, IETF, and other proprietary entities.	Standards and Protocol Support	All



Note: All features are supported on all SR OS platforms (7750 SR, 7450 ESS, and 7950 XRS) unless indicated otherwise.

2 Security

2.1 Authentication, Authorization, and Accounting

This chapter describes authentication, authorization, and accounting (AAA) used to monitor and control network access on routers. Network security is based on a multi-step process. The first step, authentication, validates a user's name and password. The second step is authorization, which allows the user to access and execute commands at various command levels based on profiles assigned to the user.

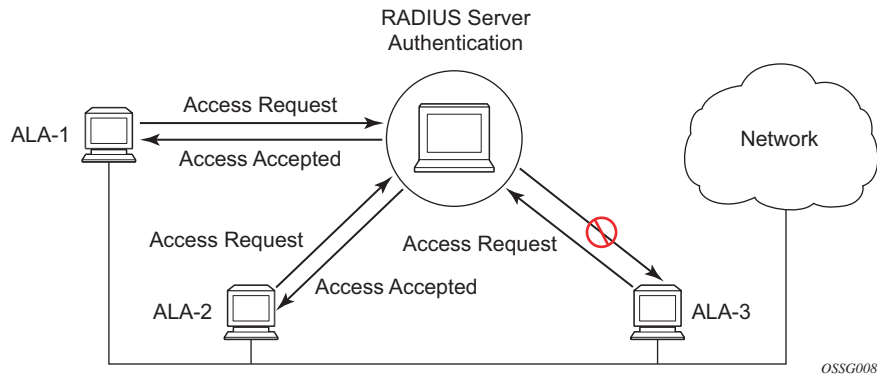
Another step, accounting, keeps track of the activity of a user who has accessed the network. The type of accounting information recorded can include a history of the commands executed, the amount of time spent in the session, the services accessed, and the data transfer size during the session. The accounting data can then be used to analyze trends, and also for billing and auditing purposes.

You can configure routers to use local, Remote Authentication Dial In User Service (RADIUS), or Terminal Access Controller Access Control System Plus (TACACS+) security to validate users who attempt to access the router by console, Telnet, or FTP. You can select the authentication order which determines the authentication method to try first, second, and third.

The router supports the following security features:

- RADIUS can be used for authentication, authorization, and accounting
- TACACS+ can be used for authentication, authorization, and accounting
- Local security can be implemented for authentication and authorization

[Figure 1](#) depicts end user access-requests sent to a RADIUS server. After validating the user names and passwords, the RADIUS server returns an access-accept message to the users on ALA-1 and ALA-2. The user name and password from ALA-3 could not be authenticated, thus access was denied.

Figure 1 RADIUS Requests and Responses

2.1.1 Authentication

Authentication validates a user name and password combination when a user attempts to log in.

When a user attempts to log in through the console, Telnet, SSH, SCP, or FTP, the client sends an access request to a RADIUS, TACACS+, or local database.

Transactions between the client and a RADIUS server are authenticated through the use of a shared secret. The secret is never transmitted over the network. User passwords are sent encrypted between the client and RADIUS server which prevents someone snooping on an insecure network to learn password information.

If the RADIUS server does not respond within a specified time, the router issues the access request to the next configured servers. Each RADIUS server must be configured identically to guarantee consistent results.

If any RADIUS server rejects the authentication request, it sends an access reject message to the router. In this case, no access request is issued to any other RADIUS servers. However, if other authentication methods such as TACACS+ and/or local are configured, then these methods are attempted. If no other authentication methods are configured, or all methods reject the authentication request, then access is denied.

For the RADIUS server selection, round-robin is used if multiple RADIUS servers are configured. Although, if the first alive server in the list cannot find a user-name, the router does not re-query the next server in the RADIUS server list and denies the access request. It may get authenticated on the next login attempt if the next selected RADIUS server has the appropriate user-name. It is recommended that the same user databases are maintained for RADIUS servers in order to avoid inconsistent behavior.

The user login is successful when the RADIUS server accepts the authentication request and responds to the router with an access accept message.

Implementing authentication without authorization for the routers does not require the configuration of VSAs (Vendor Specific Attributes) on the RADIUS server. However, users, user access permissions, and command authorization profiles must be configured on each router.

Any combination of these authentication methods can be configured to control network access from a router:

- [Local Authentication](#)
- [RADIUS Authentication](#)
- [TACACS+ Authentication](#)
- [LDAP Authentication](#)

2.1.1.1 Local Authentication

Local authentication uses user names and passwords to authenticate login attempts. The user names and passwords are local to each router not to user profiles.

By default, local authentication is enabled. When one or more of the other security methods are enabled, local authentication is disabled. Local authentication is restored when the other authentication methods are disabled. Local authentication is attempted if the other authentication methods fail and local is included in the authentication order password parameters.

Locally, user names and password management information can be configured. This is referred to as local authentication. Remote security servers such as RADIUS or TACACS+, are not enabled.

2.1.1.2 RADIUS Authentication

Remote Authentication Dial-In User Service (RADIUS) is a client/server security protocol and software that enables remote access servers to communicate with a central server to authenticate dial-in users and authorize access to the requested system or service.

RADIUS allows you to maintain user profiles in a shared central database and provides better security, allowing a company to set up a policy that can be applied at a single administered network point.

2.1.1.2.1 RADIUS Server Selection

The RADIUS server selection algorithm is used by different applications:

- RADIUS operator management
- RADIUS authentication for Enhanced Subscriber Management
- RADIUS accounting for Enhanced Subscriber Management
- RADIUS PE-discovery

In all these applications, up to 5 RADIUS servers pools (per RADIUS policy, if used) can be configured.

The RADIUS server selection algorithm can work in 2 modes, either Direct mode or Round-robin mode.

Direct Mode

The first server is used as the primary server. If this server is unreachable, the next server, based on the server index, of the server pool is used. This continues until either all servers in the pool have been tried or an answer is received.

If a server is unreachable, it will not be used again by the RADIUS application for the next 30 seconds to allow the server to recover from its unreachable state. After 30 seconds the unreachable server is available again for the RADIUS application. If in these 30 seconds the RADIUS application receives a valid response for a previously sent RADIUS packet on that unreachable server, the server will be available for the RADIUS application again, immediately after reception of that response.

Round-Robin Mode

The RADIUS application sends the next RADIUS packet to the next server in the server pool. The same server non-reachability behavior is valid as in the Direct mode.

Server Reachability Detection

A server is reachable, when the operational state UP, when a valid response is received within a timeout period which is configurable by the retry parameter on the RADIUS policy level.

A server is treated as not-reachable, when the operational state down, when the following occurs:

- A timeout — If a number of consecutive timeouts are encountered for a specific server. This number is configurable by the retry parameter on RADIUS policy level.
- A send failed — If a packet cannot be sent to the RADIUS server because the forwarding path towards the RADIUS server is broken (for example, the route is not available, the interface is shutdown, etc.), then, no retry mechanism is invoked and immediately, the next server in line is used.

A server that is down can only be used again by the RADIUS algorithm after 30 seconds, unless, during these 30 seconds a valid RADIUS reply is received for that server. Then, the server is immediately marked UP again.

The operational state of a server can also be “unknown” if the RADIUS application is not aware of the state of the RADIUS server (for example, if the server was previously down but no requests had been sent to the server, thus, it is not certain yet whether the server is actually reachable).

Application Specific Behavior

Operator Management

The server access mode is fixed to Round-Robin (Direct cannot be configured for operator management). A health-check function is available for operator management, which can optionally be disabled. The health-check polls the server once every 10 seconds with an improbable user name. If the server does not respond to this health-check, it will be marked down.

If the first server in the list cannot find a user, the next server in the RADIUS server list is not queried and access is denied. If multiple RADIUS servers are used, it is assumed they all have the same user database.

RADIUS Authentication

If the first server in the list cannot find a user, the next server in the RADIUS server list is not queried and access is denied. If multiple RADIUS servers are used, it is assumed they all have the same user database.

RADIUS Challenge/Response Interactive Authentication

Challenge-response interactive authentication is used for key authentication where the Radius server is asking for the valid response to a displayed challenge. The challenge packet includes a challenge to be displayed to the user, such as a unique generated numeric value unlikely ever to be repeated. Typically this is obtained from an external server that knows what type of authenticator is in the possession of the authorized user and can therefore choose a random or non-repeating pseudorandom number of appropriate length.

The user then enters the challenge into his device (or software) and it calculates a response, which the user enters into the client which forwards it to the RADIUS server within an access request. If the response matches the expected response, the RADIUS server allows the user access, otherwise it rejects the response.

RADIUS challenge/response mode is enabled using the CLI interactive-authentication command in the config>system>security>radius context. RADIUS interactive authentication is disabled by default. The option needs to be enabled using CLI.

Enabling interactive authentication under CLI does not mean that the system uses RADIUS challenge/response mode by default. The configured password authentication-order parameter is used. If the authentication-order parameter is local RADIUS, the system will first attempt to login the user using local authentication. If this fails, the system will revert to RADIUS and challenge/response mode. The authentication-order will precede the RADIUS interactive-authentication mode.

Even if the authentication-order is RADIUS local, the standard password prompt is always displayed. The user enters a username and password at this prompt. If RADIUS interactive-authentication is enabled the password does not have to be the correct password since authentication is accomplished using the RADIUS challenge/response method. The user can enter any password. The username and password are sent to the RADIUS server, which responds with a challenge request that is transmitted back to the node by the RADIUS server. Once the user enters the challenge response, the response is authenticated by the RADIUS server to allow node access to the user.

For example, if the system is configured with system security authentication-order set to local RADIUS, at the login prompt the user can enter the username "admin" and the corresponding password. If the password for local authentication does not match, the system falls into RADIUS authentication mode. The system checks the interactive-authentication configuration and if it is enabled it enters into challenge/response mode. It sends the username and password to the RADIUS server, and the server sends the challenge request back to the node and to the user where it appears as a challenge prompt on screen. A challenge received from the RADIUS server typically contains a string and a hardware token that can be used to generate a password on the users' local personal token generator. For example, the RADIUS server might send the challenge prompt "Enter response for challenge 12345:" to the SR OS. The string "12345" can be entered in the local token generator which generates the appropriate challenge response for the entered string. This challenge response can then be entered on the SR OS prompt for authorization.

Once the user enters the correct challenge response it is authenticated using the RADIUS server. The server authenticates the user and the user gains access to the node.

If session timeout and Idle timeout values are configured on the RADIUS server, these are used to govern the length of time before the SR OS cancels the challenge prompt. If the user is idle longer than the received idle-timeout (seconds) from the RADIUS server, and/or if the user does not press ENTER before the received session-timeout (seconds).



Note: For SSH only the session-timeout value is used. The SSH stack cannot track character input into the login prompt until the enter key is pressed.

If the idle/session attribute is not available or if the value is set to a very large number, the SR OS uses the smallest value set in "configure system login-control idle-timeout" and the idle/session timeout attribute value to terminate the prompt. If the "login-control idle-timeout" is set to 0 (equivalent to infinite), the maximum idle-timeout (24-hours) is used for the calculation.

The SR OS displays the log-in attempts/failure per user in the “show system security user user-name” screen. If the RADIUS rejects a challenge response, it counts as a failed login attempt and a new prompt is displayed. The number of failed attempts is limited by the value set for “configure system security password attempt.” An incorrect challenge response results in a failure count against the password attempts.

RADIUS Accounting

RADIUS accounting can be used for two purposes:

- CLI command accounting
- Enhanced Subscriber Management subscriber host accounting

The RADIUS accounting application will try to send all the accounting records of a subscriber host to the same RADIUS server. If that server is down, then the records are sent to the next server, and from that moment on, the RADIUS application uses that server as the destination for accounting records for that subscriber host. Enhanced Subscriber Management applies to the 7750 SR platform.

RADIUS PE-Discovery

If the first server in the list cannot find a user, the next server in the RADIUS server list is not queried and access is denied. If multiple RADIUS servers are used, it is assumed they all have the same user database.

The RADIUS PE-discovery application makes use of a 10 second time period instead of the generic 30 seconds and uses a fixed consecutive timeout value of 2 (see [Server Reachability Detection](#)).

As long as the Session-Timeout (attribute in the RADIUS user file) is specified, it is used for the polling interval. Otherwise, the configured polling interval will be used (60 seconds by default).

2.1.1.3 TACACS+ Authentication

Terminal Access Controller Access Control System, commonly referred to as TACACS is an authentication protocol that allows a remote access server to forward a user's logon password to an authentication server to determine whether access can be allowed to a given system. TACACS is an encryption protocol and therefore less secure than the later Terminal Access Controller Access Control System Plus (TACACS+) and RADIUS protocols.


TACACS+ and RADIUS have largely replaced earlier protocols in the newer or recently updated networks. TACACS+ uses Transmission Control Protocol (TCP) and RADIUS uses the User Datagram Protocol (UDP). TACACS+ is popular as TCP is thought to be a more reliable protocol. RADIUS combines authentication and authorization. TACACS+ separates these operations.

2.1.1.4 LDAP Authentication

Lightweight Directory Access Protocol (LDAP) can provide authentication, authorization, and accounting (AAA) functionality using in-band-management, and can allow users to access the full virtualized data center and networking devices. SR OS currently supports LDAP provision of a centralized authentication method with public key management. The authentication method is based on SSH public keys or keyboard authentication (username, password).

Administrators can access networking devices with one private key; public keys are usually saved locally on the SSH server. Proper key management is not feasible with locally-saved public keys on network devices or on virtual machines, as this would result in hundreds of public keys distributed on all devices. LDAPv3 provides a centralized key management system that allows for secure creation and distribution of public keys in the network. Public keys can be remotely saved on the LDAP server, which makes key management much easier, as shown in [Figure 2](#).

Figure 2 Key Management



1012

The administrator starts an SSH session through an SSH client using their private key. The SSH client for the authentication method sends a signature created with the user's private key to the router. The router authenticates the signature using the user's public key and gives access to the user. To access the public key, the router looks up the public key stored on the LDAP server instead of a locally-saved key stored on the router. Communication between the router and the LDAP server should be secured with LDAP over SSL/STL (LDAPS). After successful authentication, LDAP returns a set of public keys that can be used by the router to verify the signature.

LDAP is integrated into the SR OS as an AAA protocol alongside existing AAA protocols, such as RADIUS and TACACS+. The AAA framework provides tools and mechanisms (such as method lists, server groups, and generic attribute lists) that enable an abstract and uniform interface to AAA clients, irrespective of the actual protocol used for communication with the AAA server.

The authentication functions are:

- public key authentication—the client tries to SSH to the SR OS using public keys
Public keys can be stored locally or on the LDAP server and retrieved as needed to authenticate the user.
- password authentication—keyboard interactive
The LDAP server can be used for user authentication using keyboard interactive, as with simple user name and password authentication.

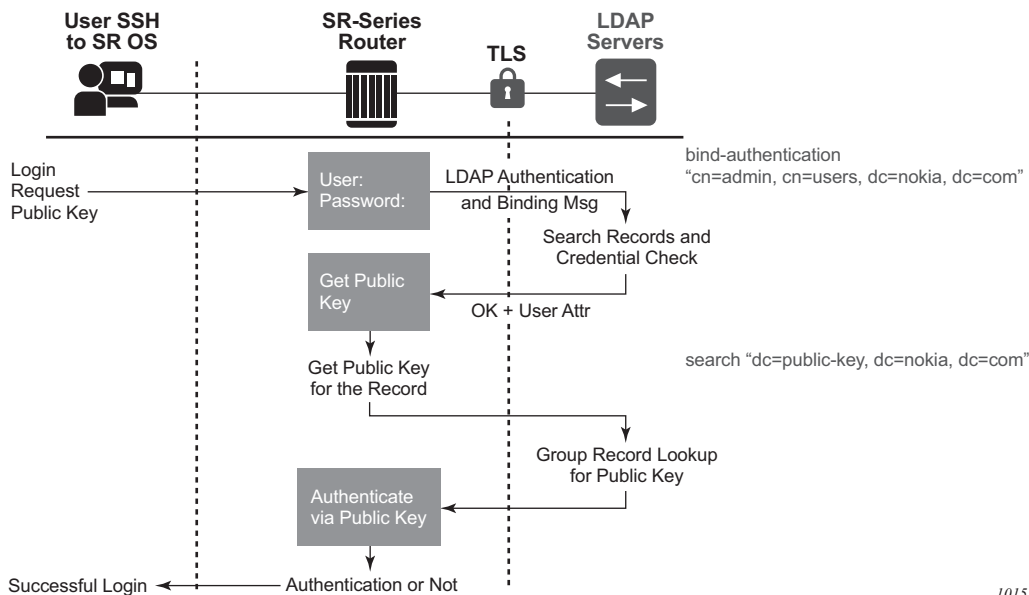
2.1.1.4.1 LDAP Authentication Process

A client starts an LDAP session by connecting to an LDAP server, called a Directory System Agent (DSA), which—by default—are on TCP port 389 and UDP port 636 for LDAP. The SR OS then sends an operation request to the server, and the server sends responses in return, as shown in [Figure 3](#). With some exceptions, the client does not need to wait for a response before sending the next request, and the server may send the responses in any order. All information is transmitted using Basic Encoding Rules (BER).

In the SR OS, the client can request the following operations:

- StartTLS
Use the LDAPv3 Transport Layer Security (TLS) extension for a secure connection.
- Bind
Authenticate and specify the LDAP protocol version.
- Search
Search for and retrieve directory entries.
- Unbind
Close the connection (not the inverse of Bind).

Figure 3 LDAP Server and SR OS Interaction for Retrieving the Public Key



The connection between the router as the LDAP client and the LDAP server should be encrypted using TLS, as all credentials between the router and LDAP are transmitted in clear text.

2.1.1.4.2 Authentication Order

The SR OS supports local and LDAP public key storage, the order of which is configured using the **config>system>security>password>authentication-order** command.



Note: The SR OS sends available authentication methods to the client and supports public key and password authentication. If the client is configured using **public-key-authentication** then it will use the public key authentication method.

If the client chooses the public key and LDAP is first in authentication order, then the SR OS will try to authenticate using public key retrieval from the LDAP server. If the public key retrieval from LDAP server fails and **exit-on-reject** was not configured, the SR OS will try the next method (**local**) in authentication order for the public key. If the next method also fails, a user authentication fail message will be sent to the client.

If the public key retrieval from the LDAP server fails and **exit-on-reject** is configured, the SR OS will not try the next method in the authentication order. A user authentication fail message will be sent to the client. At this point, the client can be configured to only use public key authentication, or use both public key authentication followed by password authentication. If the client is configured to use password authentication, it will go through the authentication order again, (for example, it will try all the configured methods in the configured **authentication-order**) as long as **exit-on-reject** is not configured.

Authentication Order Public Key Detail

There are two keys for public key authentication: a private key stored on the client and a public key stored on the server (local) or AAA server (ldap). The client uses the private key to create a signature, which only the public key can authenticate. If the signature is authenticated using the public key, then the user is also authenticated and is granted access. SR OS can locally store, using CLI, as many as 32 RSA keys and 32 ECDHA keys for a single user. In total, the SR OS can load a maximum of 128 public keys in a single authentication attempt.



Note: The client creates a signature using a single private key, but this signature can be authenticated on the SR OS with maximum of 128 public keys in a single try. If all these public keys fail to authenticate, then a failure message will be sent to the client and the number of failed attempts will be incremented.

If the client has another private key, it can create a new signature with this new private key and attempt the authentication one more time, or switch to password authentication.

The following steps outline the procedure where the client attempts to authenticate using a public key and the authentication order is configured as **ldap**, then **local**.



Note: With each increment of failed attempts, the SR OS also checks the limit for lock-out. If the limit is reached, the user is locked out.

1. The SSH client opens a session and tries to authenticate the user with private-key-1 (creating signature-1 from private-key-1).
2. The SR OS checks the authentication order.
3. The SR OS loads public keys for the user, as follows.
 - a. If **exit-on-reject** is not configured, the SR OS loads all public keys from the LDAP server and all public keys from the locally-saved location.

- b. If **exit-on-reject** is configured, the SR OS only loads all public keys from the LDAP server and not from the locally-saved location.
4. The SR OS compares received client signature-1 with signature calculated from loaded public keys and attempts to find a match.
 - a. If a match is found, the user is authenticated. The procedure ends.
 - b. If no match is found, authentication fails and the SSH client is informed. The LDAP server waits for the SSH client's reaction.
5. The SSH client reacts in one of several ways.
 - a. The connection is closed.
 - b. The password authentication method is continued. In this case, on the SR OS, the number of failed authentication attempts is not incremented.
 - c. The next public key is continued, as follows.
 - i. If it is not 21st received public key, return to step 3.
 - ii. If it is the 21st received public key, the number of failed authentication attempts is incremented and the connection is closed.

2.1.1.4.3 LDAP Authentication via Password

In addition to public key authentication, the SR OS supports password (keyboard) authentication using the LDAP server.



Note: TLS provides the encryption for password authentication.

In the following example, the client attempts to authenticate using a password and only **ldap** is configured in the authentication order.

1. The client uses telnet or SSH to reach the SR OS.
2. The SR OS retrieves the user name and password (in plain text).
3. The SR OS performs a bind operation to the LDAP server using the **config>system>security>ldap>server>blind-operation** command to set the *root-dn* and *password* variables.
4. The SR OS performs a search operation for the username on LDAP server.
 - a. If the user name is found, LDAP sends *user_distinguished_name* to the router.
 - b. If the user name is not found, the authentication fails. The attempt and failed attempt counters will be incremented.

5. The SR OS performs a bind operation to LDAP with `user_distinguished_name` and the password from step 2.
6. The LDAP server checks the password.
 - a. If the password is correct, the bind operation succeeds. The failed attempt and successful attempt counters are incremented.
 - b. If the password is incorrect, bind is unsuccessful and authentication fails. The attempt and failed attempt counters are incremented.
7. The SR OS sends a message to unbind from the LDAP server.

2.1.1.4.4 Timeout and Retry Configuration for the LDAP Server

The **retry** value is the maximum number of connection attempts that the SR OS can make to reach the current LDAP server before attempting the next server. For example, if the value is set to the default of 3, the SR OS will try to establish the connection to current server three times before attempting to establish a connection to the next server.

The **timeout** value is the number of seconds that the SR OS will wait for a response from the server with which it is attempting to establish a connection. If the server does not reply within the specified timeout value, the SR OS increments the **retry** counter by one. The SR OS attempts to establish the connection to the current server up to the configured **retry** value before moving to the next configured server.

2.1.1.4.5 TLS Behavior and LDAP

RFC 4511 section 4.14.1 states, “A client requests TLS establishment by transmitting a StartTLS request message to the server” and “The client MUST NOT send any LDAP PDUs at this LDAP message layer following this request until it receives a StartTLS Extended response”. As such, if an LDAP has a TLS profile configured and the TLS is in an operationally down state, no LDAP packets will be transmitted if TLS negotiation has not been completed, including when the TLS profile is shut down.

2.1.1.4.6 LDAP Health Check

The health check for LDAP is configured under **config>system>security>password**.

The **health-check** function, which can be disabled, is available for operator management. The health check polls the server at a specified interval (the default is 30 seconds). The SR OS health check attempts to establish a TCP connection to the LDAP server. The TCP connection is closed by an LDAP unbind message.

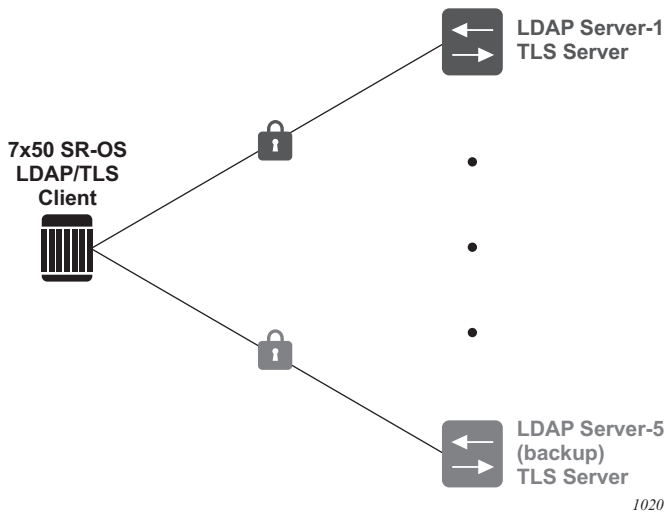
2.1.1.4.7 LDAP Redundancy and TLS

LDAP supports up to five redundant (backup) servers. Depending on the configuration of **timeout** and **retry** values, if an LDAP server is found to be out of service or operationally down, the SR OS will switch to the redundant servers. The SR OS will try the next LDAP server in the server list by choosing the next largest configured server index.

LDAP servers can use the same TLS profile or can have their own TLS profile. Each TLS profile can have a different configuration of **trust-anchor**, **cipher-list** and **cert-profile**. For security reasons, the LDAP server could be in different geographical areas and, as such, each will be assigned its own server certificate and trust anchor. The TLS profile design allows users to mix and match all components.

Redundant LDAP servers are shown in [Figure 4](#).

Figure 4 LDAP and TLS Redundancy



2.1.2 Authorization

The SR OS supports local, RADIUS, and TACACS+ authorization to control the actions of specific users. Any combination of these authorization methods can be configured to control actions of specific users:

- [Local Authorization](#)
- [RADIUS Authorization](#)
- [TACACS+ Authorization](#)

Local authorization and RADIUS authorization operate by applying a profile based on user name and password configurations once network access is granted. The profiles are configured locally as well as VSAs on the RADIUS server. See [Vendor-Specific Attributes \(VSAs\)](#).

2.1.2.1 Local Authorization

Local authorization uses user profiles and user access information after a user is authenticated. The profiles and user access information specifies the actions the user can and cannot perform.

By default, local authorization is enabled. Local authorization is disabled only when a different remote authorization method is configured, such as TACACS+ or RADIUS authorization.

You must configure profile and user access information locally.

2.1.2.2 RADIUS Authorization

RADIUS authorization grants or denies access permissions for a router. Permissions include the use of FTP, Telnet, SSH (SCP), and console access. When granting Telnet, SSH (SCP) and console access to the router, authorization can be used to limit what CLI commands the user is allowed to issue and which file systems the user is allowed or denied access.

Once a user has been authenticated using RADIUS (or another method), the router can be configured to perform authorization. The RADIUS server can be used to:

- Download the user profile to the router
- Send the profile name that the node should apply to the router.

Profiles consist of a suite of commands that the user is allowed or not allowed to execute. When a user issues a command, the authorization server looks at the command and the user information and compares it with the commands in the profile. If the user is authorized to issue the command, the command is executed. If the user is not authorized to issue the command, then the command is not executed.

Profiles must be created on each router and should be identical for consistent results. If the profile is not present, then access is denied.

Table 3 displays the following scenarios:

- Remote (RADIUS) authorization cannot be performed if authentication is done locally (on the router).
- The reverse scenario is supported if RADIUS authentication is successful and no authorization is configured for the user on the RADIUS server, then local (router) authorization is attempted, if configured in the authorization order.

When authorization is configured and profiles are downloaded to the router from the RADIUS server, the profiles are considered temporary configurations and are not saved when the user session terminates.

Table 3 Supported Authorization Configurations

	Router	RADIUS Supplied Profile
Router configured user	Supported	Not Supported
RADIUS server configured user	Supported	Supported
TACACS+ server configured user	Supported	Not Supported

When using authorization, maintaining a user database on the router is not required. User names can be configured on the RADIUS server. User names are temporary and are not saved in the configuration when the user session terminates. Temporary user login names and their associated passwords are not saved as part of the configuration.

2.1.2.3 TACACS+ Authorization

TACACS+ authorization operates in one of three ways:

- All users who authenticate via TACACS+ can use a single common default profile that is configured on the SR OS, or

- Each command attempted by a user is sent to the TACACS+ server for authorization
- The operator can configure local profiles and map **tacplus priv-lvl** based authorization to those profiles (the **use-priv-lvl** option)

To use a single common default profile to control command authorization for TACACS+ users, the operator must configure the **tacplus use-default-template** option and configure the parameters in the **user-template tacplus_default** to point to a valid local profile.

If the default template is not being used for TACACS+ authorization and the **use-priv-lvl** option is not configured, then each CLI command issued by an operator is sent to the TACACS+ server for authorization. The authorization request sent by the SR OS contains the first word of the CLI command as the value for the TACACS+ cmd and all following words become a cmd-arg. Quoted values are expanded so that the quotation marks are stripped off and the enclosed value are seen as one cmd or cmd-arg.

2.1.2.3.1 Examples

Here is a set of examples, where the following commands are typed in the CLI:

```
- "show"
- "show router"
- "show port 1/1/1"
- "configure port 1/1/1 description "my port"
```

This results in the following AVPairs:

```
cmd=show

cmd=show
cmd-arg=router

cmd=show
cmd-arg=port
cmd-arg=1/1/1

cmd=configure
cmd-arg=port
cmd-arg=1/1/1
cmd-arg=description
cmd-arg=my port
```

For TACACS+ authorization, the SR OS sends the entire CLI context in the **cmd** and **cmd-arg** values. Here is a set of examples where the CLI context is different:

```
- *A:dut-c# configure service
```



```
- *A:dut-c>config>service# vprn 555 customer 1 create
- *A:dut-c>config>service>vprn$ shutdown
```

This results in the following AVPairs:

```
cmd =configure
cmd-arg=service

cmd=configure
cmd-arg=service
cmd-arg=vprn
cmd-arg="555"
cmd-arg=customer
cmd-arg=1
cmd-arg=create

cmd=configure
cmd-arg=service
cmd-arg=vprn
cmd-arg="555"
cmd-arg=customer
cmd-arg=1
cmd-arg=create
cmd-arg=shutdown
```

2.1.3 Accounting

When enabled, RADIUS accounting sends command line accounting from the router to the RADIUS server. The router sends spars using UDP packets at port 1813 (decimal).

The router issues an accounting request packet for each event requiring the activity to be recorded by the RADIUS server. The RADIUS server acknowledges each accounting request by sending an accounting response after it has processed the accounting request. If no response is received in the time defined in the timeout parameter, the accounting request must be retransmitted until the configured retry count is exhausted. A trap is issued to alert the NMS (or trap receiver) that the server is unresponsive. The router issues the accounting request to the next configured RADIUS server (up to 5).

User passwords and authentication keys of any type are never transmitted as part of the accounting request.

2.1.3.1 RADIUS Accounting

Accounting tracks user activity to a specified host. When RADIUS accounting is enabled, the server is responsible for receiving accounting requests and returning a response to the client indicating that it has successfully received the request. Each command issued on the router generates a record sent to the RADIUS server. The record identifies the user who issued the command and the timestamp.

Accounting can be configured independently from RADIUS authorization and RADIUS authentication.

2.1.3.2 TACACS+ Accounting

The OS allows you to configure the type of accounting record packet that is to be sent to the TACACS+ server when specified events occur on the device. The **accounting record-type** parameter indicates whether TACACS+ accounting start and stop packets be sent or just stop packets be sent. Start/stop messages are only sent for individual commands, not for the session.

When a user logs in to request access to the network using Telnet or SSH, or a user enters a command for which accounting parameters are configured, or a system event occurs, such as a reboot or a configuration file reload, the router checks the configuration to see if TACACS+ accounting is required for the particular event.

If TACACS+ accounting is required, then, depending on the accounting record type specified, sends a start packet to the TACACS+ accounting server which contains information about the event.

The TACACS+ accounting server acknowledges the start packet and records information about the event. When the event ends, the device sends a stop packet. The stop packet is acknowledged by the TACACS+ accounting server.

2.2 Security Controls

You can configure routers to use RADIUS, TACACS+, and local authentication to validate users requesting access to the network. The order in which password authentication is processed among RADIUS, TACACS+ and local passwords can be specifically configured. In other words, the authentication order can be configured to process authorization through TACACS+ first, then RADIUS for authentication and

accounting. Local access can be specified next in the authentication order in the event that the RADIUS and TACACS+ servers are not operational. The security methods capabilities are listed in [Table 4](#).

Table 4 Security Methods Capabilities

Method	Authentication	Authorization	Accounting*
Local	Y	Y	N
TACACS+	Y	Y	Y
RADIUS	Y	Y	Y
* Local commands always perform account logging using the config log command.			

2.2.1 When a Server Does Not Respond

A trap is issued if a RADIUS + server is unresponsive. An alarm is raised if RADIUS is enabled with at least one RADIUS server and no response is received to either accounting or user access requests from any server.

Periodic checks to determine if the primary server is responsive again are not performed. If a server is down, it will not be contacted for 5 minutes. If a login is attempted after 5 minutes, then the server is contacted again. When a server does not respond with the health check feature enabled, the server's status is checked every 30 seconds. Health check is enabled by default. When a service response is restored from at least one server, the alarm condition is cleared. Alarms are raised and cleared on Nokia's Fault Manager or other third party fault management servers.

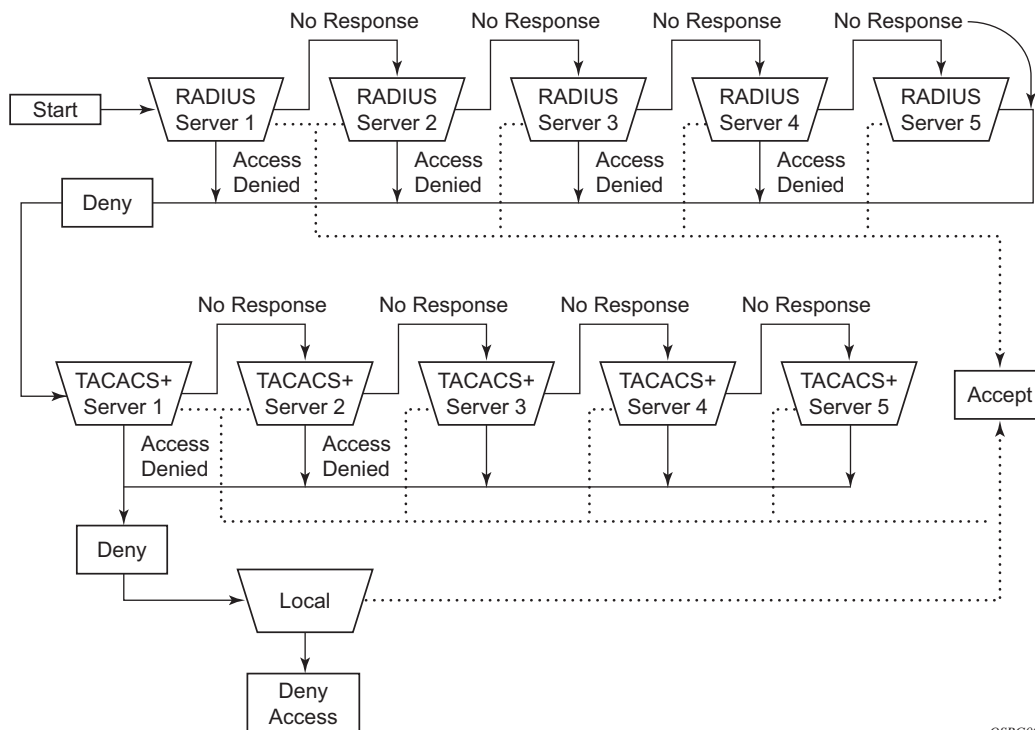
The servers are accessed in order from lowest to highest specified index (from 1 to 5) for authentication requests until a response from a server is received. A higher indexed server is only queried if no response is received, implying a lower indexed server is not available. If a response from the server is received, no other server is queried.

2.2.2 Access Request Flow

In [Figure 5](#), the authentication process is defined in the `config>system>security>password` context. The authentication order is determined by specifying the sequence in which password authentication is attempted among RADIUS, TACACS+, and local passwords. This example uses the authentication order of RADIUS, then TACACS+, and finally, local. An access request is sent to RADIUS server 1. One of two scenarios can occur. If there is no response from the server, the request is passed to the next RADIUS server with the next lowest index (RADIUS server 2) and so on, until the last RADIUS server is attempted (RADIUS server 5). If server 5 does not respond, the request is passed to the TACACS+ server 1. If there is no response from that server, the request is passed to the next TACACS+ server with the next lowest index (TACACS+ server 2) and so on.

If a request is sent to an active RADIUS server and the user name and password is not recognized, access is denied and passed on to the next authentication option, in this case, the TACACS+ server. The process continues until the request is either accepted, denied, or each server is queried. Finally, if the request is denied by the active TACACS+ server, the local parameters are checked for user name and password verification. This is the last chance for the access request to be accepted.

Figure 5 Security Flow



OSRG009

2.3 Centralized CPU Protection

The SR OS provides several rate limiting mechanisms to protect the CPM/CFM processing resources of the router:

- **Centralized CPU Protection:** A centralized rate limiting function that operates on the CPM to limit traffic destined to the CPUs. For legacy (historical) reasons, the term “centralized CPU protection” is referred to as “CPU protection” in this guide, in the CLI, and elsewhere.
- **Distributed CPU Protection:** A control traffic rate limiting protection mechanism for the CPM/CFM that operates on the line cards (hence ‘distributed’). See [Distributed CPU Protection \(DCP\)](#) for more information.

CPU protection protects the CPU of the node that it is configured on from a DOS attack by limiting the amount of traffic coming in from one of its ports and destined to the CPM (to be processed by its CPU) using a combination of the configurable limits.

Some of the limits are configured globally for the node, and some of the limits are configured in CPU Protection profiles which are assigned to interfaces.

The following limits are configured globally for the node (but take effect per port or per interface):

- **link-specific rate** — Applies to the link-specific protocols LACP (Ethernet LAG control) and LMI (ATM, Ethernet and Frame Relay). The rate is a per-link limit (each link in the system will have LACP/LMI packets limited to this rate).
- **port-overall-rate** – Applies to all control traffic each port. The rate is a per-port limit (each port in the system will have control traffic destined to the CPM limited to this rate).
- **protocol-protection** — Blocks network control traffic for unconfigured protocols. If IS-IS is not configured on an IP interface all IS-IS-related traffic will be dropped and not reach the CPU.

The following limits are configured within CPU Protection policies (1-255). CPU Protection policies are created, configured, and then assigned to interfaces.

- **overall-rate** — Applies to all control traffic destined to the CPM (all sources) received on the interface (only where the policy is applied). This is a per-interface limit. Control traffic received above this rate will be discarded.

- **per-source-rate** — Used to limit the control traffic destined to the CPM from each individual source. This per-source-rate is only applied when an object (SAP) is configured with a `cpu-protection` policy and also with the optional `mac-monitoring` or `ip-src-monitoring` keywords. A source is defined as a *SAP, Source MAC Address* tuple for `mac-monitoring` and as a *SAP, Source IP Address* tuples for `ip-src-monitoring`. Only certain protocols (as configured under *included-protocols* in the `cpu protection` policy) are limited (per source) when the `ip-src-monitoring` keyword is used.
- **out-profile-rate** – Applies to all control traffic destined to the CPM (all sources) received on the interface (only where the policy is applied). This is a per-interface limit. Control traffic received above this rate will be marked as discard eligible (such as, `out-profile/low-priority/yellow`) and is more likely to be discarded if there is contention for CPU resources.

A three-color marking mechanism uses a green, yellow and red marking function. This allows greater flexibility in how traffic limits are implemented. A CLI command within the DoS protection policy called **out-profile-rate** maps to the boundary between the green (accept) and yellow (mark as discard eligible/low priority) regions. The **overall-rate** command marks the boundary between the yellow and red (drop) regions point for the associated policy ([Figure 6](#)).

Figure 6 Profile Marking



There are two default CPU protection policies. They are modifiable, but cannot be deleted.

Policy 254:

- This is the default policy that is automatically applied to access interfaces
- Traffic above 6000 pps is discarded
- `overall-rate = 6000`
- `per-source-rate = max`
- `out-profile-rate = 6000`

Policy 255:

- This is the default policy that is automatically applied to Network interfaces

- Traffic above 3000 pps is marked as discard eligible, but is not discarded unless there is congestion in the queuing towards the CPU
- overall-rate = max
- per-source-rate = max
- out-profile-rate = 3000

All traffic destined to the CPM and that will be processed by its CPU will be subject to the limit specified. Therefore, if there is a protocol running on the violating interface, then protocol traffic on that interface will be affected. The objective of CPU protection is to limit the amount of traffic that the CPU will process at an early stage, therefore, the good and bad traffic coming in cannot be distinguished when it arrives at a rate higher than the user-configured limit.

If the overall rate is set to 1000 pps and as long as the total traffic that is destined to the CPM and intended to be processed by the CPU is less than or equal to 1000 pps, all traffic will be processed. If the rate exceeds 1000 pps, then protocol traffic is discarded (or marked as discard eligible/low priority in the case of the out-profile-rate) and traffic on the interface is affected.

This protects all the other interfaces on the system and make sure that a violation from one interface does not affect the rest of the box.

The protocol-protection configuration is not a rate (just an enable/disable configuration). When enabled, this feature causes the network processor on the CPM to discard all packets received for protocols that are not configured on the particular interface. This helps mitigate DoS attacks by filtering invalid control traffic before it hits the CPU. The system automatically populates and maintains a per-interface list of configured (such as valid) protocols (based on interface config, etc). For example, if an interface does not have IS-IS configured, then protocol-protection will discard any IS-IS packets received on that interface.

Some protocols are not bound to a specific interface, for example, BGP. The SR OS will discard packets for these protocols if the protocol is not configured anywhere in the system. Protection for the following protocols is achieved using the per-peer-queuing feature of the SR OS: BGP, T-LDP, LDP, MSDP.

Protocols controlled by the protocol-protection mechanism include:

- OSPFv2
- OSPFv3
- IS-IS
- RSVP-TE
- RIP
- PIM

- MLD
- IGMP
- L2TP
- PPPoE
- BFD
- GTP



Note: If PIM or PIM snooping is not configured on any interfaces/SAPs, then all PIM packets will be discarded. If PIM or PIM snooping is configured on an interface/SAP, then multicast PIM messages are filtered based on PIM being enabled on that particular interface. All unicast PIM messages are sent to the CPU to be processed.

The CPU protection features are supported on the following platforms:

- 7750 SR-7/SR-12
- 7750 SR-12e
- 7450 ESS-7/ESS-12
- 7950 XRS



Note: For more information about CPU protection, see “CPU Protection” and “Monitoring Attacks on the 7750 SR” sections in the SR OS Security Best Practices.

2.3.1 CPU Protection Extensions for ETH-CFM

CPU protection supports the ability to explicitly limit the amount of ETH-CFM traffic that arrives at the CPU for processing. ETH-CFM packets that are redirected to the CPU by either a Management Endpoint (MEP) or a Management Intermediate Point (MIP) will be subject to the configured limit of the associated policy. Up to four CPU protection policies may include up to ten individual eth-cfm specific entries. The eth-cfm entries allow the operator to apply a packet per second rate limit to the matching combination of level and opcode, for eth-cfm packet that are redirected to the CPU. Any eth-cfm traffic that is redirected to the CPU by a Management Point (MP) that does not match any entries of the applied policy is still subject to the overall rate limit of the policy itself. Any eth-cfm packets that are not redirected to the CPU are not subject to this function and are treated as transit data, subject to the applicable QoS policy.

The operator first creates a CPU Policy and includes the required eth-cfm entries. Overlap is allowed for the entries within a policy, first match logic is applied. This means ordering the entries in the proper sequence is important to ensure the proper behavior is achieved. Even though the number of eth-cfm entries is limited to ten, the entry numbers have a valid range from 1-100 to allow for ample space to insert policies between one and other.

Ranges are allowed when configuring the Level and the OpCode. Ranges provide the operator a simplified method for configuring multiple combinations. When more than one Level or OpCode is configured in this manner the configured rate limit is applied separately to each combination of level and OpCode match criteria. For example, if the Levels are configured as listed in [Table 5](#), with using a range of 5-7 and the OpCode is configured for 3,5 with a rate of 1. That restricts all possible combinations on that single entry to a rate of 1 packet per second. In this example six different match conditions are programmed behind the scene.

Table 5 Ranges versus Levels and OpCodes

Level	OpCode	Rate
5	3	1
5	5	1
6	3	1
6	5	1
7	3	1
7	5	1

Once the policy is created it must be applied to a SAP/Binding within a service for these rates to take affect. This means the rate is on a per SAP/Binding basis. Only a single policy may be applied to a SAP/Binding. The “eth-cfm-monitoring” option must be configured in order for the eth-cfm entries to be applied when the policy is applied to the SAP/Binding. If this option is not configured, eth-cfm entries in the policy will be ignored. It is also possible to apply a policy to a SAP/Binding configuring “eth-cfm-monitoring” which does not have an MP. In this case, although these entries are enforced, no packets are being redirect to the CPU due to the lack of an MP.

By default, rates are applied on a per peer basis. This means each individual peer is subject to the rate. However, it is suggested that the “aggregate” option be configured to apply the rate to the sum total of all peers. MIPs for example only respond to Loopback Messages and Linktrace Messages. These are typically on demand functions and per peer rate limiting is likely not required thus making the aggregate function a more appealing model.

“eth-cfm-monitoring” and “mac-monitoring” are mutually exclusive and cannot be configured on the same SAP/Binding “mac-monitoring” is used in combination with the traditional CPU protection and is not specific to the eth-cfm rate limiting feature describe here.

When an MP is configured on a SAP/Binding within a service which allows an external source to communicate with that MP, for example a User to Network Interface (UNI), it is suggested that “eth-cfm-monitoring” with the “aggregate” option be configured on all SAP/Bindings to provide the highest level of rate control.

The example below shows a sample configuration for a policy and the application of that policy to a SAP in a VPLS service configured with a MP.

Policy 1 entry 10 limits all eth-cfm traffic redirected to the CPU for all possible combinations to 1 packet per second. Policy 1 entry 20 limits all possible combinations to a rate of zero, dropping all request which match any combination. If entry 20 did not exist then only rate limiting of the entry 10 matches would occur and any other eth-cfm packets redirected to the CPU would not be bound by a CPU protection rate.

```
config>sys>security>cpu-protection#
  policy 1
    eth-cfm
      entry 10 level 5-7 opcode 3,5 rate 1
      entry 20 level 0-7 opcode 0-255 rate 0

config>service>vpls#
  sap 1/1/4:100
    cpu-protection 1 eth-cfm-monitoring aggregate
    eth-cfm
      mip
    no shutdown
```

The centralized CPU protection features are supported on the following platforms:

- 7750 SR-7/SR-12
- 7450 ESS-7/ESS-12
- 7950 XRS


2.3.2 ETH-CFM Ingress Squelching

CPU protection provides a granular method to control which ETH-CFM packets are processed. As indicated in the previous section, a unique rate can be applied to ETH-CFM packets classifying on specific MD-Level and specific OpCode and applied to both ingress (Down MEP and ingress MIP) and egress (Up MEP and egress MIP) extraction. That function is to protect the CPU upon extraction when a Management Point (MP) is configured.

It is also important to protect the ETH-CFM architecture deployed in the service provider network. The protection scheme here varies from CPU protection. This model is used to prevent ETH-CFM frames at the service provider MD-levels from gaining access to the network even when extraction is not in place. ETH-CFM squelching allows the operator to achieve this goal using a simple method to drop all ETH-CFM packets at or below the configured MD-level. The ETH-CFM squelch feature is ingress only.

Figure 7 shows a typical ETH-CFM hierarchical model with a Subscriber ME (6), Test ME (5), EVC ME (4) and an Operator ME (2). This model provides the necessary transparency at the different levels of the architecture. For security reasons, it may be necessary to prevent errant levels from entering the service provider network at the UNI, ENNI, or other untrusted interconnection points. Configuring squelching at level four on both UNI-N interconnection ensures that ETH-CFM packets matching the SAP or binding delimited configuration will silently discard ETH-CFM packets at ingress.

Figure 7 ETH-CFM Hierarchical Model



Squelching configuration uses a single MD-level [0..7] to silently drop all ETH-CFM packets matching the SAP or binding delimited configuration at and below the specified MD-level. In Figure 7, a squelch level is configured at MD-level 4. This means the configuration will silently discard MD-levels 0, 1, 2, 3 and 4, assuming there is a SAP or binding match.



Note: Extreme caution must be used when deploying this feature.

The operator is able to configure Down MEPs and ingress MIPs that conflict with the squelched levels. This also means that any existing MEP or MIP processing ingress CFM packets on a SAP on Binding where a squelching policy is configured will be interrupted as soon as this command is entered into the configuration. These MPs will not be able to receive any ingress ETH-CFM frames because squelching is processed before ETH-CFM extraction.

CPU Protection Extensions for ETH-CFM are still required in the model above because the Subscriber ME (6) and the Test ME (5) are entering the network across an untrusted connection, the UNI. ETH-CFM squelching and CPU Protection for ETH-CFM can be configured on the same SAP or binding. Squelching is first in the process order followed by CPU Protection for ETH-CFM.

MPs configured to support primary VLAN are not subjected to the squelch function. Primary VLAN based MPs, supported only on Ethernet SAPs, are extractions that take into consideration an additional VLAN beyond the SAP configuration.

The difference in the two protection mechanisms is shown in the [Table 6](#). CPU Protection is used to control access to the CPU resources when processing is required. Squelching is required when the operator is protecting the ETH-CFM architecture from external sources.

Table 6 CPU Protection and Squelching

Description	CPU Protection Extension for ETH-CFM	ETH-CFM Squelching
Ingress Filtering	Yes	Yes
Egress Filtering	Yes	No
Granularity	Specified Level AND OpCode	Level (At and below)
Rate	Configurable Rate (includes 0=drop all)	Silent Drop
Primary VLAN Support	Rate shared with SAP delineation	Not exposed to squelch
Extraction	Requires MEP or MIP to extract	No MEP or MIP required

As well as including the squelching information under the **show service service-id all**, display output the **squelch-ingress-level** key has been added to the **sap-using** and **sdp-using show** commands.

```
show service sap-using squelch-ingress-levels
=====
ETH-CFM Squelching
=====
PortId          SvcId          Squelch Level
-----
6/1/1:100.*     1              0 1 2 3 4 5 6 7
lag-1:100.*     1              0 1 2 3 4
6/1/1:200.*     2              0 1 2
lag-1:200.*     2              0 1 2 3 4 5
-----
Number of SAPs: 4
-----
show service sdp-using squelch-ingress-levels
=====
ETH-CFM Squelching
=====
SdpId          SvcId          Type Far End          Squelch Level
-----
12345:4000000000 2147483650    Spok 1.1.1.1          0 1 2 3 4
=====
```



Note: Extreme caution must be used when deploying this feature.

2.4 Distributed CPU Protection (DCP)

The SR OS provides several rate limiting mechanisms to protect the CPM/CFM processing resources of the router:

- **Centralized CPU Protection:** A centralized rate limiting function that operates on the CPM to limit traffic destined to the CPUs. See [Centralized CPU Protection](#) for more information. For legacy (historical) reasons, the term “centralized CPU protection” is referred to as “CPU protection” in this guide, in the CLI, and elsewhere.
- **Distributed CPU Protection:** A control traffic rate limiting protection mechanism for the CPM/CFM that operates on the line cards (hence ‘distributed’).

Distributed CPU Protection (DCP) offers a powerful per-protocol-per-object (examples of objects are SAPs and network interfaces) rate limiting function for control protocol traffic that is extracted from the data path and sent to the CPM. The DCP function is implemented on the router line cards that allows for high levels of scaling and granularity of control.

The DCP rate limiting is configured via policies that are applied to objects (for example, SAPs).

The basic types of policers in DCP are:

- Enforcement Policers — An instance of a policer that is policing a flow of packets comprised of a single (or small set of) protocols(s) arriving on a single object (for example, SAP). Enforcement policers perform a configurable action (for example, discard) on packets that exceed configured rate parameters. There are two basic sub-types of enforcement policers:
 - Static policers — always instantiate.
 - Dynamic policers — only instantiated (allocated from a free pool of dynamic policers) when a local monitor detects non-conformance for a set of protocols on a specific object.
- Local Monitors — A policer that is primarily used to measure the conformance of a flow comprised of multiple protocols arriving on a single object. Local monitors are used as a trigger to instantiate dynamic policers.

The use of dynamic policers reduces the number of policers required to effectively monitor and control a set of protocols across a large set of objects since the per-protocol-per-object dynamic policers are only instantiated when an attack or misconfiguration occurs, and they are only instantiated for the affected objects.

Figure 8 Per SAP per Protocol Static Rate Limiting with DCP

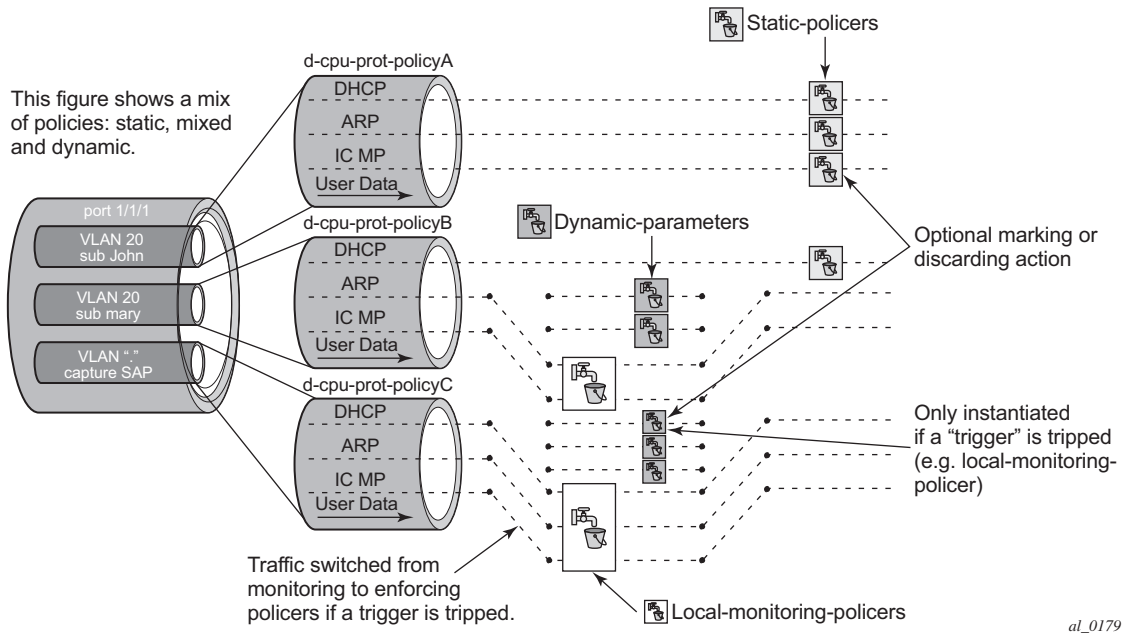
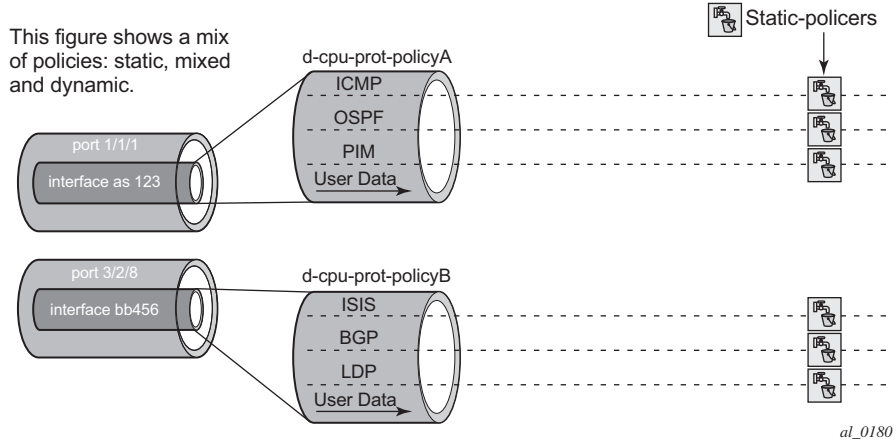


Figure 9 Per Network Interface per Protocol Static Rate Limiting with DCP



2.4.1 Applicability of Distributed CPU Protection

The system assigns a default Distributed CPU Protection (DCP) policy to newly created access and network interfaces. Originally, these policies, "_default-access-policy" and "_default-network-policy", are created empty and are modifiable by the operator. Additional DCP policies can be created for interfaces requiring a dedicated configuration.

If DCP functionality is not required on a given access or network interface, then an empty DCP policy can be created and explicitly assigned to the interface.

DCP policies can be applied to the following types of objects:

- most types of SAPs, including capture SAPs, SAPs on pseudowires, B-VPLS SAPs and VPLS template SAPs, but are not applicable to Epipe template SAPs and video ISA SAPs
- network interfaces, but not to any other type of interface, a DCP policy can be configured at the interface SAP instead

Control packets that are both forwarded (which means they could be subject to normal QoS policy policing) and also copied for extraction are not subject to Distributed CPU Protection (including in the all-unspecified bucket). This includes traffic snooping (for example, PIM in VPLS) as well as control traffic that is flooded in an R-VPLS instance and also extracted to the CPM such as ARP, ISIS and VRRP. Centralized per SAP and interface cpu-protection can be employed to rate limit or mark this traffic if desired.

Control traffic that arrives on a network interface, but inside a tunnel (for example, SDP, LSP, PW) and logically terminates on a service (that is, traffic that is logically extracted by the service rather than the network interface layer itself) will bypass the DCP function. The control packets in this case will not be subject to the DCP policy that is assigned to the network interface on which the packets arrived. This helps to avoid customer traffic in a service from impacting other services or the operator's infrastructure.

Control packets that are extracted in a vprn service, where the packets arrived into the node via a VPLS SAP (that is, r-vpls scenario), will use the DCP policy and policer instances associated with the VPLS SAP. In this case the DCP policy that an operator creates for use on VPLS SAPs, for VPLSs that have a Layer 3-interface bound to them (r-vpls), may have protocols such as OSPF, ARP, configured in the policy.

2.4.2 Log Events, Statistics, Status and SNMP support

A comprehensive set of log events are supported for DCP in order to alert the operator to potential attacks or misconfigurations and to allow tuning of the DCP settings. Refer to the NOTIFICATION-TYPE objects with “Dcp” in the names in the following MIBs for details:

- TIMETRA-CHASSIS-MIB
- TIMETRA-SAP-MIB
- TIMETRA-VRTR-MIB

The log events can also be seen in the CLI using the following **show log event-control | match Dcp** command

DCP throttles the rate of DCP events to avoid event floods when multiple parallel attacks or problems are occurring.

Many of the DCP log events can be individually enabled or disabled at the DCP policy level (in the DCP policy config) as well as globally in the system (in log event-control).

If needed when a DCP log event indicates a SAP, and that SAP is an MSAP, the operator can determine which subscriber(s) is/are on a specific MSAP by using the **show service active-subs** command and then filtering (“| match”) on the msap string.

Statistics and status related to DCP are available both via:

- CLI
- SNMP — See various tables and objects with “Dcp” or “DCpuProt” in their name in the TIMETRA-CHASSIS-MIB, TIMETRA-SECURITY-MIB, TIMETRA-SAP-MIB and TIMETRA-VRTR-MIB

2.4.3 DCP Policer Resource Management

The policer instances are a limited h/w resource on a given forwarding plane. DCP policers (static, dynamic, local-monitor) are consumed from the overall forwarding plane policer resources (from the ingress resources if ingress and egress are partitioned). Each per-protocol policer instantiated reduces the number of FP child policers available for other purposes.

When DCP is configured with dynamic enforcement, then the operator must set aside a pool of policers that can be instantiated as dynamic enforcement policers. The number of policers reserved for this function are configurable per card/FP. The policers in this pool are not available for other purposes (normal SLA enforcement).

Static enforcement policers and local monitoring policers use policers from the normal/global policer pool on the card/FP. Once a static policer is configured in a DCP policy and it is referenced by a protocol in the policy, then this policer will be instantiated for each object (SAP or network interface) that is created and references the policy. If there is no policer free on the associated card/FP, then the object will be blocked from being created. Similarly for local monitors: once a local monitoring policer is configured and referenced by a protocol, then this policer will be instantiated for each object that is created and references the policy. If there is no policer free, then the object will be blocked from being created.

Dynamic enforcement policers are allocated as needed (when the local monitor detects non-conformance) from the reserved dynamic-enforcement-policer-pool.

When a DCP policy is applied to an object on a LAG, then a set of policers is allocated on each forwarding plane (on each line card that contains a member of the LAG). The LAG mode is ignored and the policers are always shared by all ports in the LAG on that forwarding plane on the SAP/interface. In other words, with link-mode lag a set of DCP policers are not allocated per port in the LAG on the SAP.

In order to support large scale operation of DCP, and also to avoid overload conditions, a polling process is used to monitor state changes in the policers. This means there can be a delay between when an event occurs in the data plane and when the relevant state change or event notification occurs towards an operator, but in the meantime the policers are still operating and protecting the control plane.

2.4.4 Operational Guidelines and Tips

The following points offer various optional guidelines that may help an operator decide how to leverage Distributed CPU Protection.

- The rates in a policy assigned to a capture SAP should be higher than those assigned to MSAPs that will contain a single subscriber. The rates for the capture sap policy should allow for a burst of MSAP setups.
- To completely block a set of specific protocols on a given SAP, create a single static policer with a rate of 0 and map the protocols to that policer. Dynamic policers and local monitors can't be used to simultaneously allow some protocols but block others (the non-zero rates in the monitor would let all protocols slip through at a low rate).

- During normal operation it is recommended to configure “log-events” (no verbose keyword) for all static-policers, in the dynamic-parameters of all protocols and for all local-monitoring-policers. The verbose keyword can be used selectively during debug, testing, tuning and investigations.
- Packet based rate limiting is generally recommended for low rate subscriber based protocols whereas kbps rate limiting is recommended for higher rate infrastructure protocols (such as BGP).
- It is recommended to configure an exceed-action of low-priority for routing and infrastructure protocols. Marked packets are more likely to be discarded if there is congestion in the control plane of the router, but will get processed if there is no contention for CPU resources allowing for a work-conserving behavior in the CPM.
- In order to assign a different dist-cpu-protection policy to a specific MSAP (instance) or to all MSAPs for a specific msap policy, the operator can assign a new dist-cpu-protection policy to the MSAP policy and then use the **eval-msap** tool:

```
A:nodeA>tools>perform# subscriber-mgmt eval-msap  
- eval-msap {policy <msap-policy-name> | msap <sap-id>}
```



Note: Any new MSAPs will also be assigned the new dist-cpu-protection policy.

- If needed, an operator can determine which subscriber is on a specific MSAP by using the **show service active-subs** command and then filtering (“| match”) on the msap string.
- If protocol X is trusted, and using the “all-undefined” protocol is not required, then simply avoid creating protocol X in the policy configuration.
- If protocol X is trusted, but the all-undefined bucket is required, then there are two options:
 - avoid creating protocol X so that it is treated as part of the all-undefined bucket (but account for the packets from X in the all-undefined rate and local-mon rate), or
 - create protocol X and configure it to bypass.

2.5 Classification-Based Priority for Extracted Protocol Traffic

The SR OS supports a set of mechanisms to protect the router control and management planes from various types of attacks, floods, and misconfigurations. Many of the mechanisms operate by default with no need for operator configuration or intervention.

One class of mechanisms employed on the router to protect against floods of control traffic involves identifying potentially harmful or malicious traffic through the use of rate measurements. Centralized CPU protection protects and isolates interfaces from each other by default by treating unexpectedly high rate control traffic on an interface as lower priority (to be discarded if the control plane experiences congestion). Distributed CPU protection can protect and isolate at a per-protocol, per-interface granularity through configured rate profiles. These rate-based protection mechanisms make no assumptions about the contents of the packets and can be used when nothing about the packets can be trusted (for example, DSCP or source IP address, which can be spoofed).

The SR OS also supports an alternative to rate-based mechanisms for cases where the packet headers can be trusted to differentiate between good and bad control traffic. A configurable prioritization scheme can be enabled (using the **init-extract-prio-mode I3-classify** command) on a per-FP basis to initialize the drop priority of all Layer 3 extracted control traffic based on the QoS classification of the packets. This is useful, for example, in networks where the DSCP and EXP markings can be trusted as the primary method to distinguish, protect, and isolate good terminating protocol traffic from unknown or potentially harmful protocol traffic instead of using the rate-based distributed CPU protection and centralized CPU protection traffic marking/coloring mechanisms (for example, **out-profile-rate** and **exceed-action low-priority**).

The operational guidelines for deploying classification-based priority for extracted control traffic are as follows.

- Centralized CPU protection should be effectively disabled for all interfaces/SAPs on FPs configured in **I3-classify** mode by changing some CPU protection policy parameters from their default values. This is required so that centralized CPU protection does not re-mark good control traffic (traffic that was initially classified as high priority) as low priority if a flood attack occurs on the same interface. Effectively disabling centralized CPU protection can be done by ensuring that:
 - a rate value of **max** is configured for **port-overall-rate** (**max** is the default value for **port-overall-rate**)

- all objects (interfaces, MSAP policies, and SAPs) that can be assigned a CPU protection policy are referencing a policy that sets the **out-profile-rate** to **max** and the **overall-rate** to **max** (this can be done in the two default CPU protection policies if all FPs in the system are in **I3-classify** mode)
- DCP can be used in conjunction with **I3-classify** mode, but care must be taken to prevent DCP from acting on protocols where the operator wants to use QoS classification (such as DSCP or EXP) to differentiate between good and bad Layer 3 packets. On an FP with **I3-classify** mode, DCP should be configured so that BGP, LDP, and other protocols do not have their initial drop priority (color) overwritten by DCP if the QoS classification of these protocols is trusted. This can be achieved by using **exceed-action none** for those protocols in a DCP policy. For other protocols where QoS classification cannot be used to distinguish between good and bad extracted packets, DCP can be used to color the packets with a drop priority based on a configured rate.
- If any LAG member is on an FP in **I3-classify** mode, all FPs that host the other members of that LAG should also be in **I3-classify** mode.
- The QoS classification rules that are used on interfaces/SAPs on FPs in **I3-classify** mode should be configured to differentiate between good and bad control traffic. The default network ingress QoS policies do differentiate (for example, based on DSCP), but the default access ingress QoS policies do not.

The **I3-classify** mode for extracted control traffic is supported on the 7750 SR and 7950 XRS.

2.6 Vendor-Specific Attributes (VSAs)

The software supports the configuration of Nokia-specific RADIUS attributes. These attributes are known as vendor-specific attributes (VSAs) and are discussed in RFC 2138. VSAs must be configured when RADIUS authorization is enabled. It is up to the vendor to specify the format of their VSA. The attribute-specific field is dependent on the vendor's definition of that attribute. The Nokia-defined attributes are encapsulated in a RADIUS vendor-specific attribute with the vendor ID field set to 6527, the vendor ID number.



Note: The PE-record entry is required to support the RADIUS Discovery for Layer 2 VPN feature. A PE-record is only relevant if the RADIUS Discovery feature is used, not for the standard RADIUS setup.

The following RADIUS vendor-specific attributes (VSAs) are supported by Nokia.

- `timetra-access <ftp> <console> <both>` — This is a mandatory command that must be configured. This command specifies if the user has FTP and /or console (serial port, Telnet, and SSH) access.
- `timetra-profile <profile-name>` — When configuring this VSA for a user, it is assumed that the user profiles are configured on the local router and the following applies for local and remote authentication:
 1. The authentication-order parameters configured on the router must include the local keyword.
 2. The user name may or may not be configured on the router.
 3. The user must be authenticated by the RADIUS server
 4. Up to 8 valid profiles can exist on the router for a user. The sequence in which the profiles are specified is relevant. The most explicit matching criteria must be ordered first. The process stops when the first complete match is found.

If all the above mentioned conditions are not met, then access to the router is denied and a failed login event/trap is written to the security log.
- `timetra-default-action <permit-all|deny-all|none>` — This is a mandatory command that must be configured even if the `timetra-cmd` VSA is not used. This command specifies the default action when the user has entered a command and no entry configured in the `timetra-cmd` VSA for the user resulted in a match condition.
- `timetra-cmd <match-string>` — Configures a command or command subtree as the scope for the match condition.

The command and all subordinate commands in subordinate command levels are specified.

2.7 Other Security Features

This section describes the other security features supported by the SR OS.

2.7.1 Secure Shell (SSH)

Secure Shell Version 1 (SSH) is a protocol that provides a secure, encrypted Telnet-like connection to a router. A connection is always initiated by the client (the user). Authentication takes place by one of the configured authentication methods (local, RADIUS, or TACACS+). With authentication and encryption, SSH allows for a secure connection over an insecure network.

The OS allows you to configure Secure Shell (SSH) Version 2 (SSH2). SSH1 and SSH2 are different protocols and encrypt at different parts of the packets. SSH1 uses server as well as host keys to authenticate systems whereas SSH2 only uses host keys. SSH2 does not use the same networking implementation that SSH1 does and is considered a more secure, efficient, and portable version of SSH.

SSH runs on top of a transport layer (like TCP or IP), and provides authentication and encryption capabilities.

The OS has a global SSH server process to support inbound SSH and SCP sessions initiated by external SSH or SCP client applications. The SSH server supports SSHv1. This server process is separate from the SSH and SCP client commands on the routers which initiate outbound SSH and SCP sessions.

Inbound SSH sessions are counted as inbound telnet sessions for the purposes of the maximum number of inbound sessions specified by Login Control. Inbound SCP sessions are counted as inbound ftp sessions by Login Control.

When SSH server is enabled, an SSH security key is generated. The key is only valid until either the node is restarted or the SSH server is stopped and restarted (unless the preserve-key option is configured for SSH). The key size is non-configurable and set at 1024 bits. When the server is enabled, both inbound SSH and SCP sessions will be accepted provided the session is properly authenticated.

When the global SSH server process is disabled, no inbound SSH or SCP sessions will be accepted.

When using SCP to copy files from an external device to the file system, the SCP server will accept either forward slash (“/”) or backslash (“\”) characters to delimit directory and/or filenames. Similarly, the SCP client application can use either slash or backslash characters, but not all SCP clients treat backslash characters as equivalent to slash characters. In particular, UNIX systems will often times interpret the backslash character as an “escape” character which does not get transmitted to the SCP server. For example, a destination directory specified as “cf1:\dir1\file1” will be transmitted to the SCP server as “cf1:dir1file1” where the backslash escape characters are stripped by the SCP client system before transmission. On systems where the client treats the backslash like an “escape” character, a double backslash “\\” or the forward slash “/” can typically be used to properly delimit directories and the filename.

Two cipher lists, the client-cipher-list and the server-cipher-list, can be configured for negotiation of the best compatible ciphers between the client and server. The two cipher lists can be created and managed under the security ssh sub menu. The client-cipher-list is used when the SR OS is acting as ssh client and the server-cipher-list is used when the SR OS is acting as a server. The first cipher matched on the lists between the client and server is the preferred cipher for the session.

2.7.2 SSH PKI Authentication

The SR OS supports Secure Shell Version 2, but user authentication appears to be limited to using a username and password.



Note: SSHv1 is not supported when the node is running in FIPS-140-2 mode.

SSH also supports public key authentication whereby the client can provide a signed message that has been encrypted by his private key. As long as the server has been previously configured to know the client's public key, the server can authenticate the client.

Using Public Key authentication (also known as Public Key Infrastructure - PKI) can be more secure than the existing username/password method for a few reasons:

- A user will typically re-use the same password with multiple servers. If the password is compromised, the user must reconfigure the password on all affected servers.
- A password is not transmitted between the client and server using PKI. Instead the sensitive information (the private key) is kept on the client. Therefore it is less likely to be compromised.

This feature includes server side support for SSHv2 public key authentication. It does not include a key generation utility.

Support for PKI should be configured in the system level configuration where one or more public keys may be bound to a username. It should not affect any other system security or login functions.

2.7.2.1 Key Generation

Before SSH can be used with PKI, someone must generate a public/private key pair. This is typically supported by the SSH client software. For example, PuTTY supports a utility called PuTTYgen that will generate key pairs.

SSHv2 supports both RSA and DSA keys. The Digital Signature Algorithm is a U.S. Federal Government standard for digital signatures. PuTTYGen can be used to generate either type of key. The SR OS currently supports only RSA keys.

Assume the client is using PuTTY. First the user generates a key pair using PuTTYgen. The user sets the key type (SSH-1 RSA, SS-2 RSA, or SSH-2 DSA) and sets the number of bits to be used for the key (default = 1024). The user can also configure a passphrase that will be used to store the key locally in encrypted form. If the passphrase is configured the user must enter the passphrase in order to use the private key. Thus, it is a password for the private key. If the passphrase is not used the key is stored in plain text locally.

Next the user must configure the server to use his public key. This typically requires the user to add the public key to a file on the server. For example, if the server is using OpenSSH, the key must be added to the `ssh/authorized_keys` file. On the SR OS, the user can program the public Key via Telnet/SSH or SNMP.

2.7.3 Per Peer CPM Queuing

System-level security is crucial in service provider networks to address the increased threat of Denial-of-Service (DoS) attacks.

Control Processor Module Queuing (CPMQ) implements separate hardware-based queues which are allocated on a per-peer basis. CPMQ allocates a separate queue for each LDP and BGP peer and ensures that each queue is served in a round-robin fashion. This mechanism guarantees fair and “non-blocking” access to shared CPU resources across all peers. This would ensure, for example, that an LDP-based DoS attack from a given peer would be mitigated and compartmentalized so that not all CPU resources would be dedicated to the otherwise overwhelming control traffic sent by that specific peer.

CPMQ, using the **per-peer-queuing** command, ensures that service levels would not (or only partially be) impacted in case of an attack from a spoofed LDP or BGP peer IP address. SSH and Telnet supports per-peer queuing when the **login-control ttl-security** command is enabled.

2.7.4 CPM Filters and Traffic Management

Nokia routers have traffic management and queuing hardware dedicated to protecting the control plane.

CPM filters can be used to drop or accept packets, as well as allocate dedicated hardware shaping (CPM) queues for traffic directed to the control processors.

Users can allocate dedicated CPM hardware queues for certain traffic designated to the CPUs and can set the corresponding rate-limit for the queues.

CPM filters and queues control all traffic going in to the CPM from IOMs/XMAs, including all routing protocols. CPM filters apply to packets from all network and access ports, but not to packets from a management Ethernet port. CPM packet filtering and queuing is performed by network processor hardware using no resources on the main CPUs. CPM filters and queues are not configurable on one-slot chassis.

There are three filters that can be configured as part of the CPM filter policy: IP (v4) filter, IPv6 filter and MAC filter.

The SR OS filter implementation exits the filter when the first match is found and execute the actions according to the specified action. For this reason, entries must be sequenced correctly from most to least explicit. When both mac-filter and ip-filter/ipv6-filter are to be applied to a given traffic, mac-filter is applied first.

An entry of an IP(v4), IPv6, MAC CPM filters must have at least one match criteria defined to be active. A default action can be specified for CPM filter policy that applies to each of IP, IPv6, MAC filters that are in a **no shutdown** state as long as the CPM filter policy has at least one active filter entry in any of the IP(v4), IPv6, and MAC filters.

2.7.5 TTL Security for BGP and LDP

The BGP TTL Security Hack (BTSH) was originally designed to protect the BGP infrastructure from CPU utilization-based attacks. It is derived on the fact that the vast majority of ISP eBGP peerings are established between adjacent routers. Since TTL spoofing cannot be performed, a mechanism based on an expected TTL value can provide a simple and reasonably robust defense from infrastructure attacks based on forged BGP packets.

While TSH is most effective in protecting directly connected peers, it can also provide a lower level of protection to multi-hop sessions. When a multi-hop BGP session is required, the expected TTL value can be set to 255 minus the configured range-of-hops. This approach can provide a qualitatively lower degree of security for BGP (for example, a DoS attack could, theoretically, be launched by compromising a box in the path). However, BTSH will catch a vast majority of observed distributed DoS (DDoS) attacks against eBGP. For further information, refer to draft-gill-btsh-xx.txt, *The BGP TTL Security Hack (BTSH)*.

TSH can be used to protect LDP peering sessions as well. For details, see draft-chen-ldp-ttl-xx.txt, *TTL-Based Security Option for LDP Hello Message*.

The TSH implementation supports the ability to configure TTL security per BGP/LDP peer and evaluate (in hardware) the incoming TTL value against the configured TTL value. If the incoming TTL value is less than the configured TTL value, the packets are discarded and a log is generated.

2.7.6 Exponential Login Backoff

A malicious user may attempt to gain CLI access by means of a dictionary attack using a script to automatically attempt to login as an “admin” user and using a dictionary list to test all possible passwords. Using the exponential-back off feature in the **config>system>login-control** context the OS increases the delay between login attempts exponentially to mitigate attacks.

A malicious user may attempt to gain CLI access by means of a dictionary attack using a script to automatically attempt to login as an “admin” user and using a dictionary list to test all possible passwords. Using the exponential-back off feature in the **config>system>login-control** context the OS increases the delay between login attempts exponentially to mitigate attacks.

When a user tries to login to a router using a Telnet or an SSH session, there are a limited number of attempts allowed to enter the correct password. The interval between the unsuccessful attempts change after each try (1, 2 and 4 seconds). If the system is configured for user lockout, then the user will be locked out when the number of attempts is exceeded.

However, if lockout is not configured, there are three password entry attempts allowed after the first failure, at fixed 1, 2 and 4 second intervals, in the first session, and then the session terminates. Users do not have an unlimited number of login attempts per session. After each failed password attempt, the wait period becomes longer until the maximum number of attempts is reached.

The OS terminates after four unsuccessful tries. A wait period will never be longer than 4 seconds. The periods are fixed and will restart in subsequent sessions.

The **config>system>login-control>[no] exponential-backoff** command works in conjunction with the **config>system>security>password>attempts** command, which is also a system wide configuration.

For example:

```
*A:ALA-48>config>system# security password attempts
- attempts <count> [time <minutes1>] [lockout <minutes2>]
- no attempts

<count>                : [1..64]
<minutes1>             : [0..60]
```

<minutes2> : [0..1440]

Exponential backoff applies to any user and by any login method such as console, SSH and Telnet.

Refer to [Configuring Login Controls](#). The commands are described in [Login, Telnet, SSH and FTP Commands](#).

2.7.7 User Lockout

When a user exceeds the maximum number of attempts allowed (the default is 3 attempts) during a certain period of time (the default is 5 minutes), the account used during those attempts will be locked out for a pre-configured lock-out period (the default is 10 minutes).

A security or LI event log will be generated as soon as a user account has exceeded the number of allowed attempts, and the **show>system>security>user** command can be used to display the total number of failed attempts per user.

In addition to the security or LI event log, an SNMP trap is also generated so that any SNMP server (including the NSP NFM-P) can use the trap for an action.

The account will be automatically re-enabled as soon as the lock-out period has expired. The list of users who are currently locked out can be displayed with the **show>system>security>lockout** command.

A lock-out for a specific user can be administratively cleared using the **admin>user user-name>clear-lockout** command.

2.7.8 CLI Login Scripts

The SR OS supports automatic execution of CLI scripts when a user successfully logs into the router and starts a CLI session.

Users who authenticate via the local user database can use the configurable **configure>system>security>user user-name>console>login-exec file-url** login exec script.

A global login-script can be configured to execute a common script when any user logs into CLI. A per user login-script can also be configured to execute when a specific user logs into CLI. These login-scripts execute whether the user was authenticated via the local user database, TACACS+ or RADIUS. The scripts can be used, for example, to define a common set of CLI aliases that are made available on the router for all users.

To configure a global login exec script, use the **configure>system>login-control>login-scripts> global *file-url* script**.

To configure a user-specific login exec script, use the **configure>system>login-control>login-scripts>per-user>user-directory>*file-url* file-name *file-name* script**.

2.7.9 802.1x Network Access Control

The SR OS supports network access control of client devices (PCs, STBs, etc.) on an Ethernet network using the IEEE. 802.1x standard. 802.1x is known as Extensible Authentication Protocol (EAP) over a LAN network or EAPOL.


2.7.10 TCP Enhanced Authentication Option

The TCP Enhanced Authentication Option, currently covered in draft-bonica-tcp-auth-05.txt, *Authentication for TCP-based Routing and Management Protocols*, extends the previous MD5 authentication option to include the ability to change keys without tearing down the session, and allows for stronger authentication algorithms to be used.

The TCP Enhanced Authentication Option is a TCP extension that enhances security for BGP, LDP and other TCP-based protocols. This includes the ability to change keys in a BGP or LDP session seamlessly without tearing down the session. It is intended for applications where secure administrative access to both the end-points of the TCP connection is normally available.

TCP peers can use this extension to authenticate messages passed between one another. This strategy improves upon current practice, which is described in RFC 2385, *Protection of BGP Sessions via the TCP MD5 Signature Option*. Using this new strategy, TCP peers can update authentication keys during the lifetime of a TCP connection. TCP peers can also use stronger authentication algorithms to authenticate routing messages.

2.7.10.1 Packet Formats



Option Syntax

- Kind: 8 bits
The Kind field identifies the TCP Enhanced Authentication Option. This value will be assigned by IANA.
- Length: 8 bits
The Length field specifies the length of the TCP Enhanced Authentication Option, in octets. This count includes two octets representing the Kind and Length fields.
The valid range for this field is from 4 to 40 octets, inclusive.
For all algorithms specified in this memo the value will be 16 octets.
- T-Bit: 1 bit
The T-bit specifies whether TCP Options were omitted from the TCP header for the purpose of MAC calculation. A value of 1 indicates that all TCP options other than the Extended Authentication Option were omitted. A value of 0 indicates that TCP options were included.
The default value is 0.
- K-Bit: 1 bit
This bit is reserved for future enhancement. Its value MUST be equal to zero.
- Alg ID: 6 bits
The Alg ID field identifies the MAC algorithm.
- Res: 2 bits
These bits are reserved. They MUST be set to zero.
- Key ID: 6 bits
The Key ID field identifies the key that was used to generate the message digest.
- Authentication Data: Variable length
- The Authentication Data field contains data that is used to authenticate the TCP segment. This data includes, but need not be restricted to, a MAC. The length and format of the Authentication Data Field can be derived from the Alg ID.

- The Authentication for TCP-based Routing and Management Protocols draft provides an overview of the TCP Enhanced Authentication Option. The details of this feature are described in draft-bonica-tcp-auth-04.txt.

2.7.10.2 Keychain

The keychain mechanism allows for the creation of keys used to authenticate protocol communications. Each keychain entry defines the authentication attributes to be used in authenticating protocol messages from remote peers or neighbors, and it must include at least one key entry to be valid. Through the use of the keychain mechanism, authentication keys can be changed without affecting the state of the associated protocol adjacencies for OSPF, IS-IS, BGP, LDP, and RSVP-TE.

Each key within a keychain must include the following attributes for the authentication of protocol messages:

- key identifier
- authentication algorithm
- authentication key
- direction
- start time

In addition, additional attributes can be optionally specified, including:

- end time
- tolerance

[Table 7](#) shows the mapping between these attributes and the CLI command to set them.

Table 7 Keychain Mapping

Definition	CLI
The key identifier expressed as an integer (0...63)	config>system>security>keychain>direction>bi>entry config>system>security>keychain>direction>uni>receive>entry config>system>security>keychain>direction>uni>send>entry

Table 7 Keychain Mapping (Continued)

Definition	CLI
Authentication algorithm to use with key[i]	config>system>security>keychain>direction>bi>entry with algorithm <i>algorithm</i> parameter. config>system>security>keychain>direction>uni>receive>entry with algorithm <i>algorithm</i> parameter. config>system>security>keychain>direction>uni>send>entry with algorithm <i>algorithm</i> parameter.
Shared secret to use with key[i].	config>system>security>keychain>direction>uni>receive>entry with shared secret parameter config>system>security>keychain>direction>uni>send>entry with shared secret parameter config>system>security>keychain>direction>bi>entry with shared secret parameter
A vector that determines whether the key[i] is to be used to generate MACs for inbound segments, outbound segments, or both.	config>system>security>keychain>direction
Start time from which key[i] can be used.	config>system>security>keychain>direction>bi>entry>begin-time config>system>security>keychain>direction>uni>send>entry >begin-time
End time after which key[i] cannot be used by sending TCPs.	Inferred by the begin-time of the next key (youngest key rule).
Start time from which key[i] can be used.	config>system>security>keychain>direction>bi>entry>begin-time config>system>security>keychain>direction>bi>entry>tolerance config>system>security>keychain>direction>uni>receive>entry >begin-time config>system>security>keychain>direction>uni>receive>entry >tolerance
End time after which key[i] cannot be used	config>system>security>keychain>direction>uni>receive>entry>end-time

The following table details which authentication algorithm can be used in association with specific routing protocols.

[Table 8](#) shows the mapping between these attributes and the CLI command to set them.

Table 8 Security Algorithm Support Per Protocol

Protocol	Clear Text	MD5	HMAC-MD5	HMAC-SHA-1-96	HMAC-SHA-1	HMAC-SHA-256	AES-128-CMAC-96
OSPF	Yes	Yes	No	Yes	Yes	Yes	No
IS-IS	Yes	No	Yes	No	Yes	Yes	No
RSVP	Yes	No	Yes	No	Yes	No	No
BGP	No	Yes	No	Yes	No	No	Yes
LDP	No	Yes	No	Yes	No	No	Yes

2.7.11 gRPC Authentication

gRPC communication between the client and server must be authenticated and encrypted. There are two types of authentication:

- Authentication via session credentials — Session credentials operate similarly to device authentication, ensuring that the device is allowed in the network and is authorized by the provider. This type of authentication is performed using PKI and X.509.3 certificates. gRPC uses TLS for session authentication.

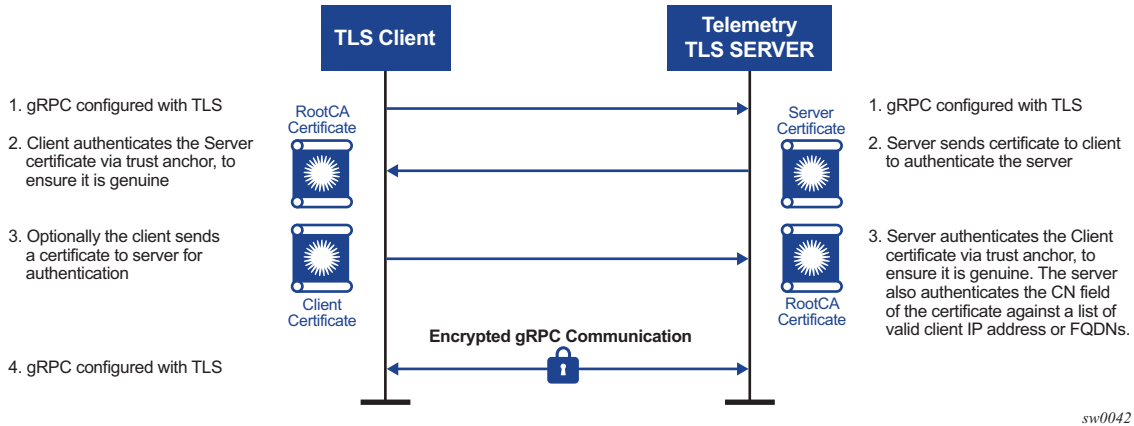
SR OS supports TLS servers for gRPC.

- Authentication using channel credentials — Channel credentials use a user name and password that are entered at the gRPC client terminal to authenticate gRPC packets using an AAA method.

Session authentication provides proof that the client and server are authorized devices and that they belong to the provider. After authentication, the session becomes encrypted using TLS, and gRPC PDUs are transmitted between the client and server.

[Figure 10](#) shows a basic session authentication using TLS.

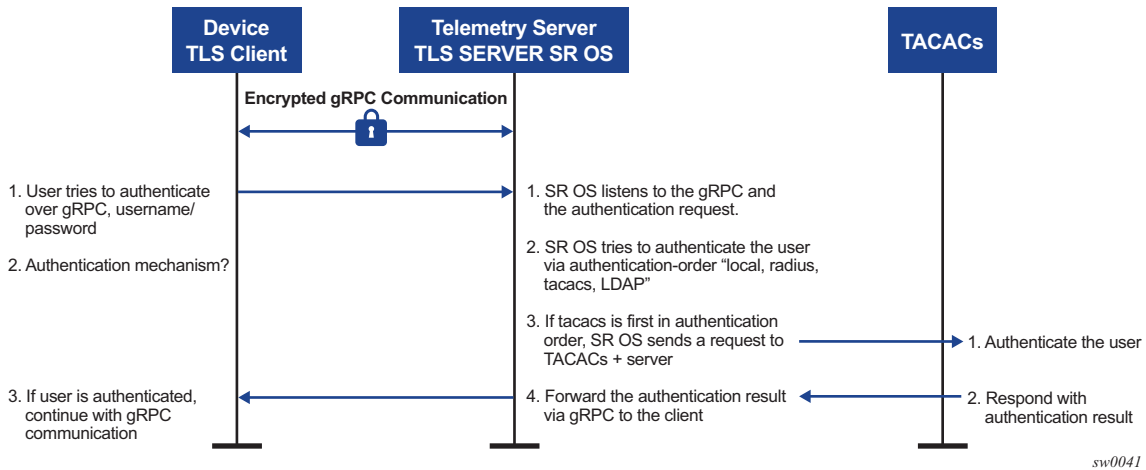
Figure 10 Session Authentication Using TLS



Channel credentials use username and password authentication. Each gRPC channel packet can contain a username and a password. Authentication is done through standard SR OS authentication order and mechanisms. All current authentication methods, including local and AAA servers, are applicable to gRPC channels. In addition, all authentication orders currently used by Telnet or SSH are compatible with gMI Call authentication.

Figure 11 shows a basic gMI Call authentication using SR OS.

Figure 11 gMI Call Authentication Using SR OS



gRPC channel packets contain the username and password in clear text, and are only encrypted using TLS. If a TLS server profile is assigned to the gRPC session, all PDUs between the server and client are encrypted. If TLS becomes operationally down, no gRPC PDUs are transmitted in clear text.

SR OS relies on existing authentication mechanisms for gRPC channels, including:

- AAA servers and local authentication orders configured using the **config>system>security>password>authentication-order** command
- password complexity rules
- requiring the user to be configured as part of gRPC access by using the **config>system>security>user>access>grpc** command
- disconnecting the gRPC session by using the **admin>disconnect gMI** command



Note: gRPC is not affected by password aging.

Security profiles can authorize bulk get, set, and subscribe gRPC commands that are received by the server. Profiles can be configured to permit or deny specific gRPC commands; for example, a profile for one user can authorize get and set commands, while a profile for another user can authorize get commands only.

2.8 Configuration Notes

This section describes security configuration caveats.

2.8.1 General

- If a RADIUS or a TACACS+ server is not configured, then password, profiles, and user access information must be configured on each router in the domain.
- If a RADIUS authorization is enabled, then VSAs must be configured on the RADIUS server.

2.9 Configuring Security with CLI

This section provides information to configure security using the command line interface.

Topics in this section include:

- [Setting Up Security Attributes](#)
 - [Configuring Authentication](#)
 - [Configuring Authorization](#)
 - [Configuring Accounting](#)
- [Security Configurations](#)
- [Configuration Tasks](#)
- [Security Configuration Procedures](#)
 - [Configuring Management Access Filters](#)
 - [Configuring IP CPM Filters Policy](#)
 - [Configuring MAC CPM Filters](#)
 - [Configuring IPv6 CPM Filters](#)
 - [Configuring CPM Queues](#)
 - [Configuring Profiles](#)
 - [Configuring Users](#)
 - [Copying and Overwriting Users and Profiles](#)
- [RADIUS Configurations](#)
 - [Configuring RADIUS Authentication](#)
 - [Configuring RADIUS Authorization](#)
 - [Configuring RADIUS Accounting](#)
- [Configuring 802.1x RADIUS Policies](#)
- [TACACS+ Configurations](#)
 - [Enabling TACACS+ Authentication](#)
 - [Configuring TACACS+ Authorization](#)
 - [Configuring TACACS+ Accounting](#)
 - [Enabling SSH](#)
- [LDAP Configurations](#)
- [Configuring Login Controls](#)

2.10 Setting Up Security Attributes

2.10.1 Configuring Authentication

Refer to the following sections to configure authentication:

- Local authentication

- [Configuring Profiles](#)
- [Configuring Users](#)

- RADIUS authentication (only)

By default, authentication is enabled locally. Perform the following tasks to configure security on each participating router:

- [Configuring Profiles](#)
- [Configuring RADIUS Authentication](#)
- [Configuring Users](#)

- RADIUS authentication

To implement only RADIUS authentication, *with* authorization, perform the following tasks on each participating router:

- [Configuring RADIUS Authentication](#)
- [Configuring RADIUS Authorization](#)

- TACACS+ authentication

To implement only TACACS+ authentication, perform the following tasks on each participating router:

- [Configuring Profiles](#)
- [Configuring Users](#)
- [Enabling TACACS+ Authentication](#)

- LDAP authentication

To implement only LDAP authentication, perform the following tasks on each participating router:

- [Configuring LDAP Authentication](#)

2.10.2 Configuring Authorization

Refer to the following sections to configure authorization.

- Local authorization

For local authorization, configure these tasks on each participating router:

- [Configuring Profiles](#)
- [Configuring Users](#)

- RADIUS authorization (only)

For RADIUS authorization (without authentication), configure these tasks on each participating router:

- [Configuring RADIUS Authorization](#)
- [Configuring Profiles](#)

For RADIUS authorization, VSAs must be configured on the RADIUS server. See [Vendor-Specific Attributes \(VSAs\)](#).

- RADIUS authorization

For RADIUS authorization (with authentication), configure these tasks on each participating router:

- [Configuring RADIUS Authorization](#)
For RADIUS authorization, VSAs must be configured on the RADIUS server. See [Vendor-Specific Attributes \(VSAs\)](#).
- [Configuring RADIUS Authentication](#)
- [Configuring Profiles](#)

- TACACS+ authorization (only)

For TACACS+ authorization (without authentication), configure these tasks on each participating router:

- [Configuring TACACS+ Authorization](#)

- TACACS+ authorization

For TACACS+ authorization (with authentication), configure these tasks on each participating router:

- [Enabling TACACS+ Authentication](#)
- [Configuring TACACS+ Authorization](#)

2.10.3 Configuring Accounting

Refer to the following sections to configure accounting.

- Local accounting is not implemented. For information about configuring accounting policies, refer to [Configuring Logging with CLI](#).
- [Configuring RADIUS Accounting](#)
- [Configuring TACACS+ Accounting](#)

2.11 Security Configurations

This section provides information to configure security and configuration examples of configuration tasks.

To implement security features, configure the following components:

- Management access filters and CPM filters
- Profiles
- User access parameters
- Password management parameters
- Enable RADIUS, TACACS+, and/or LDAP
 - One to five RADIUS, TACACS+, and/or LDAP servers
 - RADIUS, TACACS+, and/or LDAP parameters

2.12 Configuration Tasks

This section provides a brief overview of the tasks that must be performed to configure security and provides the CLI commands. [Table 9](#) depicts the capabilities of authentication, authorization, and accounting configurations. For example, authentication can be enabled locally and on RADIUS, TACACS+, and LDAP servers. Authorization can be executed locally, on a RADIUS server, or on a TACACS+ server. Accounting can be performed on a RADIUS or TACACS+ server.

Table 9 Security Configuration Requirements

Authentication	Authorization	Accounting
Local	Local	None
RADIUS	Local and RADIUS	RADIUS
TACACS+	Local	TACACS+

Table 9 Security Configuration Requirements (Continued)

Authentication	Authorization	Accounting
LDAP	None	None

2.13 Security Configuration Procedures

- [Configuring Management Access Filters](#)
- [Configuring IP CPM Filters Policy](#)
- [Configuring MAC CPM Filters](#)
- [Configuring CPM Queues](#)
- [Configuring Profiles](#)
- [Configuring Users](#)
- [Copying and Overwriting Users and Profiles](#)
- [Enabling SSH](#)

2.13.1 Configuring Management Access Filters

Creating and implementing management access filters is optional. Management access filters are software-based filters that control all traffic going in to the CPM, including all routing protocols. They apply to packets from all ports. The filters can be used to restrict management of the router by other nodes outside either specific (sub)networks or through designated ports. By default, there are no filters associated with security options. The management access filter and entries must be explicitly created on each router. These filters also apply to the management Ethernet port.

The OS implementation exits the filter when the first match is found and execute the actions according to the specified action. For this reason, entries must be sequenced correctly from most to least explicit. When both **mac-filter** and **ip-filter/ipv6-filter** are to be applied to a given traffic, **mac-filter** is applied first.

An entry may not have any match criteria defined (in which case, everything matches) but must have at least an action keyword specified CPM to be considered active complete. Entries without the action keyword are considered incomplete and will be rendered inactive. Management Access Filter must have at least one active entry defined for the filter to be active.

The following CLI commands are an example of how to configure a management access filter on the 7450 ESS. This example only accepts packets matching the criteria specified in entries 1 and 2. Non-matching packets are denied.

The following is an example of a management access filter configuration that accepts packets matching the criteria specified in IP, IPv6 and MAC entries. Non-matching packets are denied for IPv4 filter and permitted for IPv6 and MAC filters.

```
*A:Dut-C>config>system>security>mgmt-access-filter# info
-----
ip-filter
  default-action deny
  entry 10
    description "Accept SSH from mgmnt subnet"
    src-ip 192.168.5.0/26
    protocol tcp
    dst-port 22 65535
    action permit
  exit
exit
ipv6-filter
  default-action permit
  entry 10
    src-ip 3FFE::1:1/128
    next-header rsvp
    log
    action deny
  exit
exit
mac-filter
  default-action permit
  entry 12
    match frame-type ethernet_II
      svc-id 1
      src-mac 00:01:01:01:01:01 ff:ff:ff:ff:ff:ff
    exit
  action permit
  exit
exit
-----
*A:Dut-C>config>system>security>mgmt-access-filter#
```

2.13.2 Configuring IP CPM Filters Policy

The following displays a CPM filter configuration example:

```
*A:Dut-C>config>sys>security>cpm-filter# info
ip-filter
  shutdown
  entry 100 create
    action queue 50
    log 110
    match protocol icmp
      fragment true
      icmp-type dest-unreachable
      icmp-code host-unreachable
      multiple-option false
      option-present true
      src-ip 192.100.2.0/24
    exit
  exit
exit
ipv6-filter
```

```

        shutdown
        entry 30 create
        action drop
        log 190
        match next-header tcp
            dscp ef
            dst-ip 3FFE::2:2/128
            src-port 100 100
            tcp-syn true
            tcp-ack false
            flow-label 10
        exit
    exit
exit
    mac-filter
    shutdown
    entry 40 create
    action accept
    log 101
    match frame-type ethernet_II
        svc-id 12
        dst-mac 00:03:03:03:01:01 ff:ff:ff:ff:ff:ff
        etype 0x8902
        cfm-opcode gt 100
    exit
    exit
    exit
*A:Dut-C>config>sys>security>cpm-filter#

```

2.13.3 Configuring MAC CPM Filters

CPM filters and queues control all traffic going in to the CPM, including all routing protocols. They apply to packets from all network and access ports, but not to packets from a management Ethernet port. CPM packet filtering and queuing is performed by network processor hardware using no resources on the main CPUs. CPM filters and queues are not configurable on one-slot chassis.

The following displays a MAC CPM filter configuration example:

```

*A:ALA-49>config>sys>sec>cpm>mac-filter# info
-----
    entry 10 create
        description "MAC-CPM-Filter 10.10.10.100 #007"
        match
        exit
        log 101
        action drop
    exit
    entry 20 create
        description "MAC-CPM-Filter 10.10.10.100 #008"
        match
        exit
        log 101

```

```

        action drop
        exit
        no shutdown
-----
*A:ALA-49>config>sys>sec>cpm>mac-filter#

```

2.13.4 Configuring IPv6 CPM Filters

The following example displays an IPv6 CPM filter configuration:

```

A:ALA-48>config>sys>sec>cpm>ipv6-filter# info
entry 10 create
  description "IPv6 CPM Filter"
  log 101
  match next-header igp
    dst-ip 1000:1:1:1:1:1:1:1/112
    src-ip 2000:1::1/96
    flow-label 5000
  exit
exit
entry 20 create
  description "CPM-Filter 10.4.101.2 #201"
  log 101
  match next-header tcp
    dscp af11
    dst-ip 3FEE:12E1:2AC1:EA32::/64
    src-ip 3FEE:1FE1:2AC1:EA32::/64
    flow-label 5050
  exit
exit
no shutdown
A:ALA-48>config>sys>sec>cpm>ipv6-filter#

```

2.13.5 Configuring CPM Queues

CPM queues can be used to provide rate limit capabilities for traffic destined to CPM as described in an earlier section of this document.

The following example displays a CPM queue configuration:

```

A:ALA-987>config>sys>security>cpm-queue# info
-----
queue 33 create
exit
queue 101 create
  cbs 5
  mbs 5
  rate 5 cir 5
exit
queue 102 create

```

```
    cbs 5
    mbs 5
    rate 5 cir 5
exit
queue 103 create
    cbs 5
    mbs 5
    rate 5 cir 5
exit
queue 104 create
    cbs 5
    mbs 5
    rate 5 cir 5
```

```
-----
A:ALA-987>config>sys>security>cpm-queue#
```

2.13.6 IPsec Certificates Parameters

The following is an example to importing a certificate from a pem format:

```
*A:SR-7/Dut-A# admin certificate import type cert input cf3:/pre-import/R1-0cert.pem
output R1-0cert.der format pem
```

The following is an example for exporting a certificate to pem format:

```
*A:SR-7/Dut-A# admin certificate export type cert input R1-0cert.der output cf3:/
R1-0cert.pem format pem
```

The following displays an example of profile output:

```
*A:SR-7/Dut-A>config>system>security>pki# info
-----
      ca-profile "Root" create
      description "Root CA"
      cert-file "R1-0cert.der"
      crl-file "R1-0crl.der"
      no shutdown
      exit
-----
*A:SR-7/Dut-A>config>system>security>pki#
```

The following displays an example of an ike-policy with cert-auth output:

```
*A:SR-7/Dut-A>config>ipsec>ike-policy# info
-----
      ike-version 2
      auth-method cert-auth
      own-auth-method psk
-----
```

The following displays an example of a static lan-to-lan configuration using cert-auth:

```
...
interface "VPRN1" tunnel create
  sap tunnel-1.private:1 create
  ipsec-tunnel "Sanity-1" create
  security-policy 1
  local-gateway-address 30.1.1.13 peer 50.1.1.15 delivery-service 300
  dynamic-keying
  ike-policy 1
  pre-shared-key "Sanity-1"
  transform 1
  cert
    trust-anchor "R1-0"
    cert "M2cert.der"
    key "M2key.der"
```

```

        exit
    exit
no shutdown
    exit
exit
exit

```

2.13.7 Configuring Profiles

Profiles are used to deny or permit access to a hierarchical branch or specific commands. Profiles are referenced in a user configuration. A maximum of sixteen user profiles can be defined. A user can participate in up to sixteen profiles. Depending on the authorization requirements, passwords are configured locally or on the RADIUS server.

The following example displays a user profile output:

```

A:ALA-1>config>system>security# info
-----
...
    profile "ghost"
        default-action permit-all
    entry 1
        match "configure"
        action permit
    exit
    entry 2
        match "show"
    exit
    entry 3
        match "exit"
    exit
    exit
...
-----
A:ALA-1>config>system>security#

```

2.13.7.1 Parameters

Matching in authorization profiles allows the use of parameters and optional parameters. A set of angle brackets <...> indicates matching on a parameter and/or optional parameter.

The following rules govern parameter matching in the CLI:

Rule 1

Any parameter and/or optional parameter can be present in the match string.

Rule 2

When a parameter and/or optional parameter is present in the user-profile match string, all parameters or optional parameters to its left must also be stated/present.

Rule 3

The user can either specifically state or completely omit unnamed parameters in the match string, as required. However, all unnamed parameter in the CLI command must be present in the match string when matching on an unnamed parameter is used.

For example, consider the **OSPF** command:

```
*A:SwSim14# configure router ospf
- no ospf [<ospf-instance>]
- ospf [<ospf-instance>] [<router-id>]

<ospf-instance>      : [0..31]
<router-id>         : <ip-address>
```

In this case, the user can match on OSPF to allow or deny the command per user-profile, as follows:

```
Match "configure router ospf" action deny
```

Or the user can decide to only allow a certain OSPF instance for a user, as follows:

```
Match "configure router ospf <ospf-instance-value> <router-id-value>"
```



Note: Although the user's matching is based on <ospf-instance-value> that is "an unnamed value", all other unnamed values in the **OSPF** command (such as the <router-id-value>) must also be present in the match string.

Rule 4

When multiple unnamed parameters are present in the match string, the parameters must be provided in the correct order as described in the command **help** to generate the correct match behavior. For example, using the order of parameters described in the **OSPF** command usage in Rule 3 above, use the following statement for a user-profile match:

```
match "configure router ospf <ospf-instance-value> <router-id-value>
```

The desired match behavior might not be achieved if the unnamed parameters <ospf-instance-value> and <router-id-value> are out of order with respect to the help screen.

The following displays a parameter matching output:

```
config>system>security>profile# info
  entry 10
    match "show router <22> route-table "
    action permit
  exit
  entry 20
    match "configure service vprn <22>"
    action read-only
  exit
  entry 30
    match "show service id <22>"
    action permit
  exit
  entry 40
    match "configure router interface <system>"
    action deny
  exit
```

2.13.7.2 Wildcards

In addition, parameter configuration is facilitated by the availability of wildcards (.) in the OAM subtree and for commands such as “ping”, “trace-route” and “m-trace”. For example, consider the following command:

```
ping <ip-address> router 10
```

Instead of listing all the permitted IP addresses in the policy, as shown in the following example,

```
Match ping <10.0.0.1> router <10>
Action permit
Match ping <10.0.0.2> router <10>
Action permit
```

The wildcard<ip-address> parameter allows a a simpler search criterion. In the following example, the use of <.*> wildcard enables you to ping any address in the router 10 context, that is, any address in VRF 10:

```
Match ping <.*> router <10>
Action permit
```



Note: While wildcards are available and allowed for all parameters in the OAM subtree, Nokia recommends that caution is exercised when using wildcards and limit their use to commands such as 'ping', 'trace-route' and 'm-trace'. The use of wildcards in certain formats may be a security concern and result in making the IP addresses in the VRF, including the base routing table, unreachable. Or it could allow the customer to ping any IP address in the VRF, including the base routing table. This may be a potential security concern and should be avoided.

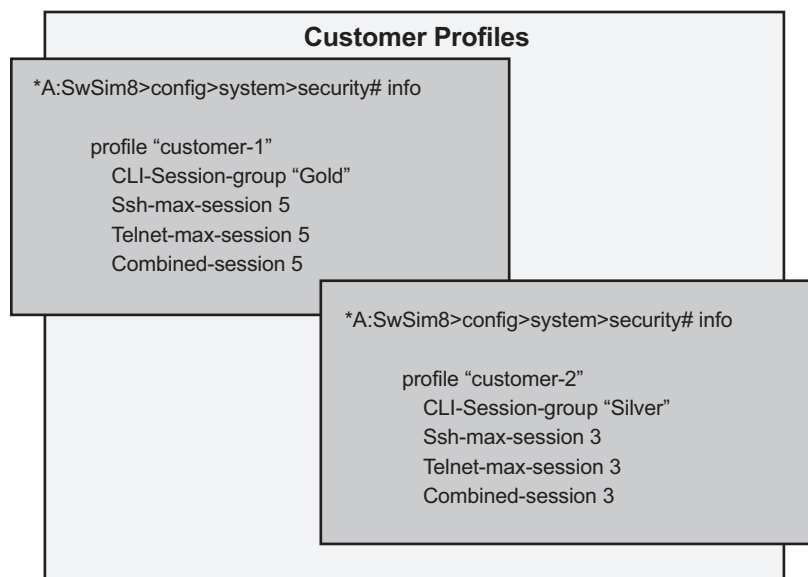
For example, the following usage is not advised:

```
Match ping <.*> router <.*>
Action permit
```

2.13.7.3 CLI Session Resource Management

SR OS has the capability to manage telnet/ssh sessions per user and at a higher level per system. At the system level, the user can configure a **cli-session-group** for different customer priorities. The **cli-session-group** is a container that sets the maximum number of CLI sessions for a class of customers, with a unique session limit for each customer. For example, as depicted in [Figure 12](#), “Gold” category customers can have a **cli-session-group** that allows them more telnet/ssh sessions compared to “Silver” category customers.

Figure 12 cli-session-group for Customer Classes

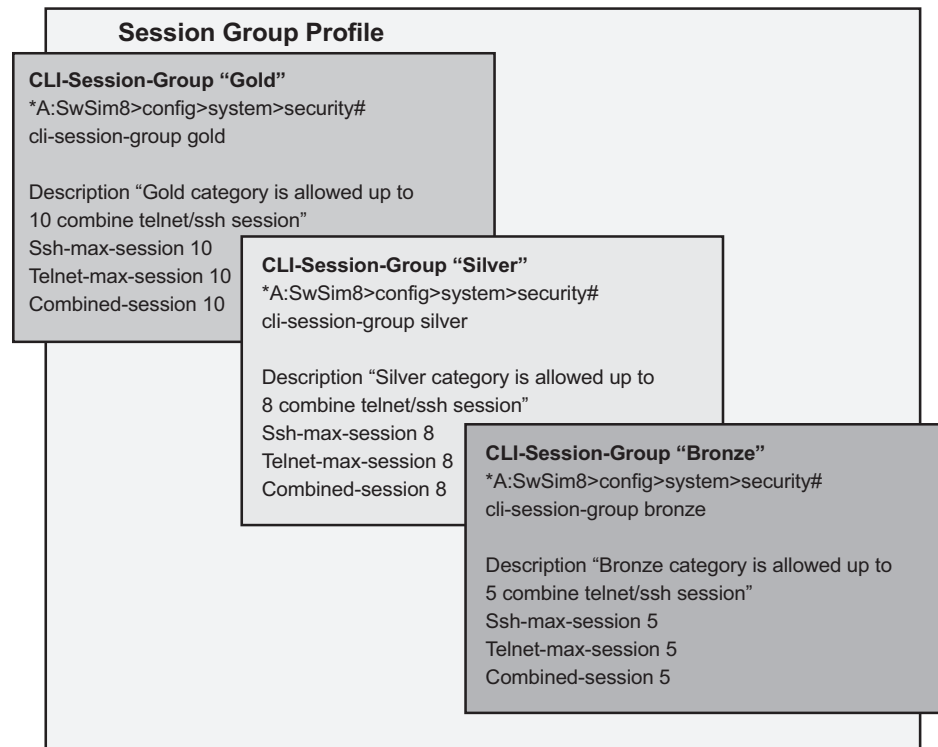


al_0777

The configured **cli-session-group** can be assigned to user-profiles. At the user profile level, each profile can be configured with its own max ssh/telnet session and it will be policed/restricted by the higher order **cli-session-group** that is assigned to it.

As depicted in [Figure 13](#), the final picture is a hierarchical configuration with top-level cli-session-groups that control each customer's total number of ssh/telnet sessions and the user-profile for each user for that customer.

Figure 13 Hierarchy of cli-session-group Profiles



al_0776

Every profile will subtract one from it's corresponding **max-session** when a TELNET or SSH session is established in the following cases:

- where multiple profiles are configured under a user
- where multiple profiles arrive from different AAA servers (Local Profile, Radius Profile or Tacacs Profile)

The first profile to run out of corresponding **max-session** will limit future TELNET or SSH sessions. In other words, while each profile for the user can have its independent **max-session**, only the lowest one will be honored. If the profile with the lowest **max-session** is removed, the next lower profile **max-session** will be honored and so on. All profiles for a user are updated when a TELNET or SSH session is established.

For information about login control, see [Configuring Login Controls](#).

Use the following CLI commands to configure CLI session resources:

```
CLI Syntax:  config>system>security>profile <name>
                [no] ssh-max-sessions session-limit
                [no] telnet-max-sessions session-limit
                [no] combined-max-session session-limit
                [no] cli-session-group session-group-name
```

2.13.8 Configuring Users

Configure access parameters for individual users. For user, define the login name for the user and, optionally, information that identifies the user.

The following displays a user configuration example:

```
A:ALA-1>config>system>security# info
-----
...
        user "49ers"
            password "$2y$10$pFoehOg/tCbBMPDJ/
kqpu.8af0AcVGy2xsR7WFqyn5fVTnwRzGmOK"
            access console ftp snmp
            restricted-to-home
            console
                member "default"
                member "ghost"
            exit
        exit
...
-----
A:ALA-1>config>system>security#
```

2.13.9 Configuring Keychains

The following displays a keychain configuration.

```
A:ALA-1>config>system>security# info
```

```

-----
...
        keychain "abc"
        direction
            bi
                entry 1 key "ZcvSElJzJx/wBZ9biCtOVQJ9YZQvVU.S" hash2 alg
algorithm aes-128-cmac-96
                begin-time 2006/12/18 22:55:20
                exit
            exit
        exit
    exit
    keychain "basasd"
    direction
        uni
            receive
                entry 1 key "Ee7xdKlYO2D0m7v3IJv/84LIu96R2fZh" hash2
algorithm aes-128-cmac-96
                tolerance forever
                exit
            exit
        exit
    exit
    exit
    exit
...
-----
A:ALA-1>config>system>security#

```

2.13.10 Copying and Overwriting Users and Profiles

You can copy a profile or user. You can copy a profile or user or overwrite an existing profile or user. The **overwrite** option must be specified or an error occurs if the destination profile or user name already exists.

2.13.10.1 User

CLI Syntax: `config>system>security# copy {user source-user | profile source-profile} to destination [overwrite]`

Example:

```

config>system>security# copy user testuser to testuserA
MINOR: CLI User "testuserA" already exists - use
  overwrite flag.
config>system>security#
config>system>security# copy user testuser to testuserA
  overwrite
config>system>security#

```

The following output displays the copied user configurations:

```
A:ALA-12>config>system>security# info
-----
...
        user "testuser"
            password "$2y$10$pFoehOg/tCbBMPDJ/
kqpu.8af0AoVGY2xsR7WFqyn5fVTnwRzGmOK"
            access snmp
            snmp
                authentication hash md5 e14672e71d3e96e7a1e19472527ee969 privacy
        none
            group "testgroup"
        exit
    exit
    user "testuserA"
        password ""
        access snmp
        console
        new-password-at-login
    exit
    snmp
        authentication hash md5 e14672e71d3e96e7a1e19472527ee969 privacy
    none
        group "testgroup"
    exit
    exit
...
-----
A:ALA-12>config>system>security# info
```



Note: The cannot-change-password flag is not replicated when a copy user command is performed. A new-password-at-login flag is created instead.

```
A:ALA-12>config>system>security>user# info
-----
password "$2y$10$pFoehOg/tCbBMPDJ/kqpu.8af0AoVGY2xsR7WFqyn5fVTnwRzGmOK"
access snmp
console
cannot-change-password
exit
snmp
authentication hash md5 e14672e71d3e96e7a1e19472527ee969 privacy none
group "testgroup"
exit
-----
A:ALA-12>config>system>security>user# exit
A:ALA-12>config>system>security# user testuserA
A:ALA-12>config>system>security>user# info
-----
password ""
access snmp
console
new-password-at-login
exit
snmp
authentication hash md5 e14672e71d3e96e7a1e19472527ee969 privacy none
```

```

group "testgroup"
exit
-----
A:ALA-12>config>system>security>user#

```

2.13.10.2 Profile

CLI Syntax: `config>system>security# copy {user source-user | profile source-profile} to destination [overwrite]`

Example: `config>system>security# copy profile default to testuser`

The following output displays the copied profiles:

```

A:ALA-49>config>system>security# info
-----
...
A:ALA-49>config>system>security# info detail
-----
...
        profile "default"
            default-action none
            entry 10
                no description
                match "exec"
                action permit
            exit
            entry 20
                no description
                match "exit"
                action permit
            exit
            entry 30
                no description
                match "help"
                action permit
            exit
            entry 40
                no description
                match "logout"
                action permit
            exit
            entry 50
                no description
                match "password"
                action permit
            exit
            entry 60
                no description
                match "show config"
                action deny
            exit
            entry 70
                no description

```



```
        match "show"
        action permit
    exit
    entry 80
        no description
        match "enable-admin"
        action permit
    exit
exit
profile "testuser"
    default-action none
    entry 10
        no description
        match "exec"
        action permit
    exit
    entry 20
        no description
        match "exit"
        action permit
    exit
    entry 30
        no description
        match "help"
        action permit
    exit
    entry 40
        no description
        match "logout"
        action permit
    exit
    entry 50
        no description
        match "password"
        action permit
    exit
    entry 60
        no description
        match "show config"
        action deny
    exit
    entry 70
        no description
        match "show"
        action permit
    exit
    entry 80
        no description
        match "enable-admin"
        action permit
    exit
exit
profile "administrative"
    default-action permit-all exit
...
-----
A:ALA-12>config>system>security#
```

2.14 RADIUS Configurations

- [Configuring RADIUS Authentication](#)
- [Configuring RADIUS Authorization](#)
- [Configuring RADIUS Accounting](#)
- [Configuring 802.1x RADIUS Policies](#)

2.14.1 Configuring RADIUS Authentication

RADIUS is disabled by default and must be explicitly enabled. The mandatory commands to enable RADIUS on the local router are **radius** and server *server-index* address *ip-address* secret *key*.

Also, the system IP address must be configured in order for the RADIUS client to work. See [Configuring a System Interface of the Router Configuration Guide](#).

The other commands are optional. The server command adds a RADIUS server and configures the RADIUS server's IP address, index, and key values. The index determines the sequence in which the servers are queried for authentication requests.

On the local router, use the following CLI commands to configure RADIUS authentication:

```
CLI Syntax:  config>system>security
                radius
                port port
                retry count
                server server-index address ip-address secret key
                timeout seconds
                no shutdown
```

The following displays a RADIUS authentication configuration example:

```
A:ALA-1>config>system>security# info
-----
                retry 5
                timeout 5
                server 1 address 10.10.10.103 secret "test1"
                server 2 address 10.10.0.1 secret "test2"
                server 3 address 10.10.0.2 secret "test3"
                server 4 address 10.10.0.3 secret "test4"
                ...
-----
A:ALA-1>config>system>security#
```

2.14.2 Configuring RADIUS Authorization

In order for RADIUS authorization to function, RADIUS authentication *must* be enabled first. See [Configuring RADIUS Authentication](#).

In addition to the local configuration requirements, VSAs must be configured on the RADIUS server. See [Vendor-Specific Attributes \(VSAs\)](#).

On the local router, use the following CLI commands to configure RADIUS authorization:

CLI Syntax:

```
config>system>security
radius
    authorization
```

The following displays a RADIUS authorization configuration example:

```
A:ALA-1>config>system>security# info
-----
...
    radius
        authorization
        retry 5
        timeout 5
        server 1 address 10.10.10.103 secret "test1"
        server 2 address 10.10.0.1 secret "test2"
        server 3 address 10.10.0.2 secret "test3"
        server 4 address 10.10.0.3 secret "test4"
    exit
...
-----
A:ALA-1>config>system>security#
```

2.14.3 Configuring RADIUS Accounting

On the local router, use the following CLI commands to configure RADIUS accounting:

CLI Syntax:

```
config>system>security
radius
    accounting
```

The following displays RADIUS accounting configuration example:

```
A:ALA-1>config>system>security# info
-----
...
    radius
```

```

shutdown
authorization
accounting
retry 5
timeout 5
server 1 address 10.10.10.103 secret "test1"
server 2 address 10.10.0.1 secret "test2"
server 3 address 10.10.0.2 secret "test3"
server 4 address 10.10.0.3 secret "test4"
exit
...
-----
A:ALA-1>config>system>security#

```

2.15 Configuring 802.1x RADIUS Policies

Use the following CLI commands to configure generic authentication parameters for clients using 802.1x EAPOL. Additional parameters are configured per Ethernet port. Refer to the Interface Configuration Guide.

To configure generic parameters for 802.1x authentication, enter the following CLI syntax.

CLI Syntax:

```

config>system>security
dot1x
    radius-plcy policy-name
        server server-index address ip-address secret
            key [port port]
        source-address ip-address
        no shutdown

```

The following displays a 802.1x configuration example:

```

A:ALA-1>config>system>security# info
-----
dot1x
    radius-plcy "dot1x_plcy" create
        server 1 address 1.1.1.1 port 65535 secret "a"
        server 2 address 1.1.1.2 port 6555 secret "a"
        source-address 1.1.1.255
        no shutdown
...
-----
A:ALA-1>config>system#

```

2.16 TACACS+ Configurations

- [Enabling TACACS+ Authentication](#)
- [Configuring TACACS+ Authorization](#)
- [Configuring TACACS+ Accounting](#)
- [Enabling SSH](#)

2.16.1 Enabling TACACS+ Authentication

To use TACACS+ authentication on the router, configure one or more TACACS+ servers on the network.

Use the following CLI commands to configure profiles:

CLI Syntax:

```
config>system>security
tacplus
    server server-index address ip-address secret
        key
    timeout seconds
    no shutdown
```

The following displays a TACACS+ authentication configuration example:

```
A:ALA-1>config>system>security>tacplus# info
-----
timeout 5
server 1 address 10.10.0.5 secret "test1"
server 2 address 10.10.0.6 secret "test2"
server 3 address 10.10.0.7 secret "test3"
server 4 address 10.10.0.8 secret "test4"
server 5 address 10.10.0.9 secret "test5"
-----
A:ALA-1>config>system>security>tacplus#
```

2.16.2 Configuring TACACS+ Authorization

In order for TACACS+ authorization to function, TACACS+ authentication *must* be enabled first. See [Enabling TACACS+ Authentication](#).

On the local router, use the following CLI commands to configure RADIUS authorization:

```
CLI Syntax:  config>system>security
               tacplus
                 authorization
                 no shutdown
```

The following displays a TACACS+ authorization configuration example:

```
A:ALA-1>config>system>security>tacplus# info
-----
      authorization
      timeout 5
      server 1 address 10.10.0.5 secret "test1"
      server 2 address 10.10.0.6 secret "test2"
      server 3 address 10.10.0.7 secret "test3"
      server 4 address 10.10.0.8 secret "test4"
      server 5 address 10.10.0.9 secret "test5"
-----
A:ALA-1>config>system>security>tacplus#
```

2.16.3 Configuring TACACS+ Accounting

On the local router, use the following CLI commands to configure TACACS+ accounting:

```
CLI Syntax:  config>system>security
               tacplus
                 accounting
```

The following displays a TACACS+ accounting configuration example:

```
A:ALA-1>config>system>security>tacplus# info
-----
      accounting
      authorization
      timeout 5
      server 1 address 10.10.0.5 secret "test1"
      server 2 address 10.10.0.6 secret "test2"
      server 3 address 10.10.0.7 secret "test3"
      server 4 address 10.10.0.8 secret "test4"
      server 5 address 10.10.0.9 secret "test5"
-----
A:ALA-1>config>system>security>tacplus#
```

2.16.4 Enabling SSH

Use the SSH command to configure the SSH server as SSH1, SSH2 or both. The default is SSH2 (SSH version 2). This command should only be enabled or disabled when the SSH server is disabled. This setting should not be changed while the SSH server is running since the actual change only takes place after SSH is disabled or enabled.

CLI Syntax:

```
config>system>security
ssh
    preserve-key
    no server-shutdown
    version ssh-version
```

The following displays a SSH server configuration as both SSH and SSH2 using a host-key:

```
A:sim1>config>system>security>ssh# info
-----
                preserve-key
                version 1-2
-----
A:sim1>config>system>security>ssh#
```

2.17 LDAP Configurations

- [Configuring LDAP Authentication](#)
- [Configuring Redundant Servers](#)
- [Enabling SSH](#)

2.17.1 Configuring LDAP Authentication

LDAP is disabled by default and must be explicitly enabled. To use LDAP authentication on the router, configure one or more LDAP servers on the network.

TLS certificates and clients must also be configured. Refer to the “TLS” section of the *System Management Guide* for more information about configuring TLS.

Use the following CLI commands to configure LDAP:

CLI Syntax:

```
config>system>security>ldap
```

```

[no] public-key-authentication
[no] retry
[no] server
[no] shutdown
[no] timeout
[no] use-default-template

config>system>security>password
authentication-order [method] exit-on-reject

config>system>security>ldap
public-key-authentication
server server-index create
address ip-address port port
bind-authentication root-dn [password
password] [hash | hash2]
ldap-server server-name
search base-dn
tls-profile tls-profile-name
no shutdown

exit
no shutdown

```

The following displays an LDAP authentication configuration example:

```

A:SwSim14>config>system>security>ldap#
-----
[no] public-key-authentication
[no] retry
[no] server
[no] shutdown
[no] timeout
[no] use-default-template
-----
*A:SwSim14>config>system>security>password#
-----
authentication-order [local | radius | tacplus | ldap] exit-on-reject
-----
*A:SwSim14>config>system>security>ldap# info
-----
public-key-authentication
server 1 create
address 1.1.1.1
bind-
authentication "cn=administrator,cn=users,dc=nacblr2,dc=example,dc=com" pass
word"
ldap-server "active-server"
search "dc=sns,dc=example,dc=com"
tls-profile "server-1-profile"
no shutdown
exit
no shutdown
-----
*A:SwSim8>config>system>security>tls# info

```



```
-----  
client-tls-profile "server-1-profile" create  
  cipher-list "to-active-server"  
  trust-anchor-profile "server-1-ca"  
no shutdown  
exit
```

2.17.2 Configuring Redundant Servers

Up to five redundant LDAP servers can be configured. The following examples show configuration of two servers, Server-1 and Server-5.

Configuration of Server-1:

```
A*:SwSim14>config>system>security>ldap# info  
public-key-authentication  
server 1 create  
  address 1.1.1.1  
  ldap-server "active-server"  
  tls-profile "server-1-profile"  
  
A*:SwSim14>config>system>security>tls# info  
client-tls-profile "server-1-profile" create  
  cert-profile "client-cert-profile"  
  cipher-list "to-active-server"  
  trust-anchor-profile "server-1-ca"  
no shutdown  
exit
```

Configuration of Server-5 (backup):

```
A*:SwSim14>config>system>security>ldap# info  
public-key-authentication  
server 5 create  
  address 5.5.5.1  
  ldap-server "backup-server-5"  
  tls-profile "server-5-profile"  
  
A*:SwSim14>config>system>security>tls# info  
client-tls-profile "server-5-profile" create  
  cert-profile "client-cert-profile"  
  cipher-list "to-backup-server-5"  
  trust-anchor-profile "server-5-ca"  
no shutdown  
exit
```

2.17.3 Enabling SSH

SSH must be enabled to use LDAP authentication. See the [2.16.4](#) subsection in the [2.16](#) section for more information.

2.18 Configuring Login Controls

Configure login control parameters for console, Telnet, and FTP sessions.

The following displays a login control configuration example:

```
A:ALA-1>config>system# info
-----
...
    login-control
        ftp
            inbound-max-sessions 5
        exit
        telnet
            inbound-max-sessions 7
            outbound-max-sessions 2
        exit
        idle-timeout 1440
        pre-login-message "Property of Service Routing Inc. Unauthorized access
                           prohibited."
        motd text "Notice to all users: Software upgrade scheduled 3/2 1:00 AM"
    exit
no exponential-backoff
...
-----
A:ALA-1>config>system#
```

2.19 Security Configuration Command Reference

2.19.1 Command Hierarchies

- Security Commands
 - LLDP Commands
 - Management Access Filter Commands
 - CLI Script Authorization Commands
 - CPM Filter Commands
 - CPM Queue Commands
 - CPU Protection Commands
 - Distributed CPU Protection Commands
 - Extracted Protocol Traffic Priority Commands
 - Security Password Commands
 - Public Key Infrastructure (PKI) Commands
 - Profile Commands
 - CLI Session Commands
 - RADIUS Commands
 - SSH Commands
 - TACPLUS Commands
 - LDAP Commands
 - User Commands
 - User Template Commands
 - Dot1x Commands
 - Keychain Commands
 - TTL Security Commands
 - gRPC Commands
- Login Control Commands

2.19.1.1 Security Commands

```
config
  — system
    — security
```

- **copy** {*user source-user* | *profile source-profile*} to *destination* [**overwrite**]
- [no] **ftp-server**
- **hash-control** [**read-version** {1 | 2 | all}] [**write-version** {1 | 2}]
- **no hash-control**
- [no] **per-peer-queuing**
- **source-address**
 - **application** *app* [*ip-int-name* | *ip-address*]
 - **no application** *app*
 - **application6** *app* *ipv6-address*
 - **no application6**
- [no] **telnet-server**
- [no] **telnet6-server**
- **vprn-network-exceptions** *number seconds*

2.19.1.1.1 LLDP Commands

- ```
configure
 — system
 — lldp
 — message-fast-tx time
 — no message-fast-tx
 — message-fast-tx-init count
 — no message-fast-tx-init
 — notification-interval time
 — no notification-interval
 — reinit-delay time
 — no reinit-delay
 — tx-credit-max count
 — no tx-credit-max
 — tx-hold-multiplier multiplier
 — no tx-hold-multiplier
 — tx-interval interval
 — no tx-interval
```

### 2.19.1.1.2 Management Access Filter Commands

- ```
config
  — system
    — security
      — [no] management-access-filter
        — [no] ip-filter
          — default-action {permit | deny}
          — [no] entry entry-id
            — action {permit | deny | deny-host-unreachable}
            — no action
            — description description-string
            — no description
            — dst-port value [mask]
            — no dst-port
```

- [no] **log**
- **protocol** *protocol-id*
- **no protocol**
- **router** {*router-instance*}
- **no router**
- **src-ip** {*ip-prefix/mask* | *ip-prefix netmask*}
- **no src-ip**
- **src-port** {*port-id* | **cpm** | **lag** *lag-id* }
- **no src-port**
- **src-port** *old-entry-number new-entry-number*
- **renum** *old-entry-number new-entry-number*
- [no] **shutdown**
- [no] **ipv6-filter**
- **default-action** {**permit** | **deny** | **deny-host-unreachable**}
- [no] **entry** *entry-id*
 - **action** {**permit** | **deny** | **deny-host-unreachable**}
 - **no action**
 - **description** *description-string*
 - **no description**
 - **dst-port** *value* [*mask*]
 - **no dst-port**
 - **flow-label** *value*
 - **no flow-label**
 - [no] **log**
 - **next-header** *next-header*
 - **no next-header**
 - **router** {*router-instance*}
 - **no router**
 - **src-ip** {*ip-prefix/mask* | *ip-prefix netmask*}
 - **no src-ip**
 - **src-port** {*port-id* | **cpm** | **lag** *lag-id* }
 - **no src-port**
- **renum** *old-entry-number new-entry-number*
- [no] **shutdown**
- [no] **mac-filter**
- **default-action** {**permit** | **deny**}
- [no] **entry** *entry-id*
 - **action** {**permit** | **deny** | **deny-host-unreachable**}
 - **no action**
 - **description** *description-string*
 - **no description**
 - [no] **log**
 - **match** **frame-type** *frame-type*
 - **no match**
 - **cfm-opcode** {**lt** | **gt** | **eq**} *opcode*
 - **cfm-opcode** **range** *start end*
 - **no cfm-opcode**
 - **dot1p** *dot1p-value* [*dot1p-mask*]
 - **dsap** *dsap-value* [*dsap-mask*]
 - **dst-mac** *ieee-address* [*ieee-address-mask*]
 - **no dst-mac**
 - **etype** *0x0600..0xfff*
 - **no etype**
 - **snap-oui** {**zero** | **non-zero**}

- **snap-pid** *snap-pid*
- **no snap-pid**
- **src-mac** *ieee-address* [*ieee-address-mask*]
- **no src-mac**
- **ssap** *ssap-value* [*ssap-mask*]
- **no ssap**
- **svc-id** *service-id*
- **no svc-id**
- **renum** *old-entry-number* *new-entry-number*
- **[no] shutdown**

2.19.1.1.3 CLI Script Authorization Commands

- ```

config
 — system
 — security
 — cli-script
 — authorization
 — cron
 — cli-user user-name
 — no cli-user
 — vsd
 — cli-user user-name
 — no cli-user
 — event-handler
 — cli-user user-name
 — no cli-user

```

### 2.19.1.1.4 CPM Filter Commands

- ```

config
  — system
    — security
      — [no] cpm-filter
        — default-action {accept | drop}
        — [no] ip-filter
          — [no] entry entry-id
            — action [accept | drop | queue queue-id ]
            — no action
            — description description-string
            — no description
            — log log-id
            — no log
            — match [protocol protocol-id]
            — no match
              — dscp dscp-name
              — no dscp
              — dst-ip {ip-address/mask | ip-address netmask |
                ip-prefix-list prefix-list-name}
  
```

- **no dst-ip**
- **dst-port** [tcp/udp port-number] [mask]
- **no dst-port**
- **fragment** {true | false}
- **no fragment**
- **icmp-code** icmp-code
- **no icmp-code**
- **icmp-type** icmp-type
- **no icmp-type**
- **ip-option** [ip-option-value] [ip-option-mask]
- **no ip-option**
- **multiple-option** {true | false}
- **no multiple-option**
- **option-present** {true | false}
- **no option-present**
- **port** tcp/udp port-number [mask]
- **port port-list** port-list-name
- **port range** tcp/udp port-number tcp/udp port-number
- **no port**
- **router**
- **src-ip** {ip-address/mask | ip-address netmask | ip-prefix-list prefix-list-name}
- **no src-ip**
- **src-port**[src-port-number] [mask]
- **no src-port**
- **tcp-ack** {true | false}
- **no tcp-ack**
- **tcp-syn** {true | false}
- **no tcp-syn**
- **renum** old-entry-id new-entry-id
- **[no] shutdown**
- **[no] ipv6-filter**
 - **[no] entry** entry-id
 - **action** [accept | drop | queue queue-id]
 - **no action**
 - **description** description-string
 - **no description**
 - **log** log-id
 - **no log**
 - **match** [next-header next-header]
 - **no match**
 - **dscp** dscp-name
 - **no dscp**
 - **dst-ip** ipv6-address/prefix-length
 - **dst-ip ipv6-prefix-list** ipv6-prefix-list-name
 - **no dst-ip**
 - **dst-port** [tcp/udp port-number] [mask]
 - **dst-port port-list** port-list-name
 - **dst-port range** tcp/udp port-number tcp/udp port-number
 - **no dst-port**
 - **flow-label** value
 - **no flow-label**

- **fragment** {true | false}
- **no fragment**
- **hop-by-hop-opt** {true | false}
- **no hop-by-hop-opt**
- **icmp-code** *icmp-code*
- **no icmp-code**
- **icmp-type** *icmp-type*
- **no icmp-type**
- **port** *tcp/udp port-number [mask]*
- **port port-list** *port-list-name*
- **port range** *start end*
- **no port**
- **router service-name** *service-name*
- **router** *router-instance*
- **no router**
- **src-ip** [*ipv6-address/prefix-length*] [**ipv6-prefix-list** *ipv6-prefix-list-name*]
- **no src-ip**
- **src-port** [*src-port-number*] [*mask*]
- **no src-port**
- **tcp-ack** {true | false}
- **no tcp-ack**
- **tcp-syn** {true | false}
- **no tcp-syn**
- **renum** *old-entry-id new-entry-id*
- [no] **shutdown**
- [no] **mac-filter**
 - [no] **entry** *entry-id*
 - **action** [accept | drop | queue *queue-id*]
 - **no action**
 - **description** *description-string*
 - **no description**
 - **log** *log-id*
 - **no log**
 - **match** [*frame-type frame-type*]
 - **no match**
 - **cfm-opcode** {lt | gt | eq} *opcode*
 - **cfm-opcode range** *start end*
 - **no cfm-opcode**
 - **dsap** *dsap-value [dsap-mask]*
 - **dst-mac** *ieee-address [ieee-address-mask]*
 - **no dst-mac**
 - **etype** *0x0600..0xfff*
 - **no etype**
 - **src-mac** *ieee-address [ieee-address-mask]*
 - **no src-mac**
 - **ssap** *ssap-value [ssap-mask]*
 - **no ssap**
 - **svc-id** *service-id*
 - **no svc-id**
 - **renum** *old-entry-number new-entry-number*
 - [no] **shutdown**

2.19.1.1.5 CPM Queue Commands

```

config
  — system
    — security
      — [no] cpm-queue
        — [no] queue queue-id
          — cbs cbs
          — no cbs
          — mbs mbs
          — no mbs
          — rate rate [cir cir]
          — no rate

```

2.19.1.1.6 CPU Protection Commands

```

config
  — system
    — security
      — cpu-protection
        — ip-src-monitoring
          — included-protocols
            — [no] dhcp
            — [no] gtp
            — [no] icmp
            — [no] igmp
        — link-specific-rate packet-rate-limit
        — no link-specific-rate
        — policy cpu-protection-policy-id [create]
        — no policy cpu-protection-policy-id
          — [no] alarm
          — description description-string
          — no description
          — eth-cfm entry entry levels levels opcodes opcodes rate
             packet-rate-limit
          — no eth-cfm
          — out-profile-rate packet-rate-limit [log-events]
          — no out-profile-rate
          — overall-rate packet-rate-limit
          — no overall-rate
          — per-source-rate packet-rate-limit
          — no per-source-rate
        — port-overall-rate packet-rate-limit [action-low-priority]
        — no port-overall-rate
        — [no] protocol-protection [allow-sham-links][block-pim-tunneled]

```

Refer to the OS Services Guide and the Multi-Service ISA Guide for command, syntax, and usage information about applying CPU Protection policies to interfaces.

CPU protection policies are applied by default (and customer policies can be applied) to a variety of entities including interfaces and SAPs. Refer to the appropriate guides for command syntax and usage for applying CPU protection policies. Examples of entities that can have CPU protection policies applied to them include:

```
config>router>if>cpu-protection policy-id
```

```
config>service>epipe>sap>cpu-protection policy-id [mac-monitoring] | [eth-cfm-monitoring [aggregate]][car]]
```

```
config>service>epipe>spoke-sdp>cpu-protection policy-id [mac-monitoring] | [eth-cfm-monitoring [aggregate]][car]]
```

```
config>service>ies>if>cpu-protection policy-id
```

```
config>service>ies>if>sap>cpu-protection policy-id [mac-monitoring] | [eth-cfm-monitoring [aggregate]][car]]
```

```
config>service>template>vpls-sap-template>cpu-protection policy-id [mac-monitoring] | [eth-cfm-monitoring [aggregate]][car]]
```

```
config>service>vpls>sap>cpu-protection policy-id [mac-monitoring] | [eth-cfm-monitoring [aggregate]][car]]
```

```
config>service>vpls>video-interface>cpu-protection policy-id
```

```
config>service>vprn>if>cpu-protection policy-id
```

```
config>service>vprn >if>sap>cpu-protection policy-id [mac-monitoring] | [eth-cfm-monitoring [aggregate]][car]]
```

```
config>service>vprn>nw-if>cpu-protection policy-id
```

```
config>service>vprn>sub-if>grp-if>sap>cpu-protection policy-id [mac-monitoring] | [eth-cfm-monitoring [aggregate]][car]]
```

```
config>subscr-mgmt>msap-policy>cpu-protection policy-id [mac-monitoring]
```

2.19.1.1.7 Distributed CPU Protection Commands

```
config
  — system
    — security
      — dist-cpu-protection
        — policy policy-name [create]
        — no policy
          — description description-string
```

- **no description**
- **[no] local-monitoring-policer** *policer-name* [**create**]
 - **[no] description** *description-string*
 - **exceed-action** {**discard** | **low-priority** | **none**}
 - **rate** {**packets** {*ppi* | **max**} **within seconds** [**initial-delay** *packets*] | **kbps** {*kilobits-per-second* | **max**} [**mbs size**] [**bytes** | **kilobytes**]}
 - **no rate**
 - **[no] log-events** [**verbose**]
- **protocol** *name* [**create**]
- **no protocol** *name*
 - **dynamic-parameters**
 - **detection-time** *seconds*
 - **exceed-action** {**discard** [**hold-down seconds**] | **low-priority** [**hold-down seconds**] | **none**}
 - **log-events** [**verbose**]
 - **no log-events**
 - **rate** {**packets** {*ppi* | **max**} **within seconds** [**initial-delay** *packets*] | **kbps** {*kilobits-per-second* | **max**} [**mbs size**] [**bytes** | **kilobytes**]}
 - **enforcement** {**static** *policer-name* | **dynamic** {*mon-policer-name* | **local-mon-bypass** }}
- **static-policer** *policer-name* [**create**]
- **no static-policer** *policer-name*
 - **description** *description-string*
 - **no description**
 - **detection-time** *seconds*
 - **no detection-time**
 - **exceed-action** {**discard** [**hold-down seconds**] | **low-priority** [**hold-down seconds**] | **none**}
 - **log-events** [**verbose**]
 - **no log-events**
 - **rate** {**packets** {*ppi* | **max**} **within seconds** [**initial-delay** *packets*] | **kbps** {*kilobits-per-second* | **max**} [**mbs size**] [**bytes** | **kilobytes**]}
 - **no rate**

- ```

config
 — card
 — fp
 — dist-cpu-protection
 — [no] dynamic-enforcement-policer-pool number-of-policers

```

### 2.19.1.1.8 Extracted Protocol Traffic Priority Commands

- ```

config
  — card
    — fp
      — init-extract-prio-mode {uniform | I3-classify}
  
```

2.19.1.1.9 Security Password Commands

```

config
  — system
    — security
      — password
        — admin-password password [hash | hash2]
        — no admin-password
        — aging days
        — no aging
        — attempts count [time minutes1] [lockout minutes2]
        — no attempts
        — authentication-order [method-1] [method-2] [method-3] [method-4]
          [exit-on-reject]
        — no authentication-order
        — complexity-rules
          — [no] allow-user-name
          — credits [lowercase credits] [uppercase credits] [numeric
            credits] [special-character credits]
          — no credits
          — minimum-classes minimum
          — no minimum-classes
          — minimum-length length
          — no minimum-length
          — repeated-characters count
          — no repeated-characters
          — required [lowercase count] [uppercase count] [numeric count]
            [special-character count]
          — no required
        — dynsvc-password password [hash | hash2]
        — no dynsvc-password
        — enable-admin-control
        — tacplus-map-to-priv-lvl admin-priv-lvl
        — no tacplus-map-to-priv-lvl
        — health-check [interval interval]
        — no health-check
        — history size
        — no history
        — minimum-age [days days] [hrs hours] [min minutes] [sec seconds]
        — no minimum-age
        — minimum-change distance
        — no minimum-change

```

2.19.1.1.10 Public Key Infrastructure (PKI) Commands

The following commands apply only to the 7450 ESS and 7750 SR:

```

config
  — system
    — security
      — pki

```

- **ca-profile** *name* [**create**]
- **no ca-profile** *name*
 - **cert-file** *filename*
 - **no cert-file**
 - **cmpv2**
 - [no] **accept-unprotected-errormsg**
 - [no] **accept-unprotected-pkiconf**
 - **http-response-timeout** *timeout*
 - **no http-response-timeout**
 - **key-list**
 - **key** *password* [**hash** | **hash2**] **reference** *reference-number*
 - **no key** **reference** *reference-number*
 - **response-signing-cert** *filename*
 - **no response-signing-cert**
 - [no] **same-recipnonce-for-pollreq**
 - **url** *url-string* [**service-id** *service-id*]
 - **no url**
- **certificate-display-format** {**ascii** | **utf8**}
- **certificate-expiration-warning** *hours* [**repeat** *repeat-hours*]
- **no certificate-expiration-warning**
- **common-name-list** *name*
 - [no] **cn** *index* **type** *type* **value** *common-name-value*
- **crl-expiration-warning** *hours* [**repeat** *repeat-hours*]
- **no crl-expiration-warning**
- **maximum-cert-chain-depth** *level*
- **no maximum-cert-chain-depth**



Note: For information about CMPv6 admin certificate commands listed in the following tree, see the Multiservice Integrated Service Adapter Guide.

admin

- **certificate**
 - **clear-ocsp-cache** [*entry-id*]
 - **crl-update** **ca** *ca-profile-name*
 - **display** **type** {**cert** | **key** | **crl** | **cert-request**} *url-string* **format** {**pkcs10** | **pkcs12** | **pkcs7-der** | **pkcs7-pem** | **pem** | **der**} [**password** [32 chars max]]
 - **export** **type** {**cert** | **key** | **crl**} **input** *filename* **output** *url-string* **format** *output-format* [**password** [32 chars max]] [**pkey** *filename*]
 - **gen-keypair** *url-string* [**size** {**512** | **1024** | **2048**}] [**type** {**rsa** | **dsa**}]
 - **gen-local-cert-req** **keypair** *url-string* **subject-dn** *subject-dn* [**domain-name** [255 chars max]] [**ip-addr** *ip-address*] **file** *url-string* [**hash-alg** *hash-algorithm*]
 - **import** **type** {**cert** | **key** | **crl**} **input** *url-string* **output** *filename* **format** *input-format* [**password** [32 chars max]]
 - **reload** **type** {**cert** | **key** | **cert-key-pair**} *filename* [**key-file** *filename*]
 - **secure-nd-export**
 - **secure-nd-import** **input** *url-string* **format** *input-format* [**password** *password*] [**key-rollover**]

2.19.1.1.11 Profile Commands

```

config
  — system
    — security
      — [no] profile user-profile-name
        — default-action {deny-all | permit-all | none | read-only-all}
        — [no] entry entry-id
          — action {deny | permit | read-only}
          — description description-string
          — no description
          — security command-string
          — no security
        — renum old-entry-number new-entry-number
        — ssh-max-sessions session-limit
        — no ssh-max-sessions
        — telnet-max-sessions session-limit
        — no telnet-max-sessions
        — combined-max-sessions session-limit
        — no combined-max-sessions

```

2.19.1.1.12 CLI Session Commands

```

config
  — system
    — security
      — cli-session-group session-group-name [create]
        — ssh-max-sessions session-limit
        — no ssh-max-sessions
        — telnet-max-sessions session-limit
        — no telnet-max-sessions
        — combined-max-sessions
        — no combined-max-sessions

```

2.19.1.1.13 RADIUS Commands

```

config
  — system
    — security
      — [no] radius
        — access-algorithm {direct | round-robin}
        — no access-algorithm
        — [no] accounting
        — accounting-port port
        — no accounting-port
        — [no] authorization
        — [no] interactive-authentication
        — port port
        — no port

```

- **retry** *count*
- **no retry**
- **server** *server-index address ip-address secret key [hash | hash2]*
- **no server** *server-index*
- **[no] shutdown**
- **timeout** *seconds*
- **no timeout**
- **[no] use-default-template**

2.19.1.1.14 SSH Commands

- ```

config
 — system
 — security
 — ssh
 — client-cipher-list protocol-version version
 — cipher index name cipher-name
 — no cipher index
 — [no] preserve-key
 — server-cipher-list protocol-version version
 — cipher index name cipher-name
 — no cipher index
 — [no] server-shutdown
 — [no] version SSH-version

```

### 2.19.1.1.15 TACPLUS Commands

- ```

config
  — system
    — security
      — [no] tacplus
        — accounting [record-type {start-stop | stop-only}]
        — no accounting
        — [no] authorization [use-priv-lvl]
        — [no] interactive-authentication
        — [no] priv-lvl-map
          — priv-lvl priv-lvl user-profile-name
          — no priv-lvl priv-lvl
        — server server-index address ip-address secret key [hash | hash2]
          [port port]
        — no server server-index
        — [no] shutdown
        — timeout seconds
        — no timeout
        — [no] use-default-template

```

2.19.1.1.16 LDAP Commands

```

config
  -- system
    -- security
      -- [no] ldap
        -- [no] public-key-authentication
        -- retry value
        -- no retry
        -- server server-index [create]
          -- address ip-address [port port]
          -- no address
          -- bind-authentication root-dn [password password] [hash |
            hash2]
          -- no bind-authentication
          -- ldap-server server-name
          -- no ldap-server
          -- search base-dn
          -- no search
          -- [no] shutdown
          -- tls-profile tls-profile-name
          -- no tls-profile
        -- no server
        -- [no] shutdown
        -- timeout seconds
        -- no timeout
        -- [no] use-default-template

```

2.19.1.1.17 User Commands

```

config
  -- system
    -- security
      -- [no] user user-name
        -- [no] access [ftp] [snmp] [console] [li] [netconf] [grpc]
        -- console
          -- [no] cannot-change-password
          -- login-exec url-prefix::source-url
          -- no login-exec
          -- member user-profile-name [user-profile-name...(up to 8 max)]
          -- no member user-profile-name
          -- [no] new-password-at-login
        -- home-directory url-prefix [directory] [directory/directory...]
        -- no home-directory
        -- password [password]
        -- public-keys
          -- ecdsa
            -- [no] ecdsa-key key-id [create]
              -- description description-string
              -- no description
              -- key-value public-key-value

```


- no **key-value**
- **rsa**
 - [no] **rsa-key** *key-id* [**create**]
 - **description** *description-string*
 - no **description**
 - **key-value** *public-key-value*
 - no **key-value**
- [no] **restricted-to-home**
- **snmp**
 - **authentication** {[none] | [[hash] {md5 *key-1* | sha *key-1* }
privacy {none | des-key | aes-128-cfb-key *key-2*}]}
 - **group** *group-name*
 - no **group**

2.19.1.1.18 User Template Commands

- ```
config
 — system
 — security
 — user-template {tacplus_default | radius_default | ldap-default}
 — [no] access [ftp] [console] [grpc]
 — console
 — login-exec url-prefix:source-url
 — no login-exec
 — home-directory url-prefix [directory] [directory/directory..]
 — no home-directory
 — profile user-profile-name
 — no profile
 — [no] restricted-to-home
```

### 2.19.1.1.19 Dot1x Commands

- ```
config
  — system
    — security
      — dot1x
        — radius-plcy name
          — retry count
          — no retry
          — server server-index address ip-address secret key ] [hash |  
hash2] [auth-port auth-port] [acct-port acct-port] [type  
server-type]
          — source-address ip-address
          — [no] shutdown
          — timeout seconds
          — no timeout
        — [no] shutdown
```

2.19.1.1.20 Keychain Commands

```

config
  — system
    — security
      — [no] keychain keychain-name
        — description description-string
        — no description
        — direction {uni | bi}
          — bi
            — entry entry-id key [authentication-key | hash-key | hash2-key] [hash | hash2] algorithm algorithm
            — no entry
              — begin-time [date] [hours-minutes] [UTC]
              — begin-time {now | forever}
              — no begin-time
              — option {basic | isis-enhanced}
              — no option
              — [no] shutdown
              — tolerance [seconds | forever]
              — no tolerance
          — uni
            — receive
              — entry entry-id key [authentication-key | hash-key | hash2-key] [hash | hash2] algorithm algorithm
              — no entry entry-id
                — begin-time [date] [hours-minutes] [UTC]
                — begin-time {now | forever}
                — no begin-time
                — end-time [date] [hours-minutes] [UTC]
                — end-time {now | forever}
                — no end-time
                — [no] shutdown
                — tolerance [seconds | forever]
                — no tolerance
            — send
              — entry entry-id key [authentication-key | hash-key | hash2-key] [hash | hash2] algorithm algorithm
              — no entry entry-id
                — begin-time [date] [hours-minutes] [UTC]
                — begin-time {now | forever}
                — no begin-time
                — [no] shutdown
          — [no] shutdown
      — tcp-option-number
        — receive option-number
        — no receive
        — send option-number
        — no send

```

2.19.1.1.21 TTL Security Commands

```
config
  — router
    — bgp
      — group
        — ttl-security min-ttl-value
        — neighbor
          — ttl-security min-ttl-value
```

```
config
  — router
    — ldp
      — tcp-session-parameters
        — peer-transport
          — ttl-security min-ttl-value
```

```
config
  — system
    — login-control
      — ssh
        — ttl-security
```

```
config
  — system
    — login-control
      — telnet
        — ttl-security
```

2.19.1.1.22 gRPC Commands

```
config
  — system
    — grpc
      — tls-server-profile name
      — no tls-server-profile
```

2.19.1.2 Login Control Commands

```
config
  — system
    — login-control
      — [no] exponential-backoff
      — ftp
        — inbound-max-sessions number-of-sessions
        — no inbound-max-sessions
      — idle-timeout {minutes | disable}
      — no idle-timeout
```

-
- **[no] login-banner**
 - **login-scripts**
 - **global** *file-url*
 - **no global**
 - **per-user** *user-directory file-url file-name file-name*
 - **no per-user**
 - **motd** {*url url-prefix: source-url | text motd-text-string*}
 - **no motd**
 - **pre-login-message** *login-text-string [name]*
 - **no pre-login-message**
 - **ssh**
 - **disable-graceful-shutdown**
 - **inbound-max-sessions**
 - **outbound-max-sessions**
 - **tty-security**
 - **telnet**
 - **enable-graceful-shutdown**
 - **inbound-max-sessions** *value*
 - **no inbound-max-sessions**
 - **outbound-max-sessions** *value*
 - **no outbound-max-sessions**
 - **tty-security**

2.19.2 Command Descriptions

This section provides the CLI command descriptions. Topics include:

- [General Security Commands](#)
- [LLDP Commands](#)
- [Login, Telnet, SSH and FTP Commands](#)
- [Management Access Filter Commands](#)
- [Password Commands](#)
- [Public Key Infrastructure \(PKI\) Commands](#)
- [Profile Management Commands](#)
- [User Management Commands](#)
- [CLI Session Management Commands](#)
- [RADIUS Client Commands](#)
- [TACACS+ Client Commands](#)
- [LDAP Client Commands](#)
- [Generic 802.1x COMMANDS](#)
- [Keychain Authentication](#)
- [CLI Script Commands](#)
- [CPM Filter Commands](#)
- [CPM Queue Commands](#)
- [TTL Security Commands](#)
- [CPU Protection Commands](#)
- [Distributed CPU Protection Commands](#)
- [Extracted Protocol Traffic Priority Commands](#)

2.19.2.1 General Security Commands

description

Syntax	description <i>description-string</i> no description
Context	config>system>security>mgmt-access-filter>ip-filter>entry config>system>security>mgmt-access-filter>ipv6-filter>entry config>sys>sec>cpm>ip-filter>entry

```

config>sys>sec>cpm>ipv6-filter>entry
config>sys>sec>cpm>mac-filter>entry
config>sys>security>keychain>direction>bi>entry
config>system>security>keychain>direction>uni>receive>entry
config>system>security>keychain>direction>uni>send>entry
config>system>security>pki>ca-profile
config>sys>security>cpu-protection>policy
config>system>security>mgmt-access-filter>mac-filter>entry
config>system>security>cpm-filter>mac-filter>entry
config>system>security>user>public-keys>ecdsa>ecdsa-key
config>system>security>user>public-keys>rsa>rsa-key

```

Description This command creates a text description stored in the configuration file for a configuration context.
This command associates a text string with a configuration context to help identify the context in the configuration file.

The **no** form of the command removes the string.

Default No description associated with the configuration context.

Parameters *string* — The description character string. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

shutdown

Syntax **[no] shutdown**

Context

```

config>system>security>mgmt-access-filter>ip-filter
config>system>security>mgmt-access-filter>ipv6-filter
config>sys>sec>cpm>ip-filter
config>system>security>keychain>direction>bi>entry
config>system>security>keychain>direction>uni>receive>entry
config>system>security>keychain>direction>uni>send>entry
config>system>security>pki>ca-profile
config>sys>sec>cpm>ipv6-filter
config>sys>sec>cpm>mac-filter>entry

```

Description The **shutdown** command administratively disables the entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics. Many entities must be explicitly enabled using the **no shutdown** command. The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they may be deleted.

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

Default no shutdown

security

Syntax	security
Context	config>system
Description	This command creates the context to configure security settings. Security commands manage user profiles and user membership. Security commands also manage user login registrations.

ftp-server

Syntax	[no] ftp-server
Context	config>system>security
Description	This command enables FTP servers running on the system. FTP servers are disabled by default. At system startup, only SSH server are enabled. The no form of the command disables FTP servers running on the system.

hash-control

Syntax	hash-control [read-version {1 2 all}] [write-version {1 2}] no hash-control
Context	config>system>security
Description	Whenever the user executes a save or info command, the system will encrypt all passwords, MD5 keys, etc., for security reasons. At present, two algorithms exist. The first algorithm is a simple, short key that can be copied and pasted in a different location when the user wants to configure the same password. However, because it is the same password and the hash key is limited to the password/key, even the casual observer will notice that it is the same key. The second algorithm is a more complex key, and cannot be copied and pasted in different locations in the configuration file. In this case, if the same key or password is used repeatedly in different contexts, each encrypted (hashed) version will be different.
Default	all — read-version set to accept both versions 1 and 2
Parameters	read-version {1 2 all} — When the read-version is configured as all , both versions 1 and 2 will be accepted by the system. Otherwise, only the selected version will be accepted when reading configuration or exec files. The presence of incorrect hash versions will abort the script/startup.

write-version {1 | 2} — Select the hash version that will be used the next time the configuration file is saved (or an info command is executed). Be careful to save the read and write version correctly, so that the file can be properly processed after the next reboot or exec.

per-peer-queuing

Syntax	[no] per-peer-queuing
Context	config>system>security
Description	<p>This command enables CPM hardware queuing per peer. This means that when a peering session is established, the router will automatically allocate a separate CPM hardware queue for that peer.</p> <p>The no form of the command disables CPM hardware queuing per peer.</p>
Default	per-peer-queuing

source-address

Syntax	source-address
Context	config>system>security
Description	<p>This command specifies the source address that should be used in all unsolicited packets sent by the application.</p> <p>This feature only applies on inband interfaces and does not apply on the out of band management interface. Packets going out the management interface will keep using that as source IP address. In other words, when the RADIUS server is reachable through both the management interface and a network interface, the management interface is used despite whatever is configured under the source-address statement.</p> <p>When a source address is specified for the ptp application, the port-based 1588 hardware timestamping assist function will be applied to PTP packets matching the IPv4 address of the router interface used to ingress the SR/ESS or IP address specified in this command. If the IP address is removed, then the port-based 1588 hardware timestamping assist function will only be applied to PTP packets matching the IPv4 address of the router interface.</p>

application

Syntax	application app [ip-int-name ip-address] no application app
Context	config>system>security>source-address

Description This command specifies the use of the source IP address specified by the **source-address** command.

Parameters *app* — Specify the application name.

Values cflowd, dns, ftp, ntp, ldap, ping, ptp, radius, sflow, snmptrap, sntp, ssh, syslog, tacplus, telnet, traceroute, mcreporter, icmp-error

ip-int-name | *ip-address* — Specifies the name of the IP interface or IP address. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

application6

Syntax **application6** *app* *ipv6-address*
no application6

Context config>system>security>source-address

Description This command specifies the application to use the source IPv6 address specified by the **source-address** command.

Parameters *app* — Specify the application name.

Values cflowd, dns, ftp, ldap, ntp, ping, radius, sflow, snmptrap, sntp, ssh, syslog, tacplus, telnet, traceroute, icmp6-error

ipv6-address — Specifies the IPv6 address.

telnet-server

Syntax **[no] telnet-server**

Context config>system>security

Description This command enables Telnet servers running on the system.

Telnet servers are off by default. At system startup, only SSH servers are enabled.

Telnet servers in networks limit a Telnet clients to three retries to login. The Telnet server disconnects the Telnet client session after three retries.

The **no** form of the command disables Telnet servers running on the system.

telnet6-server

Syntax **[no] telnet6-server**

Context config>system>security

- Description** This command enables Telnet IPv6 servers running on the system and only applies to the 7750 SR and 7950 XRS.
- Telnet servers are off by default. At system startup, only SSH server are enabled.
- The **no** form of the command disables Telnet IPv6 servers running on the system.

vprn-network-exceptions

- Syntax** `vprn-network-exceptions number seconds`
- Context** `config>system>security`
- Description** This command configures the rate to limit ICMP replies to packets with label TTL expiry received within all VPRN sentences in the system and from all network IP interfaces. This includes labeled user packets, ping and traceroute packets within VPRN.
- This feature currently also limits the same packets when received within the context of an LSP shortcut.
- This feature does not rate limit MPLS and service OAM packets (vprn-ping, vprn-trace, lsp-ping, lsp-trace, vccv-ping, and vccv-trace).
- The **no** form of the command disables the rate limiting of the reply to these packets.
- This feature only applies to the 7750 SR and 7950 XRS.
- Default** `no security vprn-network-exceptions`
- Parameters** *number* — 10 to 10,000
seconds — 1 to 60

2.19.2.2 LLDP Commands

lldp

- Syntax** `lldp`
- Context** `config>system`
- Description** This command enables the context to configure system-wide Link Layer Discovery Protocol parameters.

message-fast-tx

Syntax	message-fast-tx <i>time</i> no message-fast-tx
Context	config>system>lldp
Description	This command configures the duration of the fast transmission period.
Parameters	<i>time</i> — Specifies the fast transmission period in seconds. Values 1 to 3600 Default 1

message-fast-tx-init

Syntax	message-fast-tx-init <i>count</i> no message-fast-tx-init
Context	config>system>lldp
Description	This command configures the number of LLDPDUs to send during the fast transmission period.
Parameters	<i>count</i> — Specifies the number of LLDPDUs to send during the fast transmission period. Values 1 to 8 Default 4

notification-interval

Syntax	notification-interval <i>time</i> no notification-interval
Context	config>system>lldp
Description	This command configures the minimum time between change notifications.
Parameters	<i>time</i> — Specifies the minimum time, in seconds, between change notifications. Values 5 to 3600 Default 5

reinit-delay

Syntax **reinit-delay** *time*

no reinit-delay

Context	config>system>lldp
Description	This command configures the time before re-initializing LLDP on a port.
Parameters	<i>time</i> — Specifies the time, in seconds, before re-initializing LLDP on a port.
Values	1 to 10
Default	2

tx-credit-max

Syntax	tx-credit-max <i>count</i> no tx-credit-max
Context	config>system>lldp
Description	This command configures the maximum consecutive LLDPDUs transmitted.
Parameters	<i>count</i> — Specifies the maximum consecutive LLDPDUs transmitted.
Values	1 to 100
Default	5

tx-hold-multiplier

Syntax	tx-hold-multiplier <i>multiplier</i> no tx-hold-multiplier
Context	config>system>lldp
Description	This command configures the multiplier of the tx-interval.
Parameters	<i>multiplier</i> — Specifies the multiplier of the tx-interval.
Values	2 to 10
Default	4

tx-interval

Syntax	tx-interval <i>interval</i> no tx-interval
Context	config>system>lldp
Description	This command configures the LLDP transmit interval time.

Parameters *interval* — Specifies the LLDP transmit interval time.

Values 1 to 100

Default 5

2.19.2.3 Login, Telnet, SSH and FTP Commands

exponential-backoff

Syntax [no] exponential-backoff

Context config>system>login-control

Description This command enables the exponential-backoff of the login prompt. The exponential-backoff command is used to deter dictionary attacks, when a malicious user can gain access to the CLI by using a script to try **admin** with any conceivable password.

The **no** form of the command disables exponential-backoff.

Default no exponential-backoff

ftp

Syntax ftp

Context config>system>login-control

Description This command creates the context to configure FTP login control parameters.

idle-timeout

Syntax idle-timeout {minutes | disable}
no idle-timeout

Context config>system>login-control

Description This command configures the idle timeout for FTP, console, or Telnet sessions before the session is terminated by the system.

By default, an idle FTP, console, SSH or Telnet session times out after 30 minutes of inactivity. This timer can be set per session.

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

Default 30

- Parameters** *minutes* — The idle timeout in minutes. Allowed values are 1 to 1440. 0 implies the sessions never timeout.
- Values** 1 to 1440
- disable** — When the **disable** option is specified, a session will never timeout. To re-enable idle timeout, enter the command without the disable option.

inbound-max-sessions

- Syntax** **inbound-max-sessions** *value*
no inbound-max-sessions
- Context** config>system>login-control>ftp
- Description** This command configures the maximum number of concurrent inbound FTP sessions. This value is the combined total of inbound and outbound sessions. The **no** form of the command reverts to the default value.
- Default** 3
- Parameters** *value* — The maximum number of concurrent FTP sessions on the node.
- Values** 0 to 5

inbound-max-sessions

- Syntax** **inbound-max-sessions** *number-of-sessions*
no inbound-max-sessions
- Context** config>system>login-control>telnet
config>system>login-control>ssh
- Description** This parameter limits the number of inbound Telnet and SSH sessions. A maximum of 30 telnet and ssh connections can be established to the router. The local serial port cannot be disabled. Telnet and SSH maximum sessions can also use the combined total of both inbound sessions (SSH+Telnet). While it is acceptable to continue to internally limit the combined total of SSH and Telnet sessions to N, either SSH or Telnet sessions can use the inbound maximum sessions, if so required by the Operator. The **no** form of the command reverts to the default value.
- Default** 5

Parameters *number-of-sessions* — The maximum number of concurrent inbound Telnet sessions, expressed as an integer.

Values 0 to 50 (default = 5)
or 0 to N where N is the new total number of SSH+Telnet sessions if they are scaled

login-control

Syntax **login-control**

Context config>system

Description This command creates the context to configure the session control for console, Telnet and FTP.

login-banner

Syntax **[no] login-banner**

Context config>system>login-control

Description This command enables or disables the display of a login banner. The login banner contains the SR OS copyright and build date information for a console login attempt.

The **no** form of the command causes only the configured pre-login-message and a generic login prompt to display.

login-scripts

Syntax **login-scripts**

Context config>system>login-control

Description This command enables the context to configure CLI scripts that execute when a user (authenticated via any method including local user database, TACACS+, or RADIUS) first logs into a CLI session.

global

Syntax **global file-url**
no global

Context config>system>login-control>login-scripts

Description This command enables an operator to define a common CLI script that executes when any user logs into a CLI session. This login exec script is executed when any user (authenticated by any means including local user database, TACACS+, or RADIUS) opens a CLI session. This allows a user, for example, to define a common set of CLI aliases that are made available on the router for all users. This global login exec script is executed before any user-specific login exec files that may be configured.

This CLI script executes in the context of the user who opens the CLI session. Any commands in the script that the user is not authorized to execute will fail.

The **no** form of this command disables the execution of a global login-script.

Default no global

Parameters *file-url* — The path or directory name.

per-user

Syntax **per-user user-directory *dir-url* file-name *file-name***
no per-user

Context config>system>login-control>login-scripts

Description This command allows users to define their own login scripts that can be executed each time they first login to a CLI session. The command executes the script "*file-url / username / file-name*" when the user *username* logs into a CLI session (authenticated by any means including local user database, TACACS+, or RADIUS).

For example:

per-user user-directory "cf1:/local/users" file-name "login-script.txt"

would search for the following script when user "admin" logs in and authenticates via RADIUS:

cf1:/local/users/admin/login-script.txt

The per user login script is executed after any global script executes and before any login-exec script configured against a local user is executed. This allows users, for example, who are authenticated via TACACS+ or RADIUS to define their own login scripts.

This CLI script executes in the context of the user who opens the CLI session. Any commands in the script that the user is not authorized to execute will fail.

The **no** form of the command disables the execution of any per user login-scripts.

Default no per user

Parameters *dir-url* — The path or directory name.

file-name — The name of the file (located in the *dir-url* directory) including the extension.

motd

Syntax	motd {url <i>url-prefix: source-url</i> text <i>motd-text-string</i> } no motd
Context	config>system>login-control
Description	This command creates the message of the day displayed after a successful console login. Only one message can be configured. The no form of the command removes the message.
Default	no motd
Parameters	url <i>url-prefix: source-url</i> — When the message of the day is present as a text file, provide both url-prefix and the source-url of the file containing the message of the day. The URL prefix can be local or remote. text <i>motd-text-string</i> — The text of the message of the day. The <i>motd-text-string</i> must be enclosed in double quotes. Multiple text strings are not appended to one another. Some special characters can be used to format the message text. The \n character can be used to create multi-line messages. A \n in the message moves to the beginning of the next line by sending ASCII/UTF-8 chars 0xA (LF) and 0xD (CR) to the client terminal. An \r in the message sends the ASCII/UTF-8 char 0xD (CR) to the client terminal.

outbound-max-sessions

Syntax	outbound-max-sessions <i>value</i> no outbound-max-sessions
Context	config>system>login-control>telnet
Description	This parameter limits the number of outbound Telnet and SSH sessions. A maximum of 15 telnet and ssh connections can be established from the router. The local serial port cannot be disabled. The no form of the command reverts to the default value.
Default	5
Parameters	<i>value</i> — The maximum number of concurrent outbound Telnet sessions, expressed as an integer. Values 0 to 15

pre-login-message

Syntax	pre-login-message <i>login-text-string</i> [<i>name</i>] no pre-login-message
Context	config>system>login-control
Description	<p>This command creates a message displayed prior to console login attempts on the console via Telnet.</p> <p>Only one message can be configured. If multiple pre-login-messages are configured, the last message entered overwrites the previous entry.</p> <p>It is possible to add the name parameter to an existing message without affecting the current pre-login-message.</p> <p>The no form of the command removes the message.</p>
Default	no pre-login-message
Parameters	<p><i>login-text-string</i> — The string can be up to 900 characters. Any printable, 7-bit ASCII characters can be used. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes. Some special characters can be used to format the message text. The \n character can be used to create multi-line messages. A \n in the message moves to the beginning of the next line by sending ASCII/UTF-8 chars 0xA (LF) and 0xD (CR) to the client terminal. A \r in the message sends the ASCII/UTF-8 char 0xD (CR) to the client terminal.</p> <p>name — When the keyword <i>name</i> is defined, the configured system name is always displayed first in the login message. To remove the name from the login message, the message must be cleared and a new message entered without the name.</p>

ssh

Syntax	ssh
Context	config>system>login-control
Description	This command enables the context to configure the SSH parameters.

client-cipher-list protocol-version

Syntax	client-cipher-list protocol-version <i>version</i>
Context	config>system>security>ssh
Description	This command enables configuration the list of allowed ciphers by the SSH client.

- Parameters** *version* — Specifies the SSH version.
- Values** 1 — Specifies that the SSH server will only accept connections from clients that support SSH protocol version 1
2 — Specifies that the SSH server will accept connections from clients supporting either SSH protocol version 2

cipher

- Syntax** ***cipher index name cipher-name***
no cipher index
- Context** config>system>security>ssh>client-cipher-list
config>system>security>ssh>server-cipher-list
- Description** This command enables configuration of a cipher. Client-ciphers are used when the SR OS is acting as an SSH client. Server-ciphers are used when the SR OS is acting as an SSH server.
- Default** no cipher *index*
- Parameters** *index* — Specifies the index of the cipher in the list.
- Values** 1 to 255
- cipher-name* — Specifies the algorithm for performing encryption or decryption.
- Values** For SSHv1:
Client ciphers: des, 3des, blowfish
Server ciphers: 3des, blowfish
[Table 10](#) lists the default ciphers used for SSHv1:

Table 10 SSHv1 Default Ciphers

Cipher index value	Cipher name
10	3des
20	blowfish
30	des



Note: blowfish and des are not permitted in FIPS-140-2 mode.

Values For SSHv2:
 Client ciphers: 3des-cbc, blowfish-cbc, cast128-cbc, arcfour, aes128-cbc, aes192-cbc, aes256-cbc, rijndael-cbc, aes128-ctr, aes192-ctr, aes256-ctr
 Server ciphers: 3des-cbc, blowfish-cbc, cast128-cbc, arcfour, aes128-cbc, aes192-cbc, aes256-cbc, rijndael-cbc, aes128-ctr, aes192-ctr, aes256-ctr
[Table 11](#) lists the default ciphers used for SSHv2:

Table 11 SSHv2 Default Ciphers

Cipher index value	Cipher name
190	aes256-ctr
192	aes192-ctr
194	aes128-ctr
200	aes128-cbc
205	3des-cbc
210	blowfish-cbc
215	cast128-cbc
220	arcfour
225	aes192-cbc
230	aes256-cbc
235	rijndael-cbc



Note: blowfish-cbc, cast128-cbc, arcfour, and rijndael-cbc are not permitted in FIPS-140-2 mode.

disable-graceful-shutdown

Syntax [no] `disable-graceful-shutdown`

Context config>system>login-control>ssh

Description This command enables graceful shutdown of SSH sessions.

The **no** form of the command disables graceful shutdown of SSH sessions.

preserve-key

Syntax	[no] preserve-key
Context	config>system>security>ssh
Description	After enabling this command, private keys, public keys, and host key file will be saved by the server. It is restored following a system reboot or the ssh server restart. The no form of the command specifies that the keys will be held in memory by the SSH server and is not restored following a system reboot.
Default	no preserve-key

server-cipher-list protocol-version

Syntax	server-cipher-list protocol-version <i>version</i>
Context	config>system>security>ssh
Description	This command enables configuration the list of allowed ciphers by the SSH server.
Parameters	<i>version</i> — Specifies the SSH version. Values 1 — Specifies that the SSH server will only accept connections from clients that support SSH protocol version 1 2 — Specifies that the SSH server will accept connections from clients supporting either SSH protocol version 2

server-shutdown

Syntax	[no] server-shutdown
Context	config>system>security>ssh
Description	This command enables the SSH servers running on the system.
Default	At system startup, only the SSH server is enabled.

version

Syntax	version <i>ssh-version</i> no version
Context	config>system>security>ssh
Description	Specifies the SSH protocol version that will be supported by the SSH server.

Default	2
Parameters	<i>ssh-version</i> — Specifies the SSH version.
Values	<p>1 — Specifies that the SSH server will only accept connections from clients that support SSH protocol version 1</p> <p>2 — Specifies that the SSH server will accept connections from clients supporting either SSH protocol version 2</p> <p>1-2 — Specifies that the SSH server will accept connections from clients supporting either SSH protocol version 1, or SSH protocol version 2 or both.</p>



Note: Values “1” and “1-2” are not permitted in FIPS-140-2 mode.

telnet

Syntax	telnet
Context	config>system>login-control
Description	This command creates the context to configure the Telnet login control parameters.

enable-graceful-shutdown

Syntax	[no] enable-graceful-shutdown
Context	config>system>login-control>telnet
Description	<p>This command enables graceful shutdown of telnet sessions.</p> <p>The no form of the command disables graceful shutdown of telnet sessions.</p>

2.19.2.4 Management Access Filter Commands

management-access-filter

Syntax	[no] management-access-filter
Context	config>system>security
Description	This command creates the context to edit management access filters and to reset match criteria.

Management access filters control all traffic in and out of the CPM. They can be used to restrict management of the router by other nodes outside either specific (sub)networks or through designated ports.

Management filters, as opposed to other traffic filters, are enforced by system software.

The **no** form of the command removes management access filters from the configuration.

Default No management access filters are defined.

ip-filter

Syntax **[no] ip-filter**

Context config>system>security>mgmt-access-filter

Description This command enables the context to configure management access IP filter parameters.

ipv6-filter

Syntax **[no] ipv6-filter**

Context config>system>security>mgmt-access-filter

Description This command enables the context to configure management access IPv6 filter parameters. This command only applies to the 7750 SR and 7950 XRS.

mac-filter

Syntax **[no] mac-filter**

Context config>system>security>mgmt-access-filter

Description This command configures a management access MAC-filter.

action

Syntax **action {permit | deny | deny-host-unreachable}**
no action

Context config>system>security>mgmt-access-filter>ip-filter>entry
config>system>security>mgmt-access-filter>ipv6-filter>entry
config>system>security>mgmt-access-filter>mac-filter

Description This command creates the action associated with the management access filter match criteria entry.

The **action** keyword is required. If no **action** is defined, the filter is ignored. If multiple action statements are configured, the last one overwrites previous configured actions.

If the packet does not meet any of the match criteria the configured **default action** is applied.

Default none — The action is specified by default-action command.

Parameters *permit* — Specifies that packets matching the configured criteria will be permitted.

deny — Specifies that packets matching the configured selection criteria will be denied and that a ICMP host unreachable message will not be issued.

deny-host-unreachable — Specifies that packets matching the configured selection criteria will be denied and that a host unreachable message will not be issued.

The **deny-host-unreachable** parameter only applies to ip-filter and ipv6filter.

default-action

Syntax **default-action** {*permit* | *deny* | *deny-host-unreachable*}

Context config>system>security>mgmt-access-filter>ip-filter
config>system>security>mgmt-access-filter>ipv6-filter
config>system>security>mgmt-access-filter>mac-filter

Description This command creates the default action for management access in the absence of a specific management access filter match.

The **default-action** is applied to a packet that does not satisfy any match criteria in any of the management access filters. Whenever management access filters are configured, the **default-action** must be defined.

Default No default-action is defined.

Parameters **permit** — Specifies that packets not matching the configured selection criteria in any of the filter entries will be permitted.

deny — Specifies that packets not matching the selection criteria be denied and that an ICMP host unreachable message will not be issued.

deny-host-unreachable — Specifies that packets not matching the selection criteria be denied access and that an ICMP host unreachable message will be issued.

The **deny-host-unreachable** only applies to ip-filter and ipv6filter.

dst-port

Syntax [**no**] **dst-port** *value* [*mask*]

Context config>system>security>mgmt-access-filter>ip-filter>entry
config>system>security>mgmt-access-filter>ipv6-filter>entry

Description This command configures a source TCP or UDP port number or port range for a management access filter match criterion.

The **no** form of the command removes the source port match criterion.

Default No dst-port match criterion.

Parameters *value* — The source TCP or UDP port number as match criteria.

Values 1 to 65535 (decimal)

mask — Mask used to specify a range of source port numbers as the match criterion.

This 16 bit mask can be configured using the formats described in [Table 12](#):

Table 12 Format Styles to Configure Mask

Format Style	Format Syntax	Example
Decimal	DDDDD	63488
Hexadecimal	0xHHHH	0xF800
Binary	0bBBBBBBBBBBBBBB BB	0b1111100000000000

To select a range from 1024 up to 2047, specify 1024 0xFC00 for value and mask.

Default 65535 (exact match)

Values 1 to 65535 (decimal)

entry

Syntax **[no]** **entry** *entry-id*

Context
 config>system>security>mgmt-access-filter>ip-filter
 config>system>security>mgmt-access-filter>ipv6-filter
 config>system>security>mgmt-access-filter>mac-filter

Description This command is used to create or edit a management access IP(v4), IPv6, or MAC filter entry. Multiple entries can be created with unique *entry-id* numbers. The OS exits the filter upon the first match found and executes the actions according to the respective action command. For this reason, entries must be sequenced correctly from most to least explicit.

An entry may not have any match criteria defined (in which case, everything matches) but must have at least the keyword **action** defined to be considered complete. Entries without the **action** keyword are considered incomplete and inactive.

The **no** form of the command removes the specified entry from the management access filter.

Default No entries are defined.

Parameters *entry-id* — An entry ID uniquely identifies a match criteria and the corresponding action. It is recommended that entries are numbered in staggered increments. This allows users to insert a new entry in an existing policy without having to renumber the existing entries.

Values 1 to 9999

flow-label

Syntax **flow-label** *value*
no flow-label

Context config>system>security>mgmt-access-filter>ipv6-filter>entry

Description This command configures flow label match conditions. Flow labeling enables the labeling of packets belonging to particular traffic flows for which the sender requests special handling, such as non-default quality of service or real-time service. This command only applies to the 7750 SR and 7950 XRS.

Parameters *value* — Specify the flow identifier in an IPv6 packet header that can be used to discriminate traffic flows (See RFC 3595, *Textual Conventions for IPv6 Flow Label*.)

Values 0 to 1048575

log

Syntax [**no**] log

Context config>system>security>mgmt-access-filter>ip-filter>entry
config>system>security>mgmt-access-filter>ipv6-filter>entry
config>system>security>mgmt-access-filter>mac-filter

Description This command enables match logging. When enabled, matches on this entry will cause the Security event mafEntryMatch to be raised.

Default no log

next-header

Syntax **next-header** *next-header*
no next-header

Context config>system>security>mgmt-access-filter>ipv6-filter>entry

Description	This command specifies the next header to match. The protocol type such as TCP, UDP or OSPF is identified by its respective protocol number. Well-known protocol numbers include ICMP(1), TCP(6), UDP(17). IPv6 Extension headers are identified by the next header IPv6 numbers as per RFC2460. This command only applies to the 7750 SR and 7950 XRS.
Parameters	<i>next-header</i> — Specifies for IPv4 MAF the IP protocol field, and for IPv6 the next header type to be used in the match criteria for this Management Access Filter Entry.
	Values
	next-header: 0 to 255, protocol numbers accepted in DHB
	keywords: none, crtp, crudp, egp, eigrp, encap, ether-ip, gre, icmp, drp, igmp, igp, ip, ipv6, ipv6-icmp, ipv6-no-nxt, isis, iso-ip, l2tp, spf-igp, pim, pnni, ptp, rdp, rsvp, stp, tcp, udp, vrrp

protocol

Syntax	[no] protocol <i>protocol-id</i>
Context	config>system>security>mgmt-access-filter>ip-filter>entry
Description	This command configures an IP protocol type to be used as a management access filter match criterion. The protocol type, such as TCP, UDP, and OSPF, is identified by its respective protocol number. Well-known protocol numbers include ICMP (1), TCP (6), and UDP (17). The no form the command removes the protocol from the match criteria.
Default	No protocol match criterion is specified.
Parameters	<i>protocol</i> — The protocol number for the match criterion.
	Values 1 to 255 (decimal)

port

Syntax	port <i>tcp/udp port-number [mask]</i> port port-list <i>port-list-name</i> port range <i>tcp/udp port-number tcp/udp port-number</i> no port
Context	config>system>security>cpm-filter>ip-filter>entry>match config>system>security>cpm-filter>ipv6-filter>entry>match
Description	This command configures a TCP/UDP source or destination port match criterion in IPv4 and IPv6 CPM filter policies. A packet matches this criterion if packet's TCP/UDP (as configured by protocol/next-header match) source OR destination port matches either the specified port value or a port in the specified port range or port list.

This command is mutually exclusive with **src-port** and **dst-port** commands.

The **no** form of this command deletes the specified port match criterion.

Default	no port
Parameters	<p><i>tcp/udp port-number</i> — A source or destination port to be used as a match criterion specified as a decimal integer.</p> <p>Values 0 to 65535</p> <p><i>mask</i> — Specifies the 16 bit mask to be applied when matching the port.</p> <p>Values [0x0000 to 0xFFFF] [0 to 65535] [0b0000000000000000. to 0b1111111111111111]</p> <p>range <i>tcp/udp port-number</i> — an inclusive range of source or destination port values to be used as match criteria. <i>start</i> of the range and <i>end</i> of the range are expressed as decimal integers.</p> <p>Values start, end, port-number: 1 to 65535</p> <p>port-list <i>port-list-name</i> — A string of up to 32 characters of printable ASCII characters. If special characters are used, the string must be enclosed within double quotes.</p>

router

Syntax	<pre>router service-name service-name router {router-instance} no router</pre>
Context	<pre>config>system>security>mgmt-access-filter>ip-filter>entry config>system>security>mgmt-access-filter>ipv6-filter>entry</pre>
Description	<p>This command configures a router name or service ID to be used as a management access filter match criterion.</p> <p>The no form the command removes the router name or service ID from the match criteria.</p>
Parameters	<p><i>router-instance</i> — Specify one of the following parameters for the router instance:</p> <p><i>router-name</i> — Specifies a router name or CPM router instance, up to 32 characters to be used in the match criteria.</p> <pre>router-instance : router name router-name Base management cpm-vr-name cpm-vr-name [32 characters maximum]</pre> <p><i>service-id</i> — Specifies an existing service ID to be used in the match criteria.</p>

cpm-vr-name — Specifies a CPM router instance to be used in the match criteria

Values 1 to 2147483647

service-name service-name — Specifies an existing service name up to 64 characters in length.

renum

Syntax **renum** *old-entry-number new-entry-number*

Context config>system>security>mgmt-access-filter>ip-filter
config>system>security>mgmt-access-filter>ipv6-filter
config>system>security>mgmt-access-filter>mac-filter

Description This command renumbers existing management access filter entries for an IP(v4), IPv6, or MAC filter to re-sequence filter entries.

The exits on the first match found and executes the actions in accordance with the accompanying **action** command. This may require some entries to be re-numbered differently from most to least explicit.

Parameters *old-entry-number* — Enter the entry number of the existing entry.

Values 1 to 9999

new-entry-number — Enter the new entry number that will replace the old entry number.

Values 1 to 9999

shutdown

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

Context config>system>security>mgmt-access-filter>ip-filter
config>system>security>mgmt-access-filter>ipv6-filter
config>system>security>mgmt-access-filter>mac-filter

Description This command disables the management-access-filter.

match

Syntax **match** [**frame-type** *frame-type*]
no match

Context config>system>security>mgmt-access-filter>mac-filter>entry

Description This command configures math criteria for this MAC filter entry.

Parameters **frame-type** *frame-type* — Specifies the type of MAC frame to use as match criteria.
Values none, 802dot2-llc, ethernet_II

cfm-opcode

Syntax **cfm-opcode** {*lt* | *gt* | *eq*} *opcode*
cfm-opcode range *start end*
no cfm-opcode

Context config>system>security>mgmt-access-filter>mac-filter>entry

Description This command specifies the type of opcode checking to be performed.

If the cfm-opcode match condition is configured then a check must be made to see if the Ethertype is either IEEE802.1ag or Y1731. If the Ethertype does not match then the packet is not CFM and no match to the cfm-opcode is attempted.

The CFM (ieee802.1ag or Y1731) opcode can be assigned as a range with a start and an end number or with a (less than lt, greater than gt, or equal to eq) operator.

If no range with a start and an end or operator (lt, gt, eq) followed by an opcode with the value between 0 and 255 is defined then the command is invalid.

[Table 13](#) lists the opcode values.

Table 13 Opcode Values

CFM PDU or Organization	Acronym	Configurable Numeric Value (Range)
Reserved for IEEE 802.1 0		0
Continuity Check Message	CCM	1
Loopback Reply	LBR	2
Loopback Message	LBM	3
Linktrace Reply	LTR	4
Linktrace Message	LTM	5
Reserved for IEEE 802.1		6 – 31
Reserved for ITU		32
	AIS	33
Reserved for ITU		34
	LCK	35

Table 13 Opcode Values (Continued)

CFM PDU or Organization	Acronym	Configurable Numeric Value (Range)
Reserved for ITU		36
	TST	37
Reserved for ITU		38
	APS	39
Reserved for ITU		40
	MCC	41
	LMR	42
	LMM	43
Reserved for ITU		44
	1DM	45
	DMR	46
	DMM	47
Reserved for ITU		48 – 63
Reserved for IEEE 802.1 0		64 - 255

Defined by ITU-T Y.1731 32 - 63

Defined by IEEE 802.1.64 - 255

Default no cfm-opcode

Parameters *opcode* — Specifies the opcode checking to be performed.

start — specifies the start number.

Values 0 to 255

end — Specifies the end number.

Values 0 to 255

lt | gt | eq — keywords

dot1p

Syntax **dot1p** *dot1p-value* [*dot1p-mask*]

- Context** config>system>security>mgmt-access-filter>mac-filter>entry>match
- Description** This command configures Dot1p match conditions.
- Parameters** *dot1p-value* — Specifies the IEEE 802.1p value in decimal.
Values 0 to 7
- mask* — Specifies the 3-bit mask can be configured using the following formats:
Values 0 to 7

dsap

- Syntax** dsap dsap-value [dsap-mask]
- Context** config>system>security>mgmt-access-filter>mac-filter>entry>match
- Description** This command configures dsap match conditions.

Table 14 Management Access Filter dsap Format Style

Format Style	Format Syntax	Example
Decimal	D	4
Hexadecimal	0xH	0x4
Binary	0bBBB	0b100

- Parameters** *dsap-value* — Specifies the 8-bit dsap match criteria value in hexadecimal.
Values 0x00 to 0xFF (hex)
- mask* — This is optional and may be used when specifying a range of dsap values to use as the match criteria.
 This 8 bit mask can be configured using the formats described in [Table 15](#):

Table 15 Format Styles

Format Style	Format Syntax	Example
Decimal	DDD	240
Hexadecimal	0xHH	0xF0
Binary	0bBBBBBBBB	0b11110000

- Default** FF (hex) (exact match)
- Values** 0x00 to 0xFF

dst-mac

Syntax	dst-mac <i>ieee-address</i> [<i>ieee-address-mask</i>] no dst-mac
Context	config>system>security>mgmt-access-filter>mac-filter>entry>match
Description	This command configures the destination MAC match condition.
Parameters	<i>ieee-address</i> — Specifies the MAC address to be used as a match criterion. Values HH:HH:HH:HH:HH:HH or HH-HH-HH-HH-HH-HH where H is a hexadecimal digit <i>mask</i> — A 48-bit mask to match a range of MAC address values.

etype

Syntax	etype <i>0x0600xx0xffff</i> no etype
Context	config>system>security>mgmt-access-filter>mac-filter>entry>match
Description	Configures an Ethernet type II Ethertype value to be used as a MAC filter match criterion. The Ethernet type field is a two-byte field used to identify the protocol carried by the Ethernet frame. For example, 0800 is used to identify the IPv4 packets. The Ethernet type field is used by the Ethernet version-II frames. IEEE 802.3 Ethernet frames do not use the type field. For IEEE 802.3 frames, use the dsap, ssap or snap-pid fields as match criteria. The snap-pid field, etype field, ssap and dsap fields are mutually exclusive and may not be part of the same match criteria. Refer to the Router Configuration Guide for information about MAC Match Criteria Exclusivity Rules fields that are exclusive based on the frame format. The no form of the command removes the previously entered etype field as the match criteria.
Default	no etype
Parameters	<i>ethernet-type</i> — Specifies the Ethernet type II frame Ethertype value to be used as a match criterion expressed in hexadecimal. Values 0x0600 to 0xFFFF

snap-oui

Syntax	snap-oui { zero non-zero }
---------------	---

Context config>system>security>mgmt-access-filter>mac-filter>entry>match

Description This command configures an IEEE 802.3 LLC SNAP Ethernet Frame OUI zero or non-zero value to be used as a MAC filter match criterion.

The **no** form of the command removes the criterion from the match criteria.

Default no snap-oui

Parameters **zero** — Specifies to match packets with the three-byte OUI field in the SNAP-ID set to zero.

non-zero — Specifies to match packets with the three-byte OUI field in the SNAP-ID not set to zero.

snap-pid

Syntax **snap-pid** *snap-pid*
no snap-pid

Context config>system>security>mgmt-access-filter>mac-filter>entry>match

Description This command configures an IEEE 802.3 LLC SNAP Ethernet Frame PID value to be used as a MAC filter match criterion.

This is a two-byte protocol id that is part of the IEEE 802.3 LLC SNAP Ethernet Frame that follows the three-byte OUI field.

The snap-pid field, etype field, ssap and dsap fields are mutually exclusive and may not be part of the same match criteria. Refer to the Router Configuration Guide for information about MAC Match Criteria Exclusivity Rules fields that are exclusive based on the frame format.



Note: The snap-pid match criterion is independent of the OUI field within the SNAP header. Two packets with different three-byte OUI fields but the same PID field will both match the same filter entry based on a snap-pid match criteria.

The **no** form of the command removes the snap-pid value as the match criteria.

Default no snap-pid

Parameters *pid-value* — Specifies the two-byte snap-pid value to be used as a match criterion in hexadecimal.

Values 0x0000 to 0xFFFF

src-mac

- Syntax** **src-mac** *ieee-address* [*ieee-address-mask*]
no src-mac
- Context** config>system>security>mgmt-access-filter>mac-filter>entry>match
- Description** This command configures a source MAC address or range to be used as a MAC filter match criterion.
- The **no** form of the command removes the source mac as the match criteria.
- Default** no src-mac
- Parameters** *ieee-address* — Specifies the 48-bit IEEE mac address to be used as a match criterion.
- Values** HH:HH:HH:HH:HH:HH or HH-HH-HH-HH-HH-HH where H is a hexadecimal digit
- ieee-address-mask* — This 48-bit mask can be configured using the formats listed in [Table 16](#):

Table 16 ieee-address-mask Formats

Format Style	Format Syntax	Example
Decimal	DDDDDDDDDDDDDDDD	281474959933440
Hexadecimal	0xHHHHHHHHHHHHHH	0x0FFFFFF000000
Binary	0bBBBBBBB...B	0b11110000...B

To configure so that all packets with a source MAC OUI value of 00-03-FA are subject to a match condition then the entry should be specified as: 003FA000000 0xFFFFFFFF000000

- Default** 0xFFFFFFFFFFFFFF (exact match)
- Values** 0x0000000000000000 to 0xFFFFFFFFFFFFFF

ssap

- Syntax** **ssap** *ssap-value* [*ssap-mask*]
no ssap
- Context** config>system>security>mgmt-access-filter>mac-filter>entry>match
- Description** This command configures an Ethernet 802.2 LLC SSAP value or range for a MAC filter match criterion.
- This is a one-byte field that is part of the 802.2 LLC header of the IEEE 802.3 Ethernet Frame.

The snap-pid field, etype field, ssap and dsap fields are mutually exclusive and may not be part of the same match criteria. Refer to the Router Configuration Guide for information about MAC Match Criteria Exclusivity Rules fields that are exclusive based on the frame format.

The **no** form of the command removes the ssap match criterion.

Default no ssap

Parameters *ssap-value* — Specifies the 8-bit ssap match criteria value in hex.

Values 0x00 to 0xFF

ssap-mask — This is optional and may be used when specifying a range of ssap values to use as the match criteria.

svc-id

Syntax **svc-id** *service-id*
no svc-id

Context config>system>security>mgmt-access-filter>mac-filter>entry>match

Description This command specifies an existing svc-id to use as a match condition.

Parameters *service-id* — Specifies a service-id to match.

Values *service-id*: 1 to 2147483647
svc-name: 64 characters maximum

src-port

Syntax **src-port** {*port-id* | **cpm** | **lag** *lag-id*}
no src-port

Context config>system>security>mgmt-access-filter>ip-filter>entry
config>system>security>mgmt-access-filter>ipv6-filter>entry

Description This command restricts ingress management traffic to either the CPMCCM Ethernet port or any other logical port (for example LAG) on the device.

When the source interface is configured, only management traffic arriving on those ports satisfy the match criteria.

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

Default any interface

Parameters *port-id* — Specifies the port ID in formats shown below

Values

	<i>slot/mda/port[.channel]</i>		
<i>bundle-id</i>	<i>bundle-type-slot/mda.bundle-num</i>		
	<i>bundle</i>		keyword
	<i>type</i>		ima, fr, or ppp
	<i>bundle-num</i>		1 to 336
<i>bpgrp-id</i>	<i>bpgrp-type-bpgrp-num</i>		
	<i>bpgrp</i>		keyword
	<i>type</i>		ima or ppp
	<i>bpgrp-num</i>		1 to 2000
<i>aps-id</i>	<i>aps-group-id[.channel]</i>		
	<i>aps</i>		keyword
	<i>group-id</i>		1 to 128
<i>ccag-id</i>	<i>ccag-id.path-id[cc-type]</i>		
	<i>ccag</i>		keyword
	<i>id</i>		1 to 8
	<i>path-id</i>		a, b
	<i>cc-type</i>		.sap-net, .net-sap

cpm — Matches any traffic received on any Ethernet port

lag-id — Specifies the LAG identifier

Values 1 to 800

src-ip

Syntax **[no] src-ip** {[*ip-prefix/mask*] | [*ip-prefix*] | *ip-prefix-list prefix-list-name*}

Context config>system>security>mgmt-access-filter>ip-filter>entry

Description This command configures a source IP address range prefix to be used as a management access filter match criterion.

The **no** form of the command removes the source IP address match criterion.

Default No source IP match criterion is specified.

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

ip-prefix-list — Creates a list of IPv4 prefixes for match criteria in IPv4 ACL and CPM filter policies.

ip-prefix-list-name — A string of up to 32 characters of printable ASCII characters. If special characters are used, the string must be enclosed within double quotes.

mask — Specifies the subnet mask length expressed as a decimal integer.

Values 1 to 32 (mask length), 0.0.0.0 to 255.255.255.255 (dotted decimal)

src-ip

Syntax **[no] src-ip** {[*ip-prefix/mask*] | [*ip-prefix*] | *ip-prefix-list prefix-list-name*}

Context config>system>security>mgmt-access-filter>ipv6-filter>entry

Description This command configures a source IPv6 address range prefix to be used as a management access filter match criterion. This command only applies to the 7750 SR and 7950 XRS.

The **no** form of the command removes the source IPv6 address match criterion.

Default No source IP match criterion is specified.

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

ip-prefix-list — Creates a list of IPv4 prefixes for match criteria in IPv4 ACL and CPM filter policies.

ipv6-prefix-list-name — A string of up to 32 characters of printable ASCII characters. If special characters are used, the string must be enclosed within double quotes.

mask — Specifies the subnet mask length expressed as a decimal integer.

Values 1 to 32 (mask length), 0.0.0.0 to 255.255.255.255 (dotted decimal)

2.19.2.5 Password Commands

password

Syntax	password
Context	config>system>security
Description	This command creates the context to configure password management parameters.

admin-password

Syntax	admin-password <i>password</i> [hash hash2] no admin-password
Context	config>system>security>password
Description	This command allows a user (with admin permissions) to configure a password which enables a user to become an administrator.

This password is valid only for one session. When enabled, no authorization to TACACS+ or RADIUS is performed and the user is locally regarded as an admin user.

This functionality can be enabled in two contexts:

```
config>system>security>password>admin-password
```

```
<global> enable-admin
```

If the **admin-password** is configured in the config>system>security>password context, then any user can enter the special mode by entering the **enable-admin** command. For more information, see the description for the `enable-admin` command.

enable-admin is in the default profile. By default, all users are given access to this command.

Once the **enable-admin** command is entered, the user is prompted for a password. If the password matches, user is given unrestricted access to all the commands.

The minimum length of the password is determined by the **minimum-length** command. The complexity requirements for the password is determined by the **complexity** command.



Note: The *password* argument of this command is not sent to the servers. This is consistent with other commands that configure secrets.

The usernames and passwords in the FTP and TFTP URLs will not be sent to the authorization or accounting servers when the **file>copy source-url dest-url** command is executed.

For example:

```
file copy ftp://test:secret@131.12.31.79/test/srcfile cf1:\destfile
```

In this example, the username 'test' and password 'secret' will not be sent to the AAA servers (or to any logs). They will be replaced with '*****'.

The **no** form of the command removes the admin password from the configuration.

Default no admin-password

Parameters *password* — Configures the password which enables a user to become a system administrator. The maximum length can be up to 20 characters if unhashed, 32 characters if hashed, 54 characters if the hash2 keyword is specified.

hash — Specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified

hash2 — Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

enable-admin

Syntax **enable-admin**

Context <global>

Description Refer to the description for the `enable-admin` command. If the **admin-password** is configured in the `config>system>security>password` context, then any user can enter the special administrative mode by entering the `enable-admin` command.

enable-admin is in the default profile. By default, all users are given access to this command.

Once the **enable-admin** command is entered, the user is prompted for a password. If the password matches, user is given unrestricted access to all the commands.

The minimum length of the password is determined by the **minimum-length** command. The complexity requirements for the password is determined by the **complexity** command.

There are two ways to verify that a user is in the enable-admin mode:

- **show users** — Administrator can know which users are in this mode.

- Enter the enable-admin command again at the root prompt and an error message will be returned.

```
A:ALA-1# show users
=====
User Type From Login time Idle time
=====
admin Console -- 10AUG2006 13:55:24 0d 19:42:22
admin Telnet 10.20.30.93 09AUG2006 08:35:23 0d 00:00:00 A
-----
Number of users : 2
'A' indicates user is in admin mode
=====
A:ALA-1#
A:ALA-1# enable-admin
MINOR: CLI Already in admin mode.
A:ALA-1#
```

aging

Syntax	aging <i>days</i> no aging
Context	config>system>security>password
Description	This command configures the number of days a user password is valid before the user must change their password. This parameter can be used to force the user to change the password at the configured interval. The no form of the command reverts to the default value.
Default	No aging is enforced.
Parameters	<i>days</i> — The maximum number of days the password is valid. Values 1 to 500

attempts

Syntax	attempts <i>count</i> [time <i>minutes1</i> [lockout <i>minutes2</i>]] no attempts
Context	config>system>security>password
Description	This command configures a threshold value of unsuccessful login attempts allowed in a specified time frame. If the threshold is exceeded, the user is locked out for a specified time period. If multiple attempts commands are entered, each command overwrites the previously entered command.

The **no attempts** command resets all values to default.

Default	attempts count 3 time 5 lockout 10
Parameters	<p>count — The number of unsuccessful login attempts allowed for the specified time. This is a mandatory value that must be explicitly entered.</p> <p>Values 1 to 64</p> <p>time minutes — The period of time, in minutes, that a specified number of unsuccessful attempts can be made before the user is locked out.</p> <p>Values 0 to 60</p> <p>lockout minutes — The lockout period, in minutes, during which the user is not allowed to login.</p> <p>Values 0 to 1440, or infinite</p> <p>If the user exceeds the attempted count times in the specified time, then that user is locked out from any further login attempts for the configured lockout time period.</p> <p>Default 10</p> <p>Values 0 to 1440</p> <p>Values infinite; user is locked out and must wait until manually unlocked before any further attempts.</p>

authentication-order

Syntax	authentication-order [<i>method-1</i>] [<i>method-2</i>] [<i>method-3</i>] [<i>method-4</i>] [exit-on-reject] no authentication-order
Context	config>system>security>password
Description	<p>This command configures the sequence in which password authentication, authorization, and accounting is attempted among local passwords, RADIUS, TACACS+, and LDAP.</p> <p>The authentication order should be from the most preferred authentication method to the least preferred. The presence of all methods in the command line does not guarantee that they are all operational. Specifying options that are not available delays user authentication.</p> <p>If all (operational) methods are attempted and no authentication for a particular login has been granted, then an entry in the security log documents the failed attempt. Both the attempted login identification and originating IP address are logged with the a timestamp.</p> <p>The no form of the command reverts to the default authentication sequence.</p>
Default	authentication-order radius tacplus ldap local - The preferred order for password authentication is 1. local passwords, 2. RADIUS, 3. TACACS+, and 4. LDAP.

Parameters *method-1* — The first password authentication method to attempt.

Default local

Values local, radius, tacplus, ldap

method-2 — The second password authentication method to attempt.

Default radius

Values local, radius, tacplus, ldap

method-3 — The third password authentication method to attempt.

Default tacplus

Values local, radius, tacplus, ldap

method-4 — The fourth password authentication method to attempt.

Default ldap

Values local, radius, tacplus, ldap

local — Password authentication based on the local password database.

radius — RADIUS authentication.

tacplus — TACACS+ authentication.

ldap — LDAP authentication.

exit-on-reject — When enabled and if one of the AAA methods configured in the authentication order sends a reject, then the next method in the order will not be tried. If the **exit-on-reject** keyword is not specified and if one AAA method sends a reject, the next AAA method will be attempted. If in this process, all the AAA methods are exhausted, it will be considered as a reject.

A rejection is distinct from an unreachable authentication server. When the **exit-on-reject** keyword is specified, authorization and accounting will only use the method that provided an affirmation authentication; only if that method is no longer readable or is removed from the configuration will other configured methods be attempted. If the **local** keyword is the first authentication and:

- **exit-on-reject** is configured and the user does not exist, the user will not be authenticated
- the user is authenticated locally, then other methods, if configured, will be used for authorization and accounting
- the user is configured locally but without console access, login will be denied

complexity-rules

Syntax **complexity-rules**

Context config>system>security>password

Description This command defines a list of rules for configurable password options.

allow-user-name

Syntax	[no] allow-user-name
Context	config>system>security>password>complexity-rules
Description	The user name is allowed to be used as part of the password. The no form of the command does not allow user name to be used as password.
Default	no allow-user-name

credits

Syntax	credits [lowercase <i>credits</i>] [uppercase <i>credits</i>] [numeric <i>credits</i>] [special-character <i>credits</i>] no credits
Context	config>system>security>password>complexity-rules
Description	The maximum credits given for usage of the different character classes in the local passwords. The no form of the command resets to default.
Default	no credits
Parameters	<i>credits</i> — The number of credits that can be used for each characters class. Values 0 to 10

minimum-classes

Syntax	minimum-classes <i>minimum</i> no minimum-classes
Context	config>system>security>password>complexity-rules
Description	Force the use of at least this many different character classes The no form of the command resets to default.
Default	no minimum-classes
Parameters	<i>minimum</i> — The minimum number of classes to be configured. Values 2 to 4

minimum-length

Syntax	minimum-length <i>length</i> no minimum-length
Context	config>system>security>password
Description	<p>This command configures the minimum number of characters required for locally administered passwords, HMAC-MD5-96, HMAC-SHA-96, and des-keys configured in the system security section.</p> <p>If multiple minimum-length commands are entered each command overwrites the previous entered command.</p> <p>The no form of the command reverts to default value.</p>
Default	6
Parameters	<i>value</i> — The minimum number of characters required for a password. Values 1 to 8

repeated-characters

Syntax	repeated-characters <i>count</i> no repeated-characters
Context	config>system>security>password>complexity-rules
Description	<p>The number of times a characters can be repeated consecutively.</p> <p>The no form of the command resets to default.</p>
Default	no repeated-characters
Parameters	<i>count</i> — The minimum count of consecutively repeated characters. Values 2 to 8

required

Syntax	required [lowercase <i>count</i>] [uppercase <i>count</i>] [numeric <i>count</i>] [special-character <i>count</i>] no required
Context	config>system>security>password>complexity-rules
Description	<p>Force the minimum number of different character classes required.</p> <p>The no form of the command resets to default.</p>

Default	no required
Parameters	<i>count</i> — The minimum count of characters classes.
Values	0 to 10

dynsvc-password

Syntax	dynsvc-password <i>password</i> [hash hash2] no dynsvc-password
Context	config>system>security>password
Description	This command configures the password which enables the user to configure dynamic services.
Default	no dynsvc-password
Parameters	<p><i>password</i> — Configures the password which enables a user to become a system administrator. The maximum length can be up to 20 characters if unhashed, 32 characters if hashed, 54 characters if the hash2 keyword is specified.</p> <p>hash — Specifies the key is entered in an encrypted form. If the hash or hash2 parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the hash or hash2 parameter specified</p> <p>hash2 — Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the hash2 encrypted variable cannot be copied and pasted. If the hash or hash2 parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the hash or hash2 parameter specified.</p>

enable-admin-control

Syntax	enable-admin-control
Context	config>system>security>password
Description	Enable the user to become a system administrator.

tacplus-map-to-priv-lvl

Syntax	tacplus-map-to-priv-lvl [<i>admin-priv-lvl</i>] no tacplus-map-to-priv-lvl
Context	config>system>security>password>enable-admin-control

Description When **tacplus-map-to-priv-lvl** is enabled, and tacplus authorization is enabled with the *use-priv-lvl* option, typing **enable-admin** starts an interactive authentication exchange from the node to the TACACS+ server. The start message (service=enable) contains the user-id and the requested *admin-priv-lvl*. Successful authentication results in the use of a new profile (as configured under **config>system>security>tacplus>priv-lvl-map**).

health-check

Syntax **[no] health-check [interval interval]**

Context config>system>security>password

Description This command specifies that RADIUS, TACACS+, and LDAP servers are monitored for 3 seconds each at 30 second intervals. Servers that are not configured will have 3 seconds of idle time. If in this process a server is found to be unreachable, or a previously unreachable server starts responding, a trap will be sent based on the type of the server.

The **no** form of the command disables the periodic monitoring of the RADIUS, TACACS+, and LDAP servers. In this case, the operational status for the active server will be up if the last access was successful.

Parameters *interval* — Specifies the polling interval for RADIUS, TACACS+, and LDAP servers.

Values 6 to 1500

Default 30

history

Syntax **history size**
no history

Context config>system>security>password

Description Configure how many previous passwords a new password is matched against.

Default no history

Parameters *size* — Specifies how many previous passwords a new password is matched against.

Values 1 to 20

minimum-age

Syntax **minimum-age [days days] [hrs hours] [min minutes] [sec seconds]**
no minimum-age

Context config>system>security>password

Description	Configure the minimum required age of a password before it can be changed again.
Default	no minimum-age
Parameters	<p><i>days</i> — Specifies the minimum required days of a password before it can be changed again.</p> <p>Values 0 to1</p> <p><i>hours</i> — Specifies the minimum required hours of a password before it can be changed again.</p> <p>Values 0 to23</p> <p><i>minutes</i> — Specifies the minimum required minutes of a password before it can be changed again.</p> <p>Values 0 to59</p> <p><i>seconds</i> — Specifies the minimum required seconds of a password before it can be changed again.</p> <p>Values 0 to59</p>

minimum-change

Syntax	minimum-change <i>length</i> no minimum-change
Context	config>system>security>password
Description	<p>This command configures the minimum number of characters required to be different in the new password from a previous password.</p> <p>The no form of the command reverts to default value.</p>
Default	no min-change
Parameters	<p><i>length</i> — Specifies how many characters must be different in the new password from the old password.</p> <p>Values 2 to 20</p>

2.19.2.6 Public Key Infrastructure (PKI) Commands

The commands described in the following section apply to the 7450 ESS and 7750 SR.

pki

Syntax	pki
Context	config>system>security
Description	This command enables the context to configure certificate parameters.
Default	none

ca-profile

Syntax	ca-profile <i>name</i> [create] no ca-profile <i>name</i>
Context	config>system>security>pki
Description	This command creates a new ca-profile or enter the configuration context of an existing ca-profile . Up to 128 ca-profiles could be created in the system. A shutdown the ca-profile will not affect the current up and running ipsec-tunnel or ipsec-gw that associated with the ca-profile . But authentication afterwards will fail with a shutdown ca-profile . Executing a no shutdown command in this context will cause system to reload the configured cert-file and crl-file. A ca-profile can be applied under the ipsec-tunnel or ipsec-gw configuration. The no form of the command removes the name parameter from the configuration. A ca-profile can not be removed until all the association(ipsec-tunnel/gw) have been removed.
Parameters	<i>name</i> — Specifies the name of the ca-profile , a string up to 32 characters. create — This keyword creates a new ca-profile . The create keyword requirement can be enabled/disabled in the environment>create context.

cert-file

Syntax	cert-file <i>filename</i> no cert-file
Context	config>system>security>pki>ca-profile

Description This command specifies the filename of a file in cf3:\system-pki\cert as the CA's certificate of the ca-profile.

Notes:

- The system will perform following checks against configured cert-file when a **no shutdown** command is issued:
 - Configured cert-file must be a DER formatted X.509v3 certificate file.
 - All non-optional fields defined in section 4.1 of RFC5280 must exist and conform to the RFC 5280 defined format.
 - Check the version field to see if its value is 0x2.
 - Check The Validity field to see that if the certificate is still in validity period.
 - X509 basic constraints extension must exists, and CA Boolean must be True.
 - If Key Usage extension exists, then at least keyCertSign and cRLSign should be asserted.
 - If the certificate is not a self-signing certificate, then system will try to look for issuer's CA's certificate to verify if this certificate is signed by issuer's CA; but if there is no such CA-profile configured, then system will just proceed with a warning message.
 - If the certificate is not a self-signing certificate, then system will try to look for issuer's CA's CRL to verify that it has not been revoked; but if there is no such CA-profile configured or there is no such CRL, then system will just proceed with a warning message.
- If any of above checks fails, then the **no shutdown** command will fail.
- Changing or removing of **cert-file** is only allowed when the **ca-profile** is in a **shutdown** state.

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

Parameters *filename* — Specifies a local CF card file URL.

cmpv2

Syntax **cmpv2**

Context config>system>security>pki>ca-profile

Description This command enables the context to configure Certificate Management Protocol Version 2 (CMPv2) parameters.

accept-unprotected-errormsg

Syntax [**no**] **accept-unprotected-errormsg**

Context config>system>security>pki>ca-profile>cmpv2

- Description** This command enables the system to accept both protected and unprotected CMPv2 error message. Without this command, system will only accept protected error messages.
- The **no** form of the command causes the system to only accept protected PKI confirmation message.
- Default** no accept-unprotected-errormsg

accept-unprotected-pkiconf

- Syntax** [no] accept-unprotected-pkiconf
- Context** config>system>security>pki>ca-profile>cmpv2
- Description** This command enables the system to accept both protected and unprotected CMPv2 PKI confirmation messages. Without this command, the system will only accept protected PKI confirmation message.
- The **no** form of the command causes the system to only accept protected PKI confirmation message.
- Default** no accept-unprotected-errormsg

key-list

- Syntax** key-list
- Context** config>system>security>pki>ca-profile>cmp2
- Description** This command enables the context to configure pre-shared key list parameters.

key

- Syntax** key *password* [hash | hash2] **reference** *reference-number*
no key **reference** *reference-number*
- Context** config>system>security>pki>ca-profile>cmp2>key-list
- Description** This command specifies a pre-shared key used for CMPv2 initial registration. Multiples of key commands are allowed to be configured under this context.
- The password and reference-number is distributed by the CA via out-of-band means.
- The configured password is stored in configuration file in an encrypted form by using the SR OS hash2 algorithm.
- The **no** form of the command removes the parameters from the configuration.

Default none

Parameters *password* — Specifies a printable ASCII string, up to 64 characters in length.

hash — Specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified

hash2 — Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

reference-number — Specifies a printable ASCII string, up to 64 characters in length.

url

Syntax **url** *url-string* [**service-id** *service-id*]
no url

Context config>system>security>pki>ca-profile>cmp2

Description This command specifies HTTP URL of the CMPv2 server. The URL must be unique across all configured ca-profiles.

The URL will be resolved by the DNS server configured (if configured) in the corresponding router context.

If the *service-id* is 0 or omitted, then system will try to resolve the FQDN via DNS server configured in bof.cfg. After resolution, the system will connect to the address in management routing instance first, then base routing instance.



Note: If the service is VPRN, then the system only allows HTTP ports 80 and 8080.

Default none

Parameters *url-string* — Specifies the HTTP URL of the CMPv2 server up to 180 characters in length.

service-id *service-id* — Specifies the service instance that used to reach CMPv2 server.

Values service-id: 1 to 2147483647
base-router: 0

http-response-timeout

Syntax	http-response-timeout <i>timeout</i> no http-response-timeout
Context	config>system>security>pki>ca-profile>cmp2
Description	This command specifies the timeout value for HTTP response that is used by CMPv2. The no form of the command reverts to the default.
Default	30
Parameters	<i>timeout</i> — Specifies the HTTP response timeout in seconds. Values 1 to 3600

response-signing-cert

Syntax	response-signing-cert <i>filename</i> no response-signing-cert
Context	config>system>security>pki>ca-profile>cmp2
Description	This command specifies a imported certificate that is used to verify the CMP response message if they are protected by signature. If this command is not configured, then CA's certificate will be used.
Default	no response-signing-cert
Parameters	<i>filename</i> — Specifies the filename of the imported certificate.

same-recipnonce-for-pollreq

Syntax	[no] same-recipnonce-for-pollreq
Context	config>system>security>pki>ca-profile>cmp2
Description	This command enables the system to use same recipNonce as the last CMPv2 response for poll request. The no form of the command disables system to use same recipNonce as the last CMPv2 response for poll request.
Default	no same-recipnonce-for-pollreq

crl-file

Syntax	crl-file <i>filename</i> no crl-file
Context	config>system>security>pki>ca-profile
Description	This command specifies the name of a file in cf3:\system-pki\crl as the Certification Revoke List file of the ca-profile .
	Notes:
	<ul style="list-style-type: none"> • The system will perform following checks against configured crl-file when a no shutdown command is issued: <ul style="list-style-type: none"> – A valid cert-file of the ca-profile must be already configured. – Configured crl-file must be a DER formatted CRLv2 file. – All non-optional fields defined in section 5.1 of RFC5280 must exist and conform to the RFC5280 defined format. – Check the version field to see if its value is 0x1. – Delta CRL Indicator must NOT exists (delta CRL is not supported). – CRL's signature must be verified by using the cert-file of ca-profile. If any of above checks fail, the no shutdown command will fail. • Changing or removing the crl-file is only allowed when the ca-profile is in a shutdown state.
	The no form of the command removes the filename from the configuration.
Default	none
Parameters	<i>filename</i> — Specifies the name of CRL file stored in cf3:\system-pki\crl.

ocsp

Syntax	ocsp
Context	config>system>security>pki>ca-profile
Description	This command enables the context to configure OCSP parameters.

responder-url

Syntax	responder-url <i>url-string</i> no responder-url
Context	config>system>security>pki>ca-profile>ocsp

Description	This command specifies HTTP URL of the OCSP responder for the CA, this URL will only be used if there is no OCSP responder defined in the AIA extension of the certificate to be verified.
Default	no responder-url
Parameters	<i>url-string</i> — Specifies the HTTP URL of the OCSP responder

service

Syntax	service <i>service-id</i> no service
Context	config>system>security>pki>ca-profile>ocsp
Description	This command specifies the service or routing instance that used to contact OCSP responder. This applies to OCSP responders that either configured in CLI or defined in AIA extension of the certificate to be verified. The responder-url will also be resolved by using the DNS server configured in the configured routing instance. In case of VPRN service, system will check if the specified service-id or service-name is an existing VPRN service at the time of CLI configuration. Otherwise the configuration will fail.
Parameters	<i>service-id</i> — Specifies an existing service ID to be used in the match criteria. Values service-id: 1 to 2147483647 base-router: 0

certificate-display-format

Syntax	certificate-display-format { <i>ascii</i> <i>utf8</i> }
Context	config>system>security>pki
Description	This command specifies the display format used for the Certificates and Certificate Revocation Lists.
Default	ascii
Parameters	ascii — Specifies the ASCII format to use for the Certificates and Certificate Revocation Lists. utf8 — Specifies the UTF8 format to use for the Certificates and Certificate Revocation Lists.

certificate-expiration-warning

Syntax	certificate-expiration-warning <i>hours</i> [repeat <i>repeat-hours</i>] no certificate-expiration-warning
Context	config>system>security>pki
Description	With this command configured, the system will issues two types of warnings related to certificate expiration:

- **BeforeExp** — A warning message issued before certificate expire
- **AfterExp** — A warning message issued when certificate expire

This command specifies when system will issue **BeforeExp** message before a certificate expires. For example, with **certificate-expiration-warning 5**, the system will issue a **BeforeExp** message 5 hours before a certificate expires. An optional **repeat** *<repeat-hour>* parameter will enable the system to repeat the **BeforeExp** message every hour until the certificate expires.

If the user only wants **AfterExp**, then **certificate-expiration-warning 0** can be used to achieve this.

BeforeExp and **AfterExp** warnings can be cleared in following cases:

- The certificate is reloaded by the **admin certificate reload** command. In this case, if the reloaded file is not expired, then **AfterExp** is cleared. And, if the reloaded file is outside of configured warning window, then the **BeforeExp** is also cleared.
- When the **ca-profile/ipsec-gw/ipsec-tunnel/cert-profile** is shutdown, then **BeforeExp** and **AfterExp** of corresponding certificates are cleared.
- When **no certificate-expiration-warning** command is configured, then all existing **BeforeExp** and **AfterExp** are cleared.
- Users may change the configuration of the **certificate-expiration-warning** so that certain certificates are no longer in the warning window. **BeforeExp** of corresponding certificates are cleared.
- If the system time changes so that the new time causes the certificates to no longer be in the warning window, then **BeforeExp** is cleared. If the new time causes an expired certificate to come non-expired, then **AfterExp** is cleared.

Default	no certificate-expiration-warning
Parameters	<i>hours</i> — Specifies the amount of time before a certificate expires when system issues BeforeExp . Values 0 to 8760
	repeat <i>repeat-hours</i> — The system will repeat BeforeExp every <i>repeat-hour</i> . Values 0 to 8760

common-name-list

Syntax	common-name-list <i>name</i>
Context	config>system>security>pki
Description	This command configures a list of common names (CNs) that will be used to authenticate X.509.3 certificates. If the CN field of the X.509.3 certificate matches any of the CNs in the list, then the certificate can be used.
Parameters	<i>name</i> — Specifies the name of the CN list, up to 32 characters maximum.

cn

Syntax	[no] cn <i>index</i> type <i>value</i> <i>common-name-value</i>
Context	config>system>security>pki>common-name-list
Description	This command creates a CN list entry in text or regexp format. The no form of the command removes the specified entry.
Parameters	<i>index</i> — Specifies the index number of the entry. <i>type</i> — Specifies the type of the entry. Values ip-address, domain-name <i>common-name-value</i> — Specifies the IP address or domain name value, up to 255 characters maximum.

crl-expiration-warning

Syntax	crl-expiration-warning <i>hours</i> [repeat <i>repeat-hours</i>] no crl-expiration-warning
Context	config>system>security>pki
Description	This command specifies when system will issue BeforeExp message before a CRL expires. For example, with certificate-expiration-warning 5 , the system will issue a BeforeExp message 5 hours before a CRL expires. An optional repeat <i><repeat-hour></i> parameter will enable the system to repeat the BeforeExp message every hour until the CRL expires. If the user only wants AfterExp , then certificate-expiration-warning 0 can be used to achieve this. BeforeExp and AfterExp warnings can be cleared in following cases:

- The CRL is reloaded by the **admin certificate reload** command. In this case, if the reloaded file is not expired, then **AfterExp** is cleared. And, if the reloaded file is outside of configured warning window, then the **BeforeExp** is also cleared.
- When the **ca-profile** is shutdown, then **BeforeExp** and **AfterExp** of corresponding certificates are cleared.
- When **no crl-expiration-warning** command is configured, then all existing **BeforeExp** and **AfterExp** are cleared.
- Users may change the configuration of the **crl-expiration-warning** so that certain CRL are no longer in the warning window. **BeforeExp** of corresponding CRL are cleared.
- If the system time changes so that the new time causes the CRL to no longer be in the warning window, then **BeforeExp** is cleared. If the new time causes an expired CRL to come non-expired, then **AfterExp** is cleared.

Default no crl-expiration-warning

Parameters *hours* — Specifies the amount of time before a CRL expires when system issues **BeforeExp**.

Values 0 to 8760

repeat-hour — Specifies that the system will repeat **BeforeExp** every repeat-hour.

Values 0 to 8760

maximum-cert-chain-depth

Syntax **maximum-cert-chain-depth** *level*
no maximum-cert-chain-depth

Context config>system>security>pki

Description This command defines the maximum depth of certificate chain verification. This number is applied system wide.

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

Default 7

Parameters *level* — Specifies the maximum depth level of certificate chain verification, range from 1 to 7. the certificate under verification is not counted in. for example, if this parameter is set to 1, then the certificate under verification must be directly signed by trust anchor CA.

Values 1 to 7

shutdown

Syntax [**no**] shutdown

Context config>system>security>pki>ca-profile>

Description Use this command to enable or disable the ca-profile. The system will verify the configured cert-file and crl-file. If the verification fails, then the **no shutdown** command will fail.

The ca-profile in a **shutdown** state cannot be used in certificate authentication.

Default shutdown

certificate

Syntax **certificate**

Context admin

Description This command enables the context to configure X.509 certificate related operational parameters. For information about CMPv6 admin certificate commands, see the *Multiservice Integrated Service Adapter Guide*.

clear-ocsp-cache

Syntax **clear-ocsp-cache** [*entry-id*]

Context admin>certificate

Description This command clears the current OCSP response cache. If optional issuer and serial-number are not specified, then all current cached results are cleared.

Parameters *entry-id* — Specifies the local cache entry identifier of the certificate to clear.

Values 1 to 2000

crl-update

Syntax **crl-update ca** *ca-profile-name*

Context admin>certificate

Description This command manually triggers the Certificate Revocation List file (CRL) update for the specified ca-profile.

Using this command requires shutting down the auto-crl-update.

Default None

Parameters *ca-profile-name* — Specifies the name of the Certificate Authority profile.

display

Syntax `display type {type} url-string format {format} [password [32 chars max]]`

Context admin>certificate

Description This command displays the content of an input file in plain text.



Note: When displaying the key file content, only the key size and type are displayed.

The following list summarizes the formats supported by this command:

- System
 - system format
 - PKCS #12
 - PKCS #7 PEM encoded
 - PKCS #7 DER encoded
 - RFC4945
- Certificate Request
 - PKCS #10
- Key
 - system format
 - PKCS #12
- CRL
 - system format
 - PKCS #7 PEM encoded
 - PKCS #7 DER encoded
 - RFC4945

Default none

Parameters *file-url* — Specifies the local CF card url of the input file.

Values

url-string	<local-url> - [99 chars max]
local-url	<cflash-id>/<file-path>
cflash-id	cf1: cf2: cf3:

type — Specifies the type of input file, possible values are cert/key/crl/cert-request.

Values cert, key, crl, cert-request

format — Specifies the format of input file.

Values pkcs10, pkcs12, pkcs7-der, pkcs7-pem, pem, der

password — Specifies the password to decrypt the input file in case that it is a encrypted PKCS#12 file, up to 99 characters in length.

export

Syntax **export type** {*type*} **input** *filename* **output** *url-string* **format** *output-format* [**password** [32 chars max]] [**pkey** *filename*]

Context admin>certificate

Description This command performs certificate operations.

gen-keypair

Syntax **gen-keypair** *url-string* [**size** {512 | 1024 | 2048}] [**type** {rsa | dsa}]

Context admin>certificate

Description This command generates a RSA or DSA private key/public key pairs and store them in a local file in cf3:\system-pki\key

Parameters *url-string* — Specifies the name of the key file.

Values

url-string	<local-url> - [99 chars max]
local-url	<cf-flash-id>/<file-path>
cf-flash-id	cf1: cf2: cf3:

size — Specifies the key size in bits.
The minimum key-size is 1024 when running in FIPS-140-2 mode.

Values 512/1024/2048

Default 2048

type — Specifies the type of key.

Default rsa

gen-local-cert-req

Syntax **gen-local-cert-req** **keypair** *url-string* **subject-dn** *subject-dn* [**domain-name** [255 chars max]] [**ip-addr** *ip-address*] **file** *url-string* [**hash-alg** *hash-algorithm*]

Context	admin>certificate
Description	This command generates a PKCS#10 formatted certificate request by using a local existing key pair file.
Default	none
Parameters	<i>url-string</i> — Specifies the name of the keyfile in cf3:\system-pki\key that is used to generate a certificate request.

Values

url-string	<local-url> - [99 chars max]
local-url	<cf-flash-id>/<file-path>
cf-flash-id	cf1: cf2: cf3:

subject-dn — Specifies the distinguish name that is used as the subject in a certificate request, including:

- C-Country
- ST-State
- O-Organization name
- OU-Organization Unit name
- CN-common name

This parameter is formatted as a text string including any of the above attributes. The attribute and its value is linked by using "=", and "," is used to separate different attributes.

For example: C=US,ST=CA,O=ALU,CN=SR12

Values attr1=val1,attr2=val2... where: attrN={C|ST|O|OU|CN}, 256 chars max

domain-name — Optionally, a domain name string can be specified and included as the dNSName in the Subject Alternative Name extension of the certificate request.

ip-address — Optionally, an IPv4 address string can be specified and included as the ipAddress in the Subject Alternative Name extension of the certificate request.

cert-req-file-url — This URL could be either a local CF card path and filename to save the certificate request; or an FTP URL to upload the certificate request.

hash-alg *hash-algorithm* — Specifies the hash algorithm to be used in a certificate request.

Values sha1, sha224, sha256, sha384, sha512

import

Syntax **import type {cert | key | crl} input url-string output filename format input-format [password [32 chars max]]**

Context admin>certificate#

Description This command converts an input file(key/certificate/CRL) to a system format file. The following list summarizes the formats supported by this command:

- Certificate
 - PKCS #12
 - PKCS #7 PEM encoded
 - PKCS #7 DER encoded
 - PEM
 - DER
- Key
 - PKCS #12
 - PEM
 - DER
- CRL
 - PKCS #7 PEM encoded
 - PKCS #7 DER encoded
 - PEM
 - DER



Note: If there are multiple objects with the same type in the input file, only the first object will be extracted and converted.

Default none

Parameters **input** *url-string* — Specifies the URL for the input file. This URL could be either a local CF card URL file or a FP URL to download the input file.

output *url-string* — Specifies the name of output file up to 95 characters in length. The output directory depends on the file type like following:

- Key: cf3:\system-pki\key
- Cert: cf3:\system-pki\cert
- CRL: cf3:\system-pki\CRL

Values

url-string	<local-url> - [99 chars max]
local-url	<cflash-id>/<file-path>
cflash-id	cf1: cf2: cf3:

type — The type of input file.

Values cert, key, crl

format — Specifies the format of input file.

Values pkcs12, pkcs7-der, pkcs7-pem, pem, der

password — Specifies the password to decrypt the input file in case that it is a encrypted PKCS#12 file.

reload

Syntax	reload type {cert key cert-key-pair} filename [key-file filename]
Context	admin>certificate
Description	<p>This command reloads imported certificate or key file or both at the same time. This command is typically used to update certificate/key file without shutting down ipsec-tunnel/ipsec-gw/cert-profile/ca-profile. Note that type cert and type key will be deprecated in a future release. Use type cert-key-pair instead. Instead of type cert use type key instead.</p> <ul style="list-style-type: none"> • If the new file exists and valid, then for each tunnel using it: <ul style="list-style-type: none"> – If the key matches the certificate, then the new file will be downloaded to the MS-ISA to be used the next time. Tunnels currently up are not affected. – If the key does not match the certificate: <ul style="list-style-type: none"> • If cert and key configuration is used instead of cert-profile then the tunnel will be brought down. • If cert-profile is used, then cert-profile will be brought down. The next authentication will fail while the established tunnels are not affected. <p>If the new file does not exists or somehow invalid (bad format, does not contain right extension, etc.), then this command will abort.</p> <p>In the case of type cert-key-pair, if the new file doesn't exist or is invalid or cert and key do not match, then this command will abort with an error message.</p>
Default	none
Parameters	<p>cert — Specifies to reload a certificate file.</p> <p>key — Specifies to reload a key file.</p> <p>cert-key-pair — Specifies to reload a certificate file and its key file at the same time.</p> <p><i>file-name</i> — Specifies the file name of imported certificate or key.</p> <p><i>key-filename</i> — In case of cert-key-pair, filename is the imported filename of certificate, key-filename is the imported key file.</p>

secure-nd-export

Syntax **secure-nd-export**

Context admin>certificate

Description This command exports IPv6 Secure Neighbor Discovery (SeND) certificates to the file cf[1..3]:\system-pki\secureNdKey in PKCS #7 DER format.

secure-nd-import

Syntax **secure-nd-import** **input** *url-string* **format** *input-format* [**password** *password*] [**key-rollover**]

Context admin>certificate

Description This command imports IPv6 Secure Neighbor Discovery (SeND) certificates from a file, and saves them to cf[1..3]:\system-pki\secureNdKey in PKCS #7 DER format.

Parameters *url-string* — Specifies the name of an input file up to 99 characters in length.

Values

local-url	<cf-flash-id>\<file-path>
cf-flash-id	cf1: cf2: cf3:

input-format — Specifies the input file format.

Values pkcs12, pem, or der

password — Specifies the password to decrypt the input file if it is an encrypted PKCS#12 file.

Values 32 characters maximum

2.19.2.7 Profile Management Commands

action

Syntax **action** {**deny** | **permit** | **read-only**}

Context config>system>security>profile>entry

Description This command configures the action associated with the profile entry.

Parameters **deny** — Specifies that commands matching the entry command match criteria are to be denied.

permit — Specifies that commands matching the entry command match criteria will be permitted.

match

Syntax	match <i>command-string</i> no match
Context	config>system>security>profile>entry
Description	<p>This command configures a command or subtree commands in subordinate command levels are specified.</p> <p>Because the OS exits when the first match is found, subordinate levels cannot be modified with subsequent action commands. More specific action commands should be entered with a lower entry number or in a profile that is evaluated prior to this profile.</p> <p>All commands below the hierarchy level of the matched command are denied.</p> <p>The no form of this command removes a match condition</p>
Default	none
Parameters	<i>command-string</i> — Specifies the CLI command or CLI tree level that is the scope of the profile entry.

copy

Syntax	copy { user <i>source-user</i> <i>profile source-profile</i> } to <i>destination</i> [overwrite]
Context	config>system>security
Description	This command copies a profile or user from a source profile to a destination profile.
Parameters	<p><i>source-profile</i> — Specifies the profile to copy. The profile must exist.</p> <p><i>dest-profile</i> — Specifies the copied profile is copied to the destination profile.</p> <p>overwrite — Specifies that the destination profile configuration will be overwritten with the copied source profile configuration. A profile will not be overwritten if the overwrite command is not specified.</p>

default-action

Syntax	default-action { deny-all permit-all none read-only-all }
Context	config>system>security>profile
Description	This command specifies the default action to be applied when no match conditions are met.
Default	none
Parameters	deny-all — Sets the default of the profile to deny access to all commands.

permit-all — Sets the default of the profile to permit access to all commands.



Note: The **permit-all** parameter does not change access to security commands. Security commands are only and always available to members of the super-user profile.

none — Sets the default of the profile to no-action. This option is useful to assign multiple profiles to a user.

For example, if a user is a member of two profiles and the default action of the first profile is **permit-all**, then the second profile will never be evaluated because the **permit-all** is executed first. Set the first profile default action to **none** and if no match conditions are met in the first profile, then the second profile will be evaluated. If the default action of the last profile is **none** and no explicit match is found, then the default **deny-all** takes effect.

entry

Syntax	[no] entry <i>entry-id</i>
Context	config>system>security>profile
Description	<p>This command is used to create a user profile entry.</p> <p>More than one entry can be created with unique <i>entry-id</i> numbers. Exits when the first match is found and executes the actions according to the accompanying action command. Entries should be sequenced from most explicit to least explicit.</p> <p>An entry may not have any match criteria defined (in which case, everything matches) but must have at least the keyword action for it to be considered complete.</p> <p>The no form of the command removes the specified entry from the user profile.</p>
Default	No entry IDs are defined.
Parameters	<p><i>entry-id</i> — Specifies an entry-id that uniquely identifies a user profile command match criteria and a corresponding action. If more than one entry is configured, the <i>entry-ids</i> should be numbered in staggered increments to allow users to insert a new entry without requiring renumbering of the existing entries.</p> <p>Values 1 to 9999</p>

profile

Syntax	[no] profile <i>user-profile-name</i>
Context	config>system>security

-
- Description** This command creates a context to create user profiles for CLI command tree permissions. Profiles are used to either deny or permit user console access to a hierarchical branch or to specific commands.
- Once the profiles are created, the `user` command assigns users to one or more profiles. You can define up to 16 user profiles but a maximum of 8 profiles can be assigned to a user. The *user-profile-name* can consist of up to 32 alphanumeric characters.
- The `no` form of the command deletes a user profile.
- Default** user-profile default
- Parameters** *user-profile-name* — Specifies the user profile name entered as a character string. The string is case sensitive and limited to 32 ASCII 7-bit printable characters with no spaces.

renum

- Syntax** `renum old-entry-number new-entry-number`
- Context** config>system>security>profile
- Description** This command renumbers profile entries to re-sequence the entries.
- Since the OS exits when the first match is found and executes the actions according to accompanying action command, re-numbering is useful to rearrange the entries from most explicit to least explicit.
- Parameters** *old-entry-number* — Enter the entry number of an existing entry.
- Values** 1 to 9999
- new-entry-number* — Enter the new entry number.
- Values** 1 to 9999

2.19.2.8 User Management Commands

access

Syntax	[no] access [ftp] [snmp] [console] [li] [netconf] [grpc]
Context	config>system>security>user config>system>security>user-template
Description	<p>This command grants a user permission for FTP, SNMP, console, lawful intercept (LI), NETCONF, or gRPC access.</p> <p>If a user requires access to more than one application, then multiple applications can be specified in a single command. Multiple commands are treated additively.</p> <p>The no form of this command removes access for a specific application, and denies permission for all management access methods. To deny a single access method, enter the no form of the command followed by the method to be denied, for example, no access FTP denies FTP access.</p>
Default	no access
Parameters	<p>ftp — Specifies FTP permission.</p> <p>snmp — Specifies SNMP permission. This keyword is only configurable in the config>system>security>user context.</p> <p>console — Specifies console access (serial port or Telnet) permission.</p> <p>li — Specifies CLI command access in the lawful intercept (LI) context (applies to the 7450 ESS and 7750 SR).</p> <p>netconf — Specifies NETCONF session access for the user defined in the specified user context. When using the Base-R13 SR OS YANG data model, console access is also necessary (not required for the Nokia SR OS YANG data model).</p> <p>grpc — Specifies gRPC access.</p>

authentication

Syntax	authentication {[none] [[hash] {md5 key-1 sha key-1} privacy {none des-key aes-128-cfb-key key-2}]}
Context	config>system>security>user>snmp
Description	<p>This command configures the authentication and encryption method the user must use in order to be validated by the router. SNMP authentication allows the device to validate the managing node that issued the SNMP message and determine if the message has been tampered.</p>

The keys configured in this command must be localized keys (MD5 or DES hash of the configured SNMP engine-ID and a password). The password is not directly entered in this command (only the localized key).

Default	authentication none
Parameters	<p>none — Do not use authentication. If none is specified, then privacy cannot be configured.</p> <p>hash — When hash is not specified, then non-encrypted characters can be entered. When hash is configured, then all specified keys are stored in an encrypted format in the configuration file. The key must be entered in encrypted form when the hash parameter is used.</p> <p>md5 key — Use an HMAC-MD5-96 authentication key. The MD5 authentication key is stored in an encrypted format. The key must be entered as a full 32 hex character string.</p> <p>sha key — Use an HMAC-SHA-96 authentication key. The sha authentication key is stored in an encrypted format. The key must be entered as a full 40 hex character string.</p> <p>privacy none — Do not perform SNMP packet encryption.</p> <p>Default privacy none</p> <p>privacy des-key key-2 — Use DES for SNMP payload encryption and configure the key. The key must be a 32 hex-character string and is stored in an encrypted format. The des-key parameter is not available in FIPS-140-2 mode.</p> <p>privacy aes-128-cfb-key key-2 — Use 128 bit CFB mode AES for SNMP payload encryption and configure the key. The key must be a 32 hex-character string and is stored in an encrypted format.</p> <p>Default privacy none</p>

group

Syntax	group <i>group-name</i> no group
Context	config>system>security>user>snmp
Description	This command associates (or links) a user to a group name. The group name must be configured with the config>system>security>user >snmp>group command. The access command links the group with one or more views, security model (s), security level (s), and read, write, and notify permissions
Default	No group name is associated with a user.

Parameters *group-name* — Enter the group name (between 1 and 32 alphanumeric characters) that is associated with this user. A user can be associated with one group-name per security model.

cannot-change-password

Syntax **[no] cannot-change-password**

Context config>system>security>user>console

Description This command allows a user the privilege to change their password for both FTP and console login.

To disable a user's privilege to change their password, use the **cannot-change-password** form of the command.



Note: The **cannot-change-password** flag is not replicated when a user copy is performed. A new-password-at-login flag is created instead.

Default no cannot-change-password

console

Syntax **console**

Context config>system>security>user
config>system>security>user-template

Description This command creates the context to configure user profile membership for the console (either Telnet or CPM serial port user).

copy

Syntax **copy {user *source-user* | profile *source-profile*} to *destination* [overwrite]**

Context config>system>security

Description This command copies a specific user's configuration parameters to another (destination) user.

The password is set to a carriage return and a new password at login must be selected.

Parameters *source-user* — Specifies the user to copy. The user must already exist.
dest-user — Specifies that the copied profile is copied to a destination user.

overwrite — Specifies that the destination user configuration will be overwritten with the copied source user configuration. A configuration will not be overwritten if the **overwrite** command is not specified.

home-directory

Syntax **home-directory** *url-prefix* [*directory*] [*directory/directory...*]
no home-directory

Context config>system>security>user
config>system>security>user-template

Description This command configures the local home directory for the user for both console (file commands and '>' redirection) and FTP access.

If the URL or the specified URL/directory structure is not present, then a warning message is issued and the default is assumed.

The **no** form of the command removes the configured home directory.

Default no home-directory



Note: If restrict-to-home has been configured no file access is granted and no home-directory is created. If restrict-to-home is not applied then root becomes the user's home-directory.

Parameters *local-url-prefix* [*directory*] [*directory/directory...*] — Specifies the user's local home directory URL prefix and directory structure up to 190 characters in length.

profile

Syntax **profile** *user-profile-name*
no profile

Context config>system>security>user-template

Description This command configures the profile for the user based on this template.

Parameters *user-profile-name* — The user profile name entered as a character string. The string is case sensitive and limited to 32 ASCII 7-bit printable characters with no spaces.

login-exec

Syntax [**no**] **login-exec** *url-prefix: source-url*

Context	config>system>security>user>console config>system>security>user-template>console
Description	This command configures a user's login exec file which executes whenever the user successfully logs in to a console session. Only one exec file can be configured. If multiple login-exec commands are entered for the same user, each subsequent entry overwrites the previous entry. The no form of the command disables the login exec file for the user.
Default	no login-exec
Parameters	<i>url-prefix: source-url</i> — Enter either a local or remote URL, up to 200 characters in length, that identifies the exec file that will be executed after the user successfully logs in.

member

Syntax	member <i>user-profile-name</i> [<i>user-profile-name.....(up to 8 max)</i>] no member <i>user-profile-name</i>
Context	config>system>security>user>console
Description	This command is used to allow the user access to a profile. A user can participate in up to eight profiles. The no form of this command deletes access user access to a profile.
Default	default
Parameters	<i>user-profile-name</i> — The user profile name up to 32 characters in length.

new-password-at-login

Syntax	[no] new-password-at-login
Context	config>system>security>user>console
Description	This command forces the user to change a password at the next console login. The new password applies to FTP but the change can be enforced only by the console, SSH, or Telnet login. The no form of the command does not force the user to change passwords.
Default	no new-password-at-login

password

Syntax	password [<i>password</i>]
Context	config>system>security>user
Description	This command configures the user password for console and FTP access.

The password is stored in an encrypted format in the configuration file when specified. Passwords should be encased in double quotes (" ") at the time of the password creation. The double quote character (") is not accepted inside a password. It is interpreted as the start or stop delimiter of a string.

The password can be entered as plain text or a hashed value. SR OS can distinguish between hashed passwords and plain text passwords and take the appropriate action to store the password correctly.

```
config>system>security>user# password testuser1
```

The password is hashed by default.

For example:

```
config>system>security# user testuser1
config>system>security>user$ password xyzabcd1
config>system>security>user# exit

config>system>security# info
-----
...
        user "testuser1"
            password "$2y$10$pFoehOg/tCbBMPDJ/
kqpu.8af0AoVGY2xsR7WFqyn5fVTnwRzGmOK"
            exit
...
-----
config>system>security#
```

The **password** command allows you also to enter the password as a hashed value.

For example:

```
config>system>security# user testuser1
config>system>security>user$ password "$2y$10$pFoehOg/tCbBMPDJ/
kqpu.8af0AoVGY2xsR7WFqyn5fVTnwRzGmOK"
config>system>security>user# exit
config>system>security# info
-----
...
user "testuser1"
password "$2y$10$pFoehOg/tCbBMPDJ/kqpu.8af0AoVGY2xsR7WFqyn5fVTnwRzGmOK"
exit
...
-----
```

```
-----  
config>system>security#
```

Parameters *password* — This is the password for the user that must be entered by this user during the login procedure. The minimum length of the password is determined by the **minimum-length** command. The maximum length can be up to 20 chars if unhashed, 32 characters if hashed. The complexity requirements for the password is determined by the **complexity** command.

A password value that does not conform to the minimum-length or other password complexity rules can be configured using the **config>system>security>user>password** command, but a warning is provided in the CLI. This allows, for example, an administrator to configure a non-conformant password for a user. A user cannot configure a non-conformant password for themselves using the global **password** command.

All password special characters (#, \$, spaces, etc.) must be enclosed within double quotes.

For example: config>system>security>user# password "south#bay?"

The question mark character (?) cannot be directly inserted as input during a telnet connection because the character is bound to the **help** command during a normal Telnet/console connection.

To insert a # or ? characters, they must be entered inside a notepad or clipboard program and then cut and pasted into the Telnet session in the password field that is encased in the double quotes as delimiters for the password.

If a **password** is entered without any parameters, a password length of zero is implied: (carriage return).

public-keys

Syntax **public-keys**

Context config>system>security>user

Description This command allows the user to enter the context to configure public keys for SSH.

ecdsa

Syntax **ecdsa**

Context config>system>security>user>public-keys

Description This command allows the user to enter the context to configure ECDSA public keys.

ecdsa-key

Syntax	ecdsa-key <i>key-id</i> [create] no ecdsa-key <i>key-id</i>
Context	config>system>security>user>public-keys>ecdsa
Description	This command creates an ECDSA public key and associates it with the username. Multiple public keys can be associated with the user. The key ID is used to identify these keys for the user.
Parameters	create — Keyword used to create an ECDSA key. The create keyword requirement can be enabled/disabled in the environment>create context. <i>key-id</i> — Specifies the key identifier. Values 1 to 32

key-value

Syntax	key-value <i>public-key-value</i> no key-value
Context	config>system>security>user>public-keys>ecdsa>ecdsa-key config>system>security>user>public-keys>rsa>rsa-key
Description	This command configures a value for the RSA or ECDSA public key. The public key must be enclosed in quotation marks. For RSA, the key is between 768 and 4096 bits. For ECDSA, the key is between 1 and 1024 bits.
Default	no key-value
Parameters	<i>public-key-value</i> — Specifies the public key value up to 800 characters in length for RSA and up to 255 characters in length for ECDSA.

rsa

Syntax	rsa
Context	config>system>security>user>public-keys
Description	This command allows the user to enter the context to configure RSA public keys.

rsa-key

Syntax	rsa-key <i>key-id</i> [create] no rsa-key <i>key-id</i>
---------------	---

Context	config>system>security>user>public-keys>rsa
Description	This command creates an RSA public key and associates it with the username. Multiple public keys can be associated with the user. The key ID is used to identify these keys for the user.
Parameters	<p>create — Keyword used to create the RSA key. The create keyword requirement can be enabled/disabled in the environment>create context.</p> <p><i>key-id</i> — Specifies the key identifier.</p> <p>Values 1 to 32</p>

restricted-to-home

Syntax	[no] restricted-to-home
Context	config>system>security>user config>system>security>user-template
Description	<p>This command prevents users from navigating above their home directories for file access (either by means of CLI sessions with the file command, '>' redirection, or by means of FTP). A user is not allowed to navigate to a directory higher in the directory tree on the home directory device. The user is allowed to create and access subdirectories below their home directory.</p> <p>If a home-directory is not configured or the home directory is not available, then the user has no file access.</p> <p>The no form of the command allows the user access to navigate to directories above their home directory.</p>
Default	no restricted-to-home

snmp

Syntax	snmp
Context	config>system>security>user
Description	<p>This command creates the context to configure SNMP group membership for a specific user and defines encryption and authentication parameters.</p> <p>All SNMPv3 users must be configured with the commands available in this CLI node.</p> <p>The OS always uses the configured SNMPv3 user name as the security user name.</p>

user-template

Syntax	user-template { tacplus_default radius_default ldap-default }
Context	config>system>security
Description	This command configures default security user template parameters.
Parameters	<p>tacplus_default — Specifies the default TACACS+ user template. All parameters of the tacplus_default template except the “profile” are actively applied to all TACACS+ users if tacplus use-default-template is enabled. The “profile” parameters are applied to all TACACS+ users if tacplus authorization is enabled (without the use-priv-lvl option) and tacplus use-default-template is enabled.</p> <p>radius_default — Specifies the default RADIUS user template. The radius_default template is actively applied to a RADIUS user if radius authorization is enabled, radius use-default-template is enabled, and no VSAs are returned with the auth-accept from the RADIUS server.</p> <p>ldap_default — Specifies the default LDAP user template.</p>

user

Syntax	[no] user <i>user-name</i>
Context	config>system>security
Description	<p>This command creates a local user and a context to edit the user configuration.</p> <p>If a new <i>user-name</i> is entered, the user is created. When an existing <i>user-name</i> is specified, the user parameters can be edited.</p> <p>When creating a new user and then entering the info command, the system displays a password in the output. This is expected behavior in the hash2 scenario. However, when using that user name, there will be no password required. The user can login to the system and then <ENTER> at the password prompt, the user will be logged in.</p> <p>Unless an administrator explicitly changes the password, it will be null. The hashed value displayed uses the username and null password field, so when the username is changed, the displayed hashed value will change.</p> <p>The no form of the command deletes the user and all configuration data. Users cannot delete themselves.</p>
Default	none
Parameters	<i>user-name</i> — Specifies the name of the user up to 32 characters.

2.19.2.9 CLI Session Management Commands

cli-session-group

- Syntax** [no] **cli-session-group** *session-group-name* [**create**]
- Context** config>system>security
- Description** This command is used to configure a session group that can be used to limit the number of CLI sessions available to members of the group.
- Parameters** *session-group-name* — Specifies a particular session group.

ssh-max-sessions

- Syntax** **ssh-max-sessions** *session-limit*
no ssh-max-sessions
- Context** config>system>security>cli-session-group
config>system>security>profile
- Description** This command is used to limit the number of SSH-based CLI sessions available to all users that are part of a particular profile, or to all users of all profiles that are part of the same cli-session-group.
- The **no** form of this command disables the command and the profile/group limit is not applied on the number of sessions.
- Default** no ssh-max-sessions
- Parameters** *session-limit* — Specifies the maximum number of allowed SSH-based CLI sessions.
- Values** 0 to 50

telnet-max-sessions

- Syntax** **telnet-max-sessions** *session-limit*
no telnet-max-sessions
- Context** config>system>security>cli-session-group
config>system>security>profile
- Description** This command is used to limit the number of Telnet-based CLI sessions available to all users that are part of a particular profile, or to all users of all profiles that are part of the same cli-session-group.

The **no** form of this command disables the command and the profile/group limit is not applied on the number of sessions.

Default no telnet-max-sessions

Parameters *session-limit* — Specifies the maximum number of allowed Telnet-based CLI sessions.

Values 0 to 50

combined-max-sessions

Syntax **combined-max-sessions** *session-limit*
no combined-max-sessions

Context config>system>security>cli-session-group
config>system>security>profile

Description This command is used to limit the number of combined SSH/TELENT based CLI sessions available to all users that are part of a particular profile, or to all users of all profiles that are part of the same cli-session-group.

The **no** form of this command disables the command and the profile/group limit is not applied to the number of combined sessions.

Default no combined-max-sessions

Parameters *session-limit* — Specifies the maximum number of allowed combined SSH/TELNET based CLI sessions.

Values 0 to 50

2.19.2.10 RADIUS Client Commands

access-algorithm

Syntax	access-algorithm { direct round-robin } no access-algorithm
Context	config>system>security>radius
Description	This command indicates the algorithm used to access the set of RADIUS servers.
Default	direct
Parameters	direct — The first server will be used as primary server for all requests, the second as secondary and so on. round-robin — The first server will be used as primary server for the first request, the second server as primary for the second request, and so on. If the router gets to the end of the list, it starts again with the first server.

accounting

Syntax	[no] accounting
Context	config>system>security>radius
Description	This command enables RADIUS accounting. The no form of this command disables RADIUS accounting.
Default	no accounting

accounting-port

Syntax	accounting-port <i>port</i> no accounting-port
Context	config>system>security>radius
Description	This command specifies a UDP port number on which to contact the RADIUS server for accounting requests.
Parameters	<i>port</i> — Specifies the UDP port number. Values 1 to 65535 Default 1813

authorization

Syntax	[no] authorization
Context	config>system>security>radius
Description	This command configures RADIUS authorization parameters for the system.
Default	no authorization

interactive-authentication

Syntax	[no] interactive-authentication
Context	config>system>security>radius
Description	This command enables RADIUS interactive authentication for the system. Enabling interactive-authentication forces RADIUS to fall into challenge/response mode.
Default	no authentication

port

Syntax	port <i>port</i> no port
Context	config>system>security>radius
Description	This command configures the TCP port number to contact the RADIUS server. The no form of the command reverts to the default value.
Default	1812 (as specified in RFC 2865, <i>Remote Authentication Dial In User Service (RADIUS)</i>)
Parameters	<i>port</i> — The TCP port number to contact the RADIUS server. Values 1 to 65535

radius

Syntax	[no] radius
Context	config>system>security
Description	This command creates the context to configure RADIUS authentication on the router. Implement redundancy by configuring multiple server addresses for each router.

The **no** form of the command removes the RADIUS configuration.

retry

Syntax	retry <i>count</i> no retry
Context	config>system>security>radius config>system>security>dot1x>radius-plcy
Description	This command configures the number of times the router attempts to contact the RADIUS server for authentication if there are problems communicating with the server. The no form of the command reverts to the default value.
Default	3
Parameters	<i>count</i> — Specifies the retry count. Values 1 to 10

server

Syntax	server <i>index</i> address <i>ip-address</i> secret <i>key</i> [hash hash2] no server <i>index</i>
Context	config>system>security>radius
Description	This command adds a RADIUS server and configures the RADIUS server IP address, index, and key values. Up to five RADIUS servers can be configured at any one time. RADIUS servers are accessed in order from lowest to highest index for authentication requests until a response from a server is received. A higher indexed server is only queried if no response is received from a lower indexed server (which implies that the server is not available). If a response from a server is received, no other RADIUS servers are queried. It is assumed that there are multiple identical servers configured as backups and that the servers do not have redundant data. The no form of the command removes the server from the configuration.
Default	no server
Parameters	<i>index</i> — The index for the RADIUS server. The index determines the sequence in which the servers are queried for authentication requests. Servers are queried in order from lowest to highest index. Values 1 to 5

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

Values

ipv4-address	a.b.c.d (host bits must be 0)
ipv6-address	x:x:x:x:x:x:x (eight 16-bit pieces) x:x:x:x:x:d.d.d.d x: [0..FFFF]H d: [0..255]D

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

Values Up to 128 characters in length.

hash — Specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified

hash2 — Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

shutdown

Syntax	[no] shutdown
Context	config>system>security>radius
Description	This command administratively disables the RADIUS protocol operation. Shutting down the protocol does not remove or change the configuration other than the administrative state. The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they may be deleted. The no form of the command administratively enables the protocol which is the default state.
Default	no shutdown

timeout

Syntax	timeout <i>seconds</i> no timeout
---------------	--

Context	config>system>security>radius
Description	This command configures the number of seconds the router waits for a response from a RADIUS server. The no form of the command reverts to the default value.
Default	3
Parameters	<i>seconds</i> — Specifies the number of seconds the router waits for a response from a RADIUS server, expressed as a decimal integer. Values 1 to 90

use-default-template

Syntax	[no] use-default-template
Context	config>system>security>radius
Description	This command specifies whether the RADIUS default user template is actively applied to the RADIUS user if no VSAs are returned with the auth-accept from the RADIUS server. When enabled, the radius_default user-template is actively applied if no VSAs are returned with the auth-accept from the RADIUS server and radius authorization is enabled. The no form of the command disables the use of the RADIUS default template.
Default	no use-default-template

2.19.2.11 TACACS+ Client Commands

server

- Syntax** `server index address ip-address secret key [hash | hash2][port port]`
`no server index`
- Context** config>system>security>tacplus
- Description** This command adds a TACACS+ server and configures the TACACS+ server IP address, index, and key values.
- Up to five TACACS+ servers can be configured at any one time. TACACS+ servers are accessed in order from lowest index to the highest index for authentication requests.
- The **no** form of the command removes the server from the configuration.
- Default** No TACACS+ servers are configured.
- Parameters** *index* — The index for the TACACS+ server. The index determines the sequence in which the servers are queried for authentication requests. Servers are queried in order from the lowest index to the highest index.
- Values** 1 to 5
- address** *ip-address* — The IP address of the TACACS+ server. Two TACACS+ servers cannot have the same IP address. An error message is generated if the server address is a duplicate.
- Values**
- | | |
|--------------|-------------------------------------|
| ipv4-address | a.b.c.d (host bits must be 0) |
| ipv6-address | x:x:x:x:x:x:x (eight 16-bit pieces) |
| | x:x:x:x:x:d.d.d.d |
| | x: [0..FFFF]H |
| | d: [0..255]D |
- secret** *key* — The secret key to access the RADIUS server. This secret key must match the password on the RADIUS server.
- Values** Up to 128 characters in length.
- hash** — Specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified

hash2 — Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

port *port* — Specifies the port ID.

Values 0 to 65535

shutdown

Syntax [no] shutdown

Context config>system>security>tacplus

Description This command administratively disables the TACACS+ protocol operation. Shutting down the protocol does not remove or change the configuration other than the administrative state.

The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they may be deleted.

The **no** form of the command administratively enables the protocol which is the default state.

Default no shutdown

tacplus

Syntax [no] tacplus

Context config>system>security

Description This command creates the context to configure TACACS+ authentication on the router.

Configure multiple server addresses for each router for redundancy.

The **no** form of the command removes the TACACS+ configuration.

accounting

Syntax accounting [record-type {start-stop | stop-only}]
no accounting

Context config>system>security>tacplus

Description This command configures the type of accounting record packet that is to be sent to the TACACS+ server. The **record-type** parameter indicates whether TACACS+ accounting start and stop packets be sent or just stop packets be sent.

Default record-type stop-only

Parameters **record-type start-stop** — Specifies that a TACACS+ start packet is sent whenever the user executes a command.

record-type stop-only — Specifies that a stop packet is sent whenever the command execution is complete.

authorization

Syntax [no] authorization [use-priv-lvl]

Context config>system>security>tacplus

Description This command configures TACACS+ authorization parameters for the system.

Default no authorization

Parameters *use-priv-lvl* — Automatically performs a single authorization request to the TACACS+ server for cmd* (all commands) immediately after login, and then use the local profile associated (via the priv-lvl-map) with the priv-lvl returned by the TACACS+ server for all subsequent authorization (except enable-admin). After the initial authorization for cmd*, no further authorization requests will be sent to the TACACS+ server (except enable-admin).

interactive-authentication

Syntax [no] interactive-authentication

Context config>system>security>tacplus

Description This configuration instructs the SR OS to send no username nor password in the TACACS+ start message, and to display the *server_msg* in the GETUSER and GETPASS response from the TACACS+ server. Interactive authentication can be used to support a One Time Password scheme (e.g. S/Key). An example flow (e.g. with a telnet connection) is as follows:

- The SR OS will send an authentication start request to the TACACS+ server with no username nor password.
- TACACS+ server replies with TAC_PLUS_AUTHEN_STATUS_GETUSER and a *server_msg*.
- The SR OS displays the *server_msg*, and collects the user name.
- The SR OS sends a continue message with the user name.
- TACACS+ server replies with TAC_PLUS_AUTHEN_STATUS_GETPASS and a *server_msg*.
- The SR OS displays the *server_msg* (which may contain, for example, an S/Key for One Time Password operation), and collects the password.
- The SR OS sends a continue message with the password.

- TACACS+ server replies with PASS or FAIL.

When interactive-authentication is disabled the SR OS will send the username and password in the *tacplus* start message. An example flow (e.g. with a telnet connection) is as follows:

- TAC_PLUS_AUTHEN_TYPE_ASCII.
 - the login username in the “user” field.
 - the password in the *user_msg* field (while this is non-standard, it does not cause interoperability problems).
- TACACS+ server ignores the password and replies with TAC_PLUS_AUTHEN_STATUS_GETPASS.
- The SR OS sends a continue packet with the password in the *user_msg* field.
- TACACS+ server replies with PASS or FAIL.

When interactive-authentication is enabled, *tacplus* must be the first method specified in the authentication-order configuration.

Default no interactive-authentication

priv-lvl-map

Syntax **[no] priv-lvl-map**

Context config>system>security>tacplus

Description This command enables the context to specify a series of mappings between TACACS+ priv-lvl and locally configured profiles for authorization. These mappings are used when the *use-priv-lvl* option is specified for *tacplus* authorization.

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

Default priv-lvl-map

priv-lvl

Syntax **priv-lvl** *priv-lvl user-profile-name*
no priv-lvl *priv-lvl*

Context config>system>security>tacplus>priv-lvl-map

Description This command maps a specific TACACS+ priv-lvl to a locally configured profile for authorization. This mapping is used when the **use-priv-lvl** option is specified for TACPLUS authorization.

Parameters *priv-lvl* — Specifies the privilege level used when sending a TACACS+ ENABLE request.
Values 0 to 15
user-profile-name — Specifies the user profile for this mapping.

timeout

Syntax **timeout** *seconds*
no timeout

Context config>system>security>tacplus

Description This command configures the number of seconds the router waits for a response from a TACACS+ server.
 The **no** form of the command reverts to the default value.

Default 3

Parameters *seconds* — The number of seconds the router waits for a response from a TACACS+ server, expressed as a decimal integer.
Values 1 to 90

shutdown

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

Context config>system>security>tacplus

Description This command administratively disables the TACACS+ protocol operation. Shutting down the protocol does not remove or change the configuration other than the administrative state.
 The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they may be deleted.
 The **no** form of the command administratively enables the protocol which is the default state.

Default no shutdown

use-default-template

Syntax [**no**] **use-default-template**

Context config>system>security>tacplus

Description This command specifies whether the `tacplus_default` user-template is actively applied to the TACACS+ user. When enabled, the `tacplus_default` user-template is actively applied if `tacplus` authorization is enabled (without the `use-priv-lvl` option).

Default `use-default-template`

2.19.2.12 LDAP Client Commands

ldap

Syntax `[no] ldap`

Context `config>system>security`

Description This command configures LDAP authentication parameters for the system.

The **no** form will de-configure the LDAP client from the SR OS.

public-key-authentication

Syntax `[no] public-key-authentication`

Context `config>system>security>ldap`

Description This command enables public key retrieval from the LDAP server. If disabled (in its **no** form), password authentication will be attempted via LDAP.

Default `no public-key-authentication`

retry

Syntax `retry value`
`no retry`

Context `config>system>security>ldap`

Description This command configures the number of retries for the SR OS in its attempt to reach the current LDAP server before attempting the next server.

The **no** version of this command will configure a default **retry** value of 3.

Parameters `value` — 1 to 10

Default 3

server

- Syntax** **server** *server-index* [**create**]
no server *server-index*
- Context** config>system>security>ldap
- Description** This command configures an LDAP server. Up to five servers can be configured, which can then work in a redundant manner.

The **no** version of this command removes the server connection.
- Parameters** *server-index* — 1 to 5

address

- Syntax** **address** *ip-address* [**port** *port*]
no address
- Context** config>system>security>ldap>server
- Description** This command configures the IPv4 or IPv6 address for the LDAP server.

The **no** version of this command removes the server address.
- Parameters** *ip-address* — The IP address of the LDAP server.
- Values**
- | | |
|--------------|-------------------------------------|
| ipv4-address | a.b.c.d (host bits must be 0) |
| ipv6-address | x:x:x:x:x:x:x (eight 16-bit pieces) |
| | x:x:x:x:x:d.d.d.d |
| | x: [0..FFFF]H |
| | d: [0..255]D |
- port** *port* — Specifies the port ID. The port is the LDAP server listening port; by default it is 389 but if the listening port on LDAP server is changed, this command needs to be configured accordingly.
- Values** 1 to 65535
- Default** 389

bind-authentication

- Syntax** **bind-authentication** *root-dn* [**password** *password*] [**hash** | **hash2**]
no bind-authentication

Context config>system>security>ldap>server

Description This command configures the LDAP binding used to log into LDAP server. A string of domain components (DC) and common names (CN) can be programmed to identify the user in addition to the password field. The password is hashed. For example, "cn=admin,dc=nokia,dc=com" indicates the user admin in domain nokia.com. [Table 17](#) lists the LDAP attributes.

The **no** version of this command removes the bind-authentication.

Table 17 LDAP Attributes

Object Class	Naming Attribute Display Name	Naming Attribute LDAP Name
user	Common-Name	cn
organizationalUnit	Organizational-Unit-Name	ou
domain	Domain-Component	dc

Parameters *root-dn* — Up to 512 characters.

password *password* — Configures the password which enables a user to bind to the LDAP server. The maximum length is 128 characters.

hash — Specifies that the password is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the password is assumed to be in an unencrypted, clear text form. For security, all passwords are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified

hash2 — Specifies the password is entered in a more complex encrypted form that involves more variables than the password value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the password is assumed to be in an unencrypted, clear text form. For security, all passwords are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

ldap-server

Syntax **ldap-server** *server-name*
no ldap-server

Context config>system>security>ldap>server

Description This command configures the LDAP server name or description.

The **no** version of this command removes the LDAP server name.

Parameters *server-name* — Specifies the name of the server, up to 32 characters.

search

Syntax	search <i>base-dn</i> no search
Context	config>system>security>ldap>server
Description	<p>This command configures the LDAP search command. The search <i>base-dn</i> tells the server which part of the external directory tree to search. The search DN uses the same LDAP attribute as <i>root-dn</i>. For example, to search a public-key for an SSH generated for a Nokia vendor, one might use “dc=public-key,dc=nokia,dc=com”.</p> <p>The no version of this command remove the search DN; as such, no search will be possible on the LDAP server.</p>
Parameters	<i>base-dn</i> — Specifies the base domain name used in the search, up to 512 characters.

shutdown

Syntax	[no] shutdown
Context	config>system>security>ldap config>system>security>ldap>server
Description	<p>In the ldap context, this command enables or disabled LDAP protocol operations.</p> <p>In the server context, this command enables or disables the LDAP server. To perform no shutdown, an LDAP server address is required. To change the address, the user first needs to shut down the server.</p>

tls-profile

Syntax	tls-profile <i>tls-profile-name</i> no tls-profile
Context	config>system>security>ldap>server
Description	<p>This command attaches a TLS client profile to the LDAP client. The parameter in the TLS profile is used to encrypt the LDAP connection to the server. Each LDAP server can use its own TLS profile.</p> <p>The no version of this command removes the TLS profile from LDAP and disables the TLS encryption from LDAP.</p>
Parameters	<i>tls-profile-name</i> — Specifies the TLD profile for encryption.

timeout

Syntax	timeout <i>seconds</i> no timeout
Context	config>system>security>ldap
Description	<p>The timeout value is the number of seconds that the SR OS will wait for a response from the current server that it is trying to establish a connection with. If the server does not reply within the configured timeout value, the SR OS will increment the retry counter by 1. The SR OS attempts to establish the connection to the current server up to the configured retry value before it moves to the next configured server.</p> <p>The no version of this command configures the default timeout of 3.</p>
Parameters	<i>seconds</i> — The length of time that the SR OS waits for a response from the server.
	Values 1 to 90
	Default 3

use-default-template

Syntax	[no] use-default-template
Context	config>system>security>ldap
Description	This command specifies whether or not the default template is to be actively applied to LDAP.
Default	use-default-template

2.19.2.13 Generic 802.1x COMMANDS

dot1x

Syntax	<code>[no] dot1x</code>
Context	<code>config>system>security</code>
Description	This command creates the context to configure 802.1x network access control on the router. The no form of the command removes the 802.1x configuration.

radius-plcy

Syntax	<code>[no] radius-plcy</code>
Context	<code>config>system>security> dot1x</code>
Description	This command creates the context to configure RADIUS server parameters for 802.1x network access control on the router.



Note: The RADIUS server configured under the `config>system>security>dot1x>radius-plcy` context authenticates clients who get access to the data plane of the router as opposed to the RADIUS server configured under the **config>system>radius** context which authenticates CLI login users who get access to the management plane of the router.

The **no** form of the command removes the RADIUS server configuration for 802.1x.

retry

Syntax	<code>retry count</code> <code>no retry</code>
Context	<code>config>system>security> dot1x</code>
Description	This command configures the number of times the router attempts to contact the RADIUS server for authentication if there are problems communicating with the server. The no form of the command reverts to the default value.
Default	3
Parameters	<code>count</code> — The retry count. Values 1 to 10

server

Syntax	server <i>server-index</i> address <i>ip-address</i> secret <i>key</i> [hash hash2] [auth-port <i>auth-port</i>] [acct-port <i>acct-port</i>] [type <i>server-type</i>] no server <i>index</i>
Context	config>system>security> dot1x>radius-plcy
Description	<p>This command adds a Dot1x server and configures the Dot1x server IP address, index, and key values.</p> <p>Up to five Dot1x servers can be configured at any one time. Dot1x servers are accessed in order from lowest to highest index for authentication requests until a response from a server is received. A higher indexed server is only queried if no response is received from a lower indexed server (which implies that the server is not available). If a response from a server is received, no other Dot1x servers are queried. It is assumed that there are multiple identical servers configured as backups and that the servers do not have redundant data.</p> <p>The no form of the command removes the server from the configuration.</p>
Default	no server
Parameters	<p>server-index — The index for the Dot1x server. The index determines the sequence in which the servers are queried for authentication requests. Servers are queried in order from lowest to highest index.</p> <p>Values 1 to 5</p> <p>address <i>ip-address</i> — The IP address of the Dot1x server. Two Dot1x servers cannot have the same IP address. An error message is generated if the server address is a duplicate.</p> <p>secret <i>key</i> — The secret key to access the Dot1x server. This secret key must match the password on the Dot1x server.</p> <p>Values Up to 128 characters in length.</p> <p>hash — Specifies the key is entered in an encrypted form. If the hash or hash2 parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the hash or hash2 parameter specified</p> <p>hash2 — Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the hash2 encrypted variable cannot be copied and pasted. If the hash or hash2 parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the hash or hash2 parameter specified.</p> <p>acct-port <i>acct-port</i> — The UDP port number on which to contact the RADIUS server for accounting requests.</p> <p>auth-port <i>auth-port</i> — Specifies a UDP port number to be used as a match criteria.</p> <p>Values 1 to 65535</p>

type *server-type* — Specifies the server type.

Values authorization, accounting, combined

source-address

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

Context config>system>security> dot1x>radius-plcy

Description This command configures the NAS IP address to be sent in the RADIUS packet.
The **no** form of the command reverts to the default value.

Default By default the System IP address is used in the NAS field.

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

Values 0.0.0.0 to 255.255.255.255

shutdown

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

Context config>system>security>dot1x
config>system>security>dot1x>radius-plcy

Description This command administratively disables the 802.1x protocol operation. Shutting down the protocol does not remove or change the configuration other than the administrative state.
The operational state of the entity is disabled as well as the operational state of any entities contained within.
The **no** form of the command administratively enables the protocol which is the default state.

Default shutdown

timeout

Syntax **timeout** *seconds*
no timeout

Context config>system>security> dot1x>radius-plcy

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

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

Default 3

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

Values 1 to 90

2.19.2.14 Keychain Authentication

keychain

Syntax	<code>[no] keychain <i>keychain-name</i></code>
Context	<code>config>system>security</code>
Description	<p>This command enables the context to configure keychain parameters. A keychain must be configured on the system before it can be applied to a session.</p> <p>The no form of the command removes the keychain nodal context and everything under it from the configuration. If the keychain to be removed is in use when the no keychain command is entered, the command will not be accepted and an error indicating that the keychain is in use will be printed.</p>
Default	none
Parameters	<p><i>keychain-name</i> — Specifies a keychain name which identifies this particular keychain entry.</p> <p>Values An ASCII string up to 32 characters.</p>

direction

Syntax	<code>direction</code>
Context	<code>config>system>security>keychain</code>
Description	This command specifies the data type that indicates the TCP stream direction to apply the keychain.
Default	none

bi

Syntax	<code>bi</code>
Context	<code>config>system>security>keychain>direction</code>
Description	This command configures keys for both send and receive stream directions.
Default	none

uni

Syntax	uni
Context	config>system>security>keychain>direction
Description	This command configures keys for send or receive stream directions.
Default	none

receive

Syntax	receive
Context	config>system>security>keychain>direction>uni
Description	This command enables the receive nodal context. Entries defined under this context are used to authenticate TCP segments that are being received by the router.
Default	none

send

Syntax	send
Context	config>system>security>keychain>direction>uni
Description	This command specifies the send nodal context to sign TCP segments that are being sent by the router to another device.
Default	none

entry

Syntax	entry <i>entry-id</i> key [<i>authentication-key</i> <i>hash-key</i> <i>hash2-key</i>] [hash hash2] algorithm <i>algorithm</i> no entry <i>entry-id</i>
Context	config>system>security>keychain>direction>bi config>system>security>keychain>direction>uni>receive config>system>security>keychain>direction>uni>send
Description	This command defines a particular key in the keychain. Entries are defined by an entry-id. A keychain must have valid entries for the TCP Enhanced Authentication mechanism to work.

The **no** form of the command removes the entry from the keychain. If the entry is the active entry for sending, then this will cause a new active key to be selected (if one is available using the youngest key rule). If it is the ONLY possible send key, then the system will reject the command with an error indicating the configured key is the only available send key.

If the key is one of the eligible keys for receiving, it will be removed. If the key is the ONLY possible eligible key, then the command will not be accepted, and an error indicating that this is the only eligible key will be output.

The **no** form of the command deletes the entry.

Default There are no default entries.

Parameters *entry-id* — Specifies an entry that represents a key configuration to be applied to a keychain.

Values 0 to 63

key — Specifies a key ID which is used along with *keychain-name* and **direction** to uniquely identify this particular key entry.

authentication-key — Specifies the *authentication-key* that will be used by the encryption algorithm. The key is used to sign and authenticate a protocol packet.

The *authentication-key* can be any combination of letters or numbers.

Values A key must be 160 bits for algorithm hmac-sha-1-96 and must be 128 bits for algorithm aes-128-cmac-96. If the key given with the entry command amounts to less than this number of bits, then it is padded internally with zero bits up to the correct length.

algorithm-algorithm — Specifies an enumerated integer that indicates the encryption algorithm to be used by the key defined in the keychain.

Values aes-128-cmac-96 — Specifies an algorithm based on the AES standard for TCP authentication.
hmac-sha-1-96 — Specifies an algorithm based on SHA-1 for RSVP-TE and TCP authentication.
message-digest — MD5 hash used for TCP authentication.
hmac-md5 — MD5 hash used for IS-IS and RSVP-TE.
password — Specifies a simple password authentication for OSPF, IS-IS, and RSVP-TE.
hmac-sha-1 — Specifies the sha-1 algorithm for OSPF, IS-IS, and RSVP-TE.
hmac-sha-256 — Specifies the sha-256 algorithm for OSPF and IS-IS.

hash-key | *hash2-key* — Specifies the hash key. The key can be any combination of ASCII characters up to 33 for the *hash-key* and 96 characters for the *hash2-key* in length (encrypted). If spaces are used in the string, enclose the entire string in quotation marks (“ ”).

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

hash — Specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified

hash2 — Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

begin-time

Syntax	begin-time <i>date hours-minutes</i> [UTC] begin-time {now forever} no begin-time
Context	config>system>security>keychain>direction>bi>entry config>system>security>keychain>direction>uni>receive>entry config>system>security>keychain>direction>uni>send>entry
Description	This command specifies the calendar date and time after which the key specified by the keychain authentication key is used to sign and/or authenticate the protocol stream. If no date and time is set, the begin-time is represented by a date and time string with all NULLs and the key is not valid by default.
Parameters	<i>date hours-minutes</i> — Specifies the date and time for the key to become active. Values date: YYYY/MM/DD hours-minutes: hh:mm[:ss] now — Specifies the key should become active immediately. forever — Specifies that the key should always be active.

end-time

Syntax	end-time <i>date hours-minutes</i> [UTC] end-time {now forever} no end-time
Context	config>system>security>keychain>direction>uni>receive>entry config>system>security>keychain>direction>uni>send>entry
Description	This command specifies the calendar date and time after which the key specified by the authentication key is no longer eligible to sign and/or authenticate the protocol stream.
Default	forever

- Parameters**
- date* — Specifies the calendar date after which the key specified by the authentication key is no longer eligible to sign and/or authenticate the protocol stream in the YYYY/MM/DD format. When no year is specified the system assumes the current year.
 - hours-minutes* — Specifies the time after which the key specified by the authentication key is no longer eligible to sign and/or authenticate the protocol stream in the hh:mm[:ss] format. Seconds are optional, and if not included, assumed to be 0.
 - UTC** — Indicates that time is given with reference to Coordinated Universal Time in the input.
 - now** — Specifies a time equal to the current system time.
 - forever** — Specifies a time beyond the current epoch.

tolerance

- Syntax** **tolerance** [*seconds* | **forever**]
no tolerance
- Context** config>system>security>keychain>direction>bi>entry
config>system>security>keychain>direction>uni>receive>entry
config>system>security>keychain>direction>uni>send>entry
- Description** This command configures the amount of time that an eligible receive key should overlap with the active send key or to never expire.
- Parameters** *seconds* — Specifies the duration that an eligible receive key overlaps with the active send key.
- Values** 0 to 4294967294 seconds
- forever** — Specifies that an eligible receive key overlap with the active send key forever.

option

- Syntax** **option** {**basic** | **isis-enhanced**}
no option
- Context** config>system>security>keychain>direction>bi>entry
config>system>security>keychain>direction>uni>send>entry
- Description** This command configures allows options to be associated with the authentication key.
- Parameters** **basic** — Specifies that IS-IS should use RFC 5304 encoding of the authentication information. It is only applicable if used with the IS-IS protocol. All other protocols should ignore this configuration command.
- isis-enhanced** — Specifies that IS-IS should use RFC 5310 encoding of the authentication information. It is only applicable if used with the IS-IS protocol. All other protocols should ignore this configuration command.

tcp-option-number

Syntax	tcp-option-number
Context	config>system>security>keychain
Description	This command enables the context to configure the TCP option number to be placed in the TCP packet header.

receive

Syntax	receive <i>option-number</i> no receive
Context	config>system>security>keychain>tcp-option-number
Description	This command configures the TCP option number accepted in TCP packets received.
Default	254
Parameters	<i>option-number</i> — Specifies an enumerated integer that indicates the TCP option number to be used in the TCP header. Values 253, 254, 253&254

send

Syntax	send <i>option-number</i> no send
Context	config>system>security>keychain>tcp-option-number
Description	This command configures the TCP option number accepted in TCP packets sent.
Default	254
Parameters	<i>option-number</i> — Specifies an enumerated integer that indicates the TCP option number to be used in the TCP header. Values 253, 254

2.19.2.15 CLI Script Commands

cli-script

Syntax	cli-script
Context	config>system>security
Description	This command enables the context to configure CLI scripts.

authorization

Syntax	authorization
Context	config>system>security>cli-script
Description	This command enables the context to authorize CLI script execution.

cron

Syntax	cron
Context	config>system>security>cli-script>authorization
Description	This command enables the context to configure authorization for the Cron job-scheduler.

vsd

Syntax	[no] vsd
Context	config>system>security>cli-script>authorization
Description	This command enables the context to configure authorization for the VSD server. The no form of the command removes all authorizations for the VSD server.

event-handler

Syntax	event-handler
Context	config>system>security>cli-script>authorization

Description This command enables the context to configure authorization for the Event Handling System (EHS). EHS allows user-controlled programmatic exception handling by allowing a CLI script to be executed upon the detection of a log event.

cli-user

Syntax **cli-user** *user-name*
no cli-user

Context config>system>security>cli-script>authorization>event-handler
config>system>security>cli-script>authorization>cron
config>system>security>cli-script>authorization>vsd

Description This command configures The user context under which various types of CLI scripts should execute in order to authorize the script commands. TACACS+ and RADIUS users and authorization are not permitted for **cli-script** authorization.

The **no** form of this command configures scripts to execute with no restrictions and without performing authorization.

Default no cli-user

Parameters *user-name* — The name of a user in the local node database. TACACS+ or RADIUS users can not be used. The user configuration should reference a valid local profile for authorization.

2.19.2.16 CPM Filter Commands

cpm-filter

Syntax	cpm-filter
Context	config>system>security
Description	This command enables the context to configure a CPM filter. A CPM filter is a hardware filter done by the P chip on the CPM and CFM that applies to all the traffic going to the CPM or CFM CPU. It can be used to drop, accept packets, as well as allocate dedicated hardware queues for the traffic. The no form of the command disables the CPM filter.

default-action

Syntax	default-action {accept drop}
Context	config>system>security>cpm-filter
Description	This command specifies the action to take on the traffic when the filter entry matches. If there are no filter entry defined, the packets received will either be dropped or forwarded based on that default action.
Default	accept
Parameters	accept — Specifies that packets matching the filter entry are forwarded. drop — Specifies that packets matching the filter entry are dropped.

ip-filter

Syntax	[no] ip-filter
Context	config>system>security>cpm-filter
Description	This command enables the context to configure CPM IP filter parameters.
Default	shutdown

ipv6-filter

Syntax	[no] ipv6-filter
Context	config>system>security>cpm-filter

Description This command enables the context to configure CPM IPv6 filter parameters. This command applies only to the 7750 SR and 7950 XRS.

Default shutdown

mac-filter

Syntax [no] **mac-filter**

Context config>system>security>cpm-filter

Description This command enables the context to configure CPM MAC-filter parameters.

Default shutdown

entry

Syntax **entry** *entry-id*

Context config>sys>sec>cpm>ip-filter
config>sys>sec>cpm>ipv6-filter
config>sys>sec>cpm>mac-filter

Description This command specifies a particular CPM filter match entry. Every CPM filter must have at least one filter match entry. Entries are created and deleted by user.

The default match criteria is match none.

Parameters *entry-id* — Identifies a CPM filter entry as configured on this system.

Values 1 to 6144 for ip-filter and ipv6-filter
1 to 2048 for mac-filter

action

Syntax **action** [accept | drop | queue *queue-id*]
no action

Context config>sys>sec>cpm>ip-filter>entry
config>sys>sec>cpm>ipv6-filter>entry
config>sys>sec>cpm>mac-filter>entry

Description This command specifies the action to take for packets that match this filter entry.

Default drop

Parameters **accept** — Specifies packets matching the entry criteria will be forwarded.

drop — Specifies packets matching the entry criteria will be dropped.

queue *queue-id* — Specifies packets matching the entry criteria will be forward to the specified CPM hardware queue.

log

Syntax	log <i>log-id</i>
Context	config>sys>sec>cpm>ip-filter>entry config>sys>sec>cpm>ipv6-filter>entry config>sys>sec>cpm>mac-filter>entry
Description	This command specifies the log in which packets matching this entry should be entered. The value zero indicates that logging is disabled. The no form of the command deletes the log ID.
Parameters	<i>log-id</i> — Specifies the log ID where packets matching this entry should be entered.

match

Syntax	match [protocol <i>protocol-id</i>] no match
Context	config>sys>sec>cpm>ip-filter>entry
Description	This command enables the context to enter match criteria for the filter entry. When the match criteria have been satisfied the action associated with the match criteria is executed. If more than one match criteria (within one match statement) are configured then all criteria must be satisfied (AND function) before the action associated with the match is executed. A match context may consist of multiple match criteria, but multiple match statements cannot be entered per entry. The no form of the command removes the match criteria for the <i>entry-id</i> .
Parameters	protocol — Configures an IP protocol to be used as an IP filter match criterion. The protocol type such as TCP or UDP is identified by its respective protocol number.

protocol-id — Configures the decimal value representing the IP protocol to be used as an IP filter match criterion. Well known protocol numbers include ICMP(1), TCP(6), UDP(17). The **no** form the command removes the protocol from the match criteria.

Values 1 to 255 (values can be expressed in decimal, hexadecimal, or binary)
keywords - none, crtp, crudp, egp, eigrp, encap, ether-ip, gre, icmp, idrp, igmp, igp, ip, ipv6, ipv6-frag, ipv6-icmp, ipv6-no-nxt, ipv6-opts, ipv6-route, isis, iso-ip, l2tp, ospf-igp, pnni, ptp, rdp, rsvp, stp, tcp, udp, vrrp, * — udp/tcp wildcard

Table 18 IP Protocol Names

Protocol	Protocol ID	Description
icmp	1	Internet Control Message
igmp	2	Internet Group Management
ip	4	IP in IP (encapsulation)
tcp	6	Transmission Control
egp	8	Exterior Gateway Protocol
igp	9	any private interior gateway (used by Cisco for their IGRP)
udp	17	User Datagram
rdp	27	Reliable Data Protocol
ipv6	41	IPv6
ipv6-route	43	Routing Header for IPv6
ipv6-frag	44	Fragment Header for IPv6
idrp	45	Inter-Domain Routing Protocol
rsvp	46	Reservation Protocol
gre	47	General Routing Encapsulation
ipv6-icmp	58	ICMP for IPv6
ipv6-no-nxt	59	No Next Header for IPv6
ipv6-opts	60	Destination Options for IPv6
iso-ip	80	ISO Internet Protocol
eigrp	88	EIGRP
ospf-igp	89	OSPFIGP

Table 18 IP Protocol Names (Continued)

Protocol	Protocol ID	Description
ether-ip	97	Ethernet-within-IP Encapsulation
encap	98	Encapsulation Header
pnni	102	PNNI over IP
pim	103	Protocol Independent Multicast
vrrp	112	Virtual Router Redundancy Protocol
l2tp	115	Layer Two Tunneling Protocol
stp	118	Spanning Tree Protocol
ptp	123	Performance Transparency Protocol
isis	124	ISIS over IPv4
crtp	126	Combat Radio Transport Protocol
crudp	127	Combat Radio User Datagram

match

Syntax `match [next-header next-header]`

`no match`

Context `config>sys>sec>cpm>ipv6-filter>entry`

Description This command specifies match criteria for the IP filter entry. This command applies only the the 775 SR and 7950 XRS.

The **no** form of this command removes the match criteria for the *entry-id*.

Parameters *next-header next-header* — Specifies the next header to match.

The protocol type such as TCP / UDP / OSPF is identified by its respective protocol number. Well-known protocol numbers include ICMP(1), TCP(6), UDP(17).

Values

next-header: 1 to 42, 45 to 49, 52 to 59, 61 to 255 protocol numbers accepted in DHB

keywords: none, crtp, crudp, egp, eigrp, encap, ether-ip, gre, icmp, drp, igmp, igp, ip, ipv6, ipv6-icmp, ipv6-no-nxt, isis, iso-ip, l2tp, spf-igp, pim, pnni, ptp, rdp, rsvp, stp, tcp, udp, vrrp

* — udp/tcp wildcard

action

Syntax	action {permit deny} no action
Context	config>system>security>mgmt-access-filter>mac-filter
Description	<p>This command creates the action associated with the management access filter match criteria entry.</p> <p>The action keyword is required. If no action is defined, the filter is ignored. If multiple action statements are configured, the last one overwrites previous configured actions.</p> <p>If the packet does not meet any of the match criteria the configured default action is applied.</p>
Default	none — The action is specified by default-action command.
Parameters	<p><i>permit</i> — Specifies that packets matching the configured criteria will be permitted.</p> <p>deny — Specifies that packets matching the configured selection criteria will be denied and that a ICMP host unreachable message will not be issued.</p>

default-action

Syntax	default-action {permit deny}
Context	config>system>security>mgmt-access-filter>mac-filter
Description	<p>This command creates the default action for management access in the absence of a specific management access filter match.</p> <p>The default-action is applied to a packet that does not satisfy any match criteria in any of the management access filters. Whenever management access filters are configured, the default-action must be defined.</p>
Default	No default-action is defined.
Parameters	<p>permit — Specifies that packets not matching the configured selection criteria in any of the filter entries will be permitted.</p> <p>deny — Specifies that packets not matching the selection criteria be denied and that an ICMP host unreachable message will not be issued.</p>

dscp

Syntax	dscp <i>dscp-name</i> no dscp
Context	config>sys>sec>cpm>ip-filter>entry>match

```
config>sys>sec>cpm>ipv6-filter>entry>match
config>sys>sec>cpm>mac-filter>entry>match
```

Description This command configures a DiffServ Code Point (DSCP) name to be used as an IP filter match criterion.

The **no** form of the command removes the DSCP match criterion.

Default no dscp

Parameters *dscp-name* — Configures a dscp name that has been previously mapped to a value using the **dscp-name** command. The DiffServ code point may only be specified by its name.

dst-ip

Syntax **dst-ip** *ipv6-address/prefix-length*
dst-ip **ipv6-prefix-list** *ipv6-prefix-list-name*
no dst-ip

Context config>sys>sec>cpm>ip-filter>entry>match
config>sys>sec>cpm>ipv6-filter>entry>match

Description This command configures a destination IP address range to be used as an IP filter match criterion.

To match on the destination IP address, specify the address and its associated mask, for example, 10.1.0.0/16. The conventional notation of 10.1.0.0 255.255.0.0 may also be used.

The **no** form of the command removes the destination IP address match criterion.

Default no dst-ip

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

Values 0.0.0.0 to 255.255.255.255

ip-prefix-list — Creates a list of IPv4 prefixes for match criteria in IPv4 ACL and CPM filter policies.

ip-prefix-list-name — A string of up to 32 characters of printable ASCII characters. If special characters are used, the string must be enclosed within double quotes.

mask — Specifies the subnet mask length expressed as a decimal integer.

Values 1 to 32

netmask — Specifies the dotted quad equivalent of the mask length.

Values 0.0.0.0 to 255.255.255.255

dst-ip

Syntax	dst-ip [<i>ipv6-address /prefix-length</i>] [<i>ipv6-prefix-list ipv6-prefix-list-name</i>] no dst-ip										
Context	config>sys>sec>cpm>ipv6-filter>entry>match										
Description	<p>This command configures a destination IPv6 address range to be used as an IPv6 filter match criterion.</p> <p>To match on the destination IPv6 address, specify the address.</p> <p>The no form of the command removes the destination IP address match criterion.</p> <p>This command only applies to the 7750 SR and 7950 XRS.</p>										
Default	no dst-ip										
Parameters	<p><i>ipv6-address/prefix-length</i> — Specifies the IPv6 address for the IPv6 match criterion in dotted decimal notation. An IPv6 IP address is written as eight 4-digit (16-bit) hexadecimal numbers separated by colons. One string of zeros per address can be left out, so that 1010::700:0:217A is the same as 1010:0:0:0:700:0:217A.</p> <p>Values</p> <table border="0" style="margin-left: 20px;"> <tr> <td>x:x:x:x:x:x:x</td> <td>(eight 16-bit pieces)</td> </tr> <tr> <td>x:x:x:x:x:d.d.d.d</td> <td></td> </tr> <tr> <td style="padding-left: 100px;">x:</td> <td>[0 to .FFFF]H</td> </tr> <tr> <td style="padding-left: 100px;">d:</td> <td>[0 to 255]D</td> </tr> <tr> <td style="padding-left: 100px;">prefix-length:</td> <td>1 to 128</td> </tr> </table> <p>ipv6-prefix-list — Creates a list of IPv4 prefixes for match criteria in IPv4 ACL and CPM filter policies.</p> <p><i>ipv6-prefix-list-name</i> — Specifies a string of up to 32 characters of printable ASCII characters. If special characters are used, the string must be enclosed within double quotes.</p>	x:x:x:x:x:x:x	(eight 16-bit pieces)	x:x:x:x:x:d.d.d.d		x:	[0 to .FFFF]H	d:	[0 to 255]D	prefix-length:	1 to 128
x:x:x:x:x:x:x	(eight 16-bit pieces)										
x:x:x:x:x:d.d.d.d											
x:	[0 to .FFFF]H										
d:	[0 to 255]D										
prefix-length:	1 to 128										

dst-port

Syntax	dst-port [<i>tcp/udp port-number</i>] [<i>mask</i>] dst-port port-list <i>port-list-name</i> dst-port range <i>tcp/udp port-number tcp/udp port-number</i> no dst-port
Context	config>sys>sec>cpm>ip-filter>entry>match config>sys>sec>cpm>ipv6-filter>entry>match

Description This command specifies the TCP/UDP port or port name to match the destination-port of the packet.



Note: An entry containing Layer 4 match criteria will not match non-initial (2nd, 3rd, etc) fragments of a fragmented packet since only the first fragment contains the Layer 4 information.

The **no** form of the command removes the destination port match criterion.

Default no dst-port

Parameters *tcp/udp port-numb-number* — Specifies the destination port number to be used as a match criteria expressed as a decimal integer.

Values 0 to 65535 (accepted in decimal hex or binary)

port-list-name — Specifies the port list name to be used as a match criteria for the destination port.

mask — Specifies the 16 bit mask to be applied when matching the destination port.

Values [0x0000..0xFFFF] | [0..65535] |
[0b0000000000000000..0b1111111111111111]

flow-label

Syntax **flow-label** *value*
no flow-label

Context config>sys>sec>cpm>ipv6-filter>entry>match

Description This command configures flow label match conditions. Flow labeling enables the labeling of packets belonging to particular traffic flows for which the sender requests special handling, such as non-default quality of service or real-time service.

Parameters *value* — Specify the flow identifier in an IPv6 packet header that can be used to discriminate traffic flows (See RFC 3595, *Textual Conventions for IPv6 Flow Label*.)

Values 0 to 1048575

fragment

Syntax **fragment** {true | false}
no fragment

Context config>sys>sec>cpm>ip-filter>entry>match
config>sys>sec>cpm>ipv6-filter>entry>match

Description This command specifies fragmented or non-fragmented IP packets as an IP filter match criterion.



Note: An entry containing Layer 4 match criteria will not match non-initial (2nd, 3rd, etc) fragments of a fragmented packet since only the first fragment contains the Layer 4 information.

This command enables match on existence of IPv6 Fragmentation Extension Header in the IPv6 filter policy. To match first fragment of an IP fragmented packet, specify additional Layer 4 matching criteria in a filter policy entry. The **no** version of this command ignores IPv6 Fragmentation Extension Header presence/absence in a packet when evaluating match criteria of a given filter policy entry.

The **no** form of the command removes the match criterion.

This command enables match on existence of IPv6 Fragmentation Extension Header in the IPv6 filter policy. To match first fragment of an IP fragmented packet, specify additional Layer 4 matching criteria in a filter policy entry. The **no** version of this command ignores IPv6 Fragmentation Extension Header presence/absence in a packet when evaluating match criteria of a given filter policy entry.

Default no fragment

Parameters **true** — Specifies to match on all fragmented IP packets. A match will occur for all packets that have either the MF (more fragment) bit set or have the Fragment Offset field of the IP header set to a non-zero value. For IPv6, packet matches if it contains IPv6 Fragmentation Extension Header.

false — Specifies to match on all non-fragmented IP packets. Non-fragmented IP packets are packets that have the MF bit set to zero and have the Fragment Offset field also set to zero. For IPv6, packet matches if it does not contain IPv6 Fragmentation Extension Header.

hop-by-hop-opt

Syntax **hop-by-hop-opt {true | false}**
no hop-by-hop-opt

Context config>sys>sec>cpm>ipv6-filter>entry>match

Description This command enables match on existence of Hop-by-Hop Options Extension Header in the IPv6 filter policy. This command applies to the 7750 SR and 7950 XRS.

The **no** form of this command ignores Hop-by-Hop Options Extension Header presence/absence in a packet when evaluating match criteria of a given filter policy entry.

Default no hop-by-hop-opt

- Parameters** **true** — Match if a packet contains Hop-by-Hop Options Extension Header.
false — Match if a packet does not contain Hop-by-Hop Options Extension Header.

icmp-code

- Syntax** **icmp-code** *icmp-code*
no icmp-code
- Context** config>sys>sec>cpm>ip-filter>entry>match
config>sys>sec>cpm>ipv6-filter>entry>match
- Description** This command configures matching on ICMP code field in the ICMP header of an IP packet as an IP filter match criterion.



Note: An entry containing Layer 4 match criteria will not match non-initial (2nd, 3rd, etc) fragments of a fragmented packet since only the first fragment contains the Layer 4 information.

The behavior of the **icmp-code** value is dependent on the configured **icmp-type** value, thus a configuration with only an **icmp-code** value specified will have no effect. To match on the **icmp-code**, an associated **icmp-type** must also be specified.

The **no** form of the command removes the criterion from the match entry.

- Default** no icmp-code
- Parameters** *icmp-code* — Specifies the ICMP code values that must be present to match.
Values 0 to 255

icmp-type

- Syntax** **icmp-type** *icmp-type*
no icmp-type
- Context** config>sys>sec>cpm>ip-filter>entry>match
config>sys>sec>cpm>ipv6-filter>entry>match
- Description** This command configures matching on ICMP type field in the ICMP header of an IP packet as an IP filter match criterion.



Note: An entry containing Layer 4 match criteria will not match non-initial (2nd, 3rd, etc) fragments of a fragmented packet since only the first fragment contains the Layer 4 information.

The **no** form of the command removes the criterion from the match entry.

Default no icmp-type

Parameters *icmp-type* — Specifies the ICMP type values that must be present to match.

Values 0 to 255

ip-option

Syntax **ip-option** *ip-option-value ip-option-mask*
no ip-option

Context config>sys>sec>cpm>ip-filter>entry>match

Description This command configures matching packets with a specific IP option or a range of IP options in the IP header as an IP filter match criterion.

The option-type octet contains 3 fields:

- 1 bit copied flag (copy options in all fragments)
- 2 bits option class,
- 5 bits option number.

The **no** form of the command removes the match criterion.

Default no ip-option

Parameters *ip-option-value* — Enter the 8 bit option-type as a decimal integer. The mask is applied as an AND to the option byte, the result is compared with the option-value.
The decimal value entered for the match should be a combined value of the eight bit option type field and not just the option number. Thus to match on IP packets that contain the Router Alert option (option number =20), enter the option type of 148 (10010100).

Values 0 to 255

ip-option-mask — Specifies a range of option numbers to use as the match criteria.
This 8 bit mask can be configured using the formats described in [Table 19](#):

Table 19 ip-option-mask Formats

Format Style	Format Syntax	Example
Decimal	DDD	20
Hexadecimal	0xHH	0x14
Binary	0BBBBBBBB	0b0010100

Default 255 (decimal) (exact match)

Values 1 to 255 (decimal)

multiple-option

Syntax **multiple-option {true | false}**
no multiple-option

Context config>sys>sec>cpm>ip-filter>entry>match

Description This command configures matching packets that contain more than one option fields in the IP header as an IP filter match criterion.

The **no** form of the command removes the checking of the number of option fields in the IP header as a match criterion.

Default no multiple-option

Parameters **true** — Specifies matching on IP packets that contain more that one option field in the header.

false — Specifies matching on IP packets that do not contain multiple option fields present in the header.

option-present

Syntax **option-present {true | false}**
no option-present

Context config>sys>sec>cpm>ip-filter>entry>match

Description This command configures matching packets that contain the option field or have an option field of zero in the IP header as an IP filter match criterion.

The **no** form of the command removes the checking of the option field in the IP header as a match criterion.

- Default** no option-present
- Parameters** **true** — Specifies matching on all IP packets that contain the option field in the header. A match will occur for all packets that have the option field present. An option field of zero is considered as no option present.
- false** — Specifies matching on IP packets that do not have any option field present in the IP header (an option field of zero). An option field of zero is considered as no option present.

router

- Syntax** **router service-name service-name**
router router-instance
no router
- Context** config>sys>sec>cpm>ip-filter>entry>match
config>sys>sec>cpm>ipv6-filter>entry>match
- Description** This command specifies a router name or a service-id to be used in the match criteria.
- Default** no router
- Parameters** *router-instance* — Specify one of the following parameters for the router instance:
router-name — Specifies a router name up to 32 characters to be used in the match criteria.
service-id — Specifies an existing service ID to be used in the match criteria.
- Values** 1 to 2147483647
- service-name service-name* — Specifies an existing service name up to 64 characters in length.

src-ip

- Syntax** **src-ip [ip-address/mask | ip-prefix-list prefix-list-name]**
no src-ip
- Context** config>sys>sec>cpm>ip-filter>entry>match
- Description** This command specifies the IP address to match the source IP address of the packet.
- To match on the source IP address, specify the address and its associated mask, such as 10.1.0.0/16. The conventional notation of 10.1.0.0 255.255.0.0 may also be used.
- The **no** form of the command removes the source IP address match criterion.
- Default** no src-ip

Parameters *ip-address/mask* — Specifies the IP address for the match criterion in dotted decimal notation. An IP address is written as eight 4-digit (16-bit) hexadecimal numbers separated by colons. One string of zeros per address can be left out, so that 1010::700:0:217A is the same as 1010:0:0:0:700:0:217A.

Values

ipv4-address	a.b.c.d (host bits must be 0) x:x:x:x:x:d.d.d.d[-interface] x: [0..FFFF]H d: [0..255]D interface: 32 characters maximum, mandatory for link local addresses
mask:	Specifies the 16 bit mask to be applied when matching the source IP address. 1 to 32

ip-prefix-list — Creates a list of IPv4 prefixes for match criteria in IPv4 ACL and CPM filter policies.

ip-prefix-list-name — A string of up to 32 characters of printable ASCII characters. If special characters are used, the string must be enclosed within double quotes.

src-ip

Syntax **src-ip** [*ip-address/mask* | **ipv6-prefix-list** *ipv6-prefix-list-name*]
no src-ip

Context config>sys>sec>cpm>ipv6-filter>entry>match

Description This command specifies the IPv6 address to match the source IPv6 address of the packet.

To match on the source IP address, specify the address and its associated mask, such as 10.1.0.0/16. The conventional notation of 10.1.0.0 255.255.0.0 may also be used.

The **no** form of the command removes the source IP address match criterion.

This command only applies to the 7750 SR and 7950 XRS.

Default no src-ip

Parameters *ip-address/mask* — Specifies the IP address for the match criterion in dotted decimal notation. An IP address is written as eight 4-digit (16-bit) hexadecimal numbers separated by colons. One string of zeros per address can be left out, so that 1010::700:0:217A is the same as 1010:0:0:0:700:0:217A.

Values

ipv6-address	x:x:x:x:x:x[-interface] x:x:x:x:x:d.d.d.d[-interface]
--------------	--

x: [0..FFFF]H
 d: [0..255]D
 interface: 32 characters maximum, mandatory for link local addresses
 mask: Specifies eight 16-bit hexadecimal pieces representing bit match criteria.
 Values x:x:x:x:x:x (eight 16-bit pieces)

ipv6-prefix-list — Creates a list of IPv6 prefixes for match criteria in IPv6 ACL and CPM filter policies.

ipv6-prefix-list-name — A string of up to 32 characters of printable ASCII characters. If special characters are used, the string must be enclosed within double quotes.

src-port

Syntax **src-port** *src-port-number* [*mask*]

Context config>sys>sec>cpm>ip-filter>entry>match
config>sys>sec>cpm>ipv6-filter>entry>match

Description This command specifies the TCP/UDP port to match the source port of the packet.



Note: An entry containing Layer 4 match criteria will not match non-initial (2nd, 3rd, etc) fragments of a fragmented packet since only the first fragment contains the Layer 4 information.

Default no src-port

Parameters *src-port-number* — The source port number to be used as a match criteria expressed as a decimal integer.

Values 0 to 65535

mask — Specifies the 16 bit mask to be applied when matching the source port.

Values 0 to 128

tcp-ack

Syntax **tcp-ack** {**true** | **false**}
no tcp-ack

Context config>sys>sec>cpm>ip-filter>entry>match

```
config>sys>sec>cpm>ipv6-filter>entry>match
```

Description This command configures matching on the ACK bit being set or reset in the control bits of the TCP header of an IP or IPv6 packet as an IP filter match criterion.



Note: An entry containing Layer 4 match criteria will not match non-initial (2nd, 3rd, etc) fragments of a fragmented packet since only the first fragment contains the Layer 4 information.

The **no** form of the command removes the criterion from the match entry.

Default no tcp-ack

Parameters **true** — Specifies matching on IP or IPv6 packets that have the ACK bit set in the control bits of the TCP header of an IP or IPv6 packet.

false — Specifies matching on IP or IPv6 packets that do not have the ACK bit set in the control bits of the TCP header of the IP or IPv6 packet.

tcp-syn

Syntax **tcp-syn {true | false}**
no tcp-syn

Context config>sys>sec>cpm>ip-filter>entry>match
config>sys>sec>cpm>ipv6-filter>entry>match
config>sys>sec>cpm>ipv6-filter>entry>match

Description This command configures matching on the SYN bit being set or reset in the control bits of the TCP header of an IP or IPv6 packet as an IP filter match criterion.



Note: An entry containing Layer 4 match criteria will not match non-initial (2nd, 3rd, etc) fragments of a fragmented packet since only the first fragment contains the Layer 4 information.

The SYN bit is normally set when the source of the packet wants to initiate a TCP session with the specified destination IP or IPv6 address.

The **no** form of the command removes the criterion from the match entry.

Default no tcp-syn

Parameters **true** — Specifies matching on IP or IPv6 packets that have the SYN bit set in the control bits of the TCP header.

false — Specifies matching on IP or IPv6 packets that do not have the SYN bit set in the control bits of the TCP header.

renum

Syntax	renum <i>old-entry-id</i> <i>new-entry-id</i>
Context	config>sys>sec>cpm>ip-filter config>sys>sec>cpm>ipv6-filter>entry>match config>sys>sec>cpm>mac-filter>entry>match
Description	<p>This command renumbers existing IP(IPv4), IPv6, or MAC filter entries to re-sequence filter entries.</p> <p>This may be required in some cases since the OS exits when the first match is found and execute the actions according to the accompanying action command. This requires that entries be sequenced correctly from most to least explicit.</p>
Parameters	<p><i>old-entry-id</i> — Enter the entry number of an existing entry.</p> <p>Values 1 to 6144 for ip-filter and ipv6-filter 1 to 2048 for mac-filter</p> <p><i>new-entry-id</i> — Enter the new entry-number to be assigned to the old entry.</p> <p>Values 1 to 6144 for ip-filter and ipv6-filter 1 to 2048 for mac-filter</p>

shutdown

Syntax	[no] shutdown
Context	config>sys>sec>cpm>ip-filter config>sys>sec>cpm>ipv6-filter config>sys>sec>cpm>mac-filter
Description	<p>This command enables IPv4, IPv6 or MAC CPM filter.</p> <p>The no form of this command disable the filter.</p>
Default	shutdown

2.19.2.17 CPM Queue Commands

cpm-queue

- Syntax** `cpm-queue`
- Context** `config>system>security`
- Description** This command enables the context to configure a CPM queue.

queue

- Syntax** `queue queue-id`
- Context** `config>system>security>cpm-queue`
- Description** This command allows users to allocate dedicated CPM. The first available queue is 33.

cbs

- Syntax** `cbs cbs`
`no cbs`
- Context** `config>system>cpm-queue>queue`
- Description** This command specifies the amount of buffer that can be drawn from the reserved buffer portion of the queue's buffer pool.
- Parameters** *cbs* — Specifies the committed burst size in kbytes.

mbs

- Syntax** `mbs mbs`
`no mbs`
- Context** `config>system>security>cpm-queue>queue`
- Description** This command specifies the maximum queue depth to which a queue can grow.
- Parameters** *mbs* — Specifies the maximum burst size in kbytes.

rate

Syntax	rate <i>rate</i> [cir <i>cir</i>] no rate
Context	config>system>security>cpm-queue>queue
Description	This command specifies the maximum bandwidth that will be made available to the queue in kilobits per second (kb/s).
Parameters	<i>rate</i> — Specifies the administrative Peak Information Rate (PIR) for the queue. <i>cir cir</i> — Specifies the amount of bandwidth committed to the queue.

2.19.2.18 TTL Security Commands

ttl-security

Syntax	ttl-security <i>min-ttl-value</i> no ttl-security
Context	config>router>bgp>group config>router>bgp>group>neighbor config>router>ldp>tcp-session-params>peer-transport config>system>login-control>ssh config>system>login-control>telnet
Description	This command configures TTL security parameters for incoming packets. When the feature is enabled, LDP will accept incoming IP packets from a peer only if the TTL value in the packet is greater than or equal to the minimum TTL value configured for that peer. Per-peer-queueing must be enabled in order for TTL protection to operate. The no form of the command disables TTL security.
Parameters	<i>min-ttl-value</i> — Specify the minimum TTL value for an incoming BGP packet. Values 1 to 255

ttl-security

Syntax	ttl-security <i>min-ttl-value</i> no ttl-security
Context	config>router>ldp>tcp-session-params>peer-transport

-
- Description** This command configures TTL security parameters for incoming packets. When the feature is enabled, BGP will accept incoming IP packets from a peer only if the TTL value in the packet is greater than or equal to the minimum TTL value configured for that peer. Per-peer-queueing must be enabled in order for TTL protection to operate.
- The **no** form of the command disables TTL security.
- Default** no ttl-security
- Parameters** *min-ttl-value* — Specifies the minimum TTL value for an incoming LDP packet.
- Values** 1 to 255

ttl-security

- Syntax** **ttl-security** *min-ttl-value*
no ttl-security
- Context** config>system>login-control>ssh
config>system>login-control>telnet
- Description** This command configures TTL security parameters for incoming packets. When the feature is enabled, SSH/Telnet will accept incoming IP packets from a peer only if the TTL value in the packet is greater than or equal to the minimum TTL value configured for that peer. Per-peer-queueing must be enabled in order for TTL protection to operate.
- The **no** form of the command disables TTL security.
- Parameters** *min-ttl-value* — Specify the minimum TTL value for an incoming BGP packet.
- Values** 1 to 255

2.19.2.19 gRPC Commands

grpc

Syntax	grpc
Context	config>system
Description	This command enters the context to configure gRPC parameters.

tls-server-profile

Syntax	tls-server-profile <i>name</i> no tls-server-profile
Context	config>system>grpc
Description	This command adds a configured TLS server profile to the gRPC session. The TLS server is used for encryption of the gRPC session. gRPC will not transmit any PDUs if there is a TLS server profile assigned to it and the TLS connection is down. The no form of the command removes the specified TLS server profile from the gRPC session.
Parameters	<i>name</i> — Specifies the name of the TLS server profile configured under the config>system>security>tls context.

2.19.2.20 CPU Protection Commands

cpu-protection

- Syntax** `cpu-protection`
- Context** `config>sys>security`
- Description** This command enters the context to configure CPU protection parameters.

included-protocols

- Syntax** `included-protocols`
- Context** `config>sys>security>cpu-protection> ip>included-protocols`
- Description** This context allows configuration of which protocols are included for ip-src-monitoring. This is system-wide configuration that applies to cpu protection globally.

dhcp

- Syntax** `[no] dhcp`
- Context** `config>sys>security>cpu-protection> ip>included-protocols`
- Description** This command includes the extracted IPv4 DHCP packets for ip-src-monitoring. IPv4 DHCP packets will be subject to the per-source-rate of CPU protection policies.
- Default** `dhcp` (Note this is different from the other protocols)

gtp

- Syntax** `[no] gtp`
- Context** `config>sys>security>cpu-protection> ip>included-protocols`
- Description** This command includes the extracted IPV4 GTP packets for ip-src-monitoring. IPv4 GTP packets will be subject to the per-source-rate of CPU protection policies.
- Default** `no gtp`

icmp

Syntax	[no] icmp
Context	config>sys>security>cpu-protection> ip>included-protocols
Description	This command includes the extracted IPv4 ICMP packets for ip-src-monitoring. IPv4 ICMP packets will be subject to the per-source-rate of CPU protection policies.
Default	no icmp

igmp

Syntax	[no] igmp
Context	config>sys>security>cpu-protection> ip>included-protocols
Description	This command includes the extracted IPv4 IGMP packets for ip-src-monitoring. IPv4 IGMP packets will be subject to the per-source-rate of CPU protection policies.
Default	no igmp

link-specific-rate

Syntax	link-specific-rate <i>packet-rate-limit</i> no link-specific-rate
Context	config>sys>security>cpu-protection
Description	This command configures a link-specific rate for CPU protection. This limit is applied to all ports within the system. The CPU will receive no more than the configured packet rate for all link level protocols such as LACP from any one port. The measurement is cleared each second and is based on the ingress port.
Default	15000
Parameters	<i>packet-rate-limit</i> — Specifies a packet arrival rate limit, in packets per second, for link level protocols.
Values	1 to 65535, max (no limit)

policy

Syntax	policy <i>cpu-protection-policy-id</i> [create] no policy <i>cpu-protection-policy-id</i>
Context	config>sys>security>cpu-protection

Description	<p>This command configures CPU protection policies.</p> <p>The no form of the command deletes the specified policy from the configuration.</p> <p>Policies 254 and 255 are reserved as the default access and network interface policies, and cannot be deleted. The parameters within these policies can be modified. An event will be logged (warning) when the default policies are modified.</p>
Default	<p>Policy 254 (default access interface policy):</p> <ul style="list-style-type: none"> • per-source-rate: max (no limit) • overall-rate: 6000 • out-profile-rate: 6000 • alarm <p>Policy 255 (default network interface policy):</p> <ul style="list-style-type: none"> • per-source-rate: max (no limit) • overall-rate : max (no limit) • out-profile-rate: 3000 • alarm
Parameters	<p><i>cpu-protection-policy-id</i> — Assigns a policy ID to the specific CPU protection policy.</p> <p>Values 1 to 255</p> <p>create — Keyword used to create CPU protection policy. The create keyword requirement can be enabled/disabled in the environment>create context.</p>

alarm

Syntax	[no] alarm
Context	config>sys>security>cpu-protection>policy
Description	<p>This command enables the generation of an event when a rate is exceeded. The event includes information about the offending source. Only one event is generated per monitor period.</p> <p>The no form of the command disables the notifications.</p>
Default	no alarm

eth-cfm

Syntax	[no] eth-cfm
Context	config>sys>security>cpu-protection>policy

Description Provides the construct under which the different entries within CPU policy can define the match criteria and overall arrival rate of the Ethernet Configuration and Fault Management (ETH-CFM) packets at the CPU.

entry

Syntax **entry** *entry levels levels opcodes opcodes rate packet-rate-limit*
no entry

Context config>sys>security>cpu-protection>eth-cfm>

Description Builds the specific match and rate criteria. Up to ten entries may exist in up to four CPU protection policies.

The **no** form of the command reverses the match and rate criteria configured.

Default no entry

Parameters *rate* — Specifies a packet rate limit in frames per second, where a '0' means drop all.

Values 1 to 100

level — Specifies a domain level.

Values

all	Wildcard entry level
range	0 to 7: within specified range, multiple ranges allowed
number	0 to 7: specific level number, may be combined with range

opcode — Specifies an operational code that identifies the application.

Values

range	0 to 255: within specified range, multiple ranges allowed
number	0 to 255: specific level number, may be combined with range

out-profile-rate

Syntax **out-profile-rate** *packet-rate-limit [log-event]*
no out-profile-rate

Context config>sys>security>cpu-protection>policy

Description	This command applies a packet arrival rate limit for the entire SAP/interface, above which packets will be marked as discard eligible, in other words, out-profile/low-priority/yellow. The rate defined is a global rate limit for the interface regardless of the number of traffic flows. It is a per-SAP/interface rate. The no form of the command sets out-profile-rate parameter back to the default value.
Default	3000 for cpu-protection-policy-id 1-253 6000 for cpu-protection-policy-id 254 (default access interface policy) 3000 for cpu-protection-policy-id 255 (default network interface policy)
Parameters	<i>packet-rate-limit</i> — Specifies a packet arrival rate limit in packets per second. Values 1 to 65535, max (max indicates no limit) log-events — Issues a tmnxCpmProtViolSapOutProf, tmnxCpmProtViolIfOutProf, or tmnxCpmProtViolSdpBindOutProf log event and tracks violating interfaces when the out-profile-rate is exceeded. Supported on CPM3 and above only.

overall-rate

Syntax	overall-rate <i>packet-rate-limit</i> no overall-rate
Context	config>sys>security>cpu-protection>policy
Description	This command applies a maximum packet arrival rate limit (applied per SAP/interface) for the entire SAP/interface, above which packets will be discarded immediately. The rate defined is a global rate limit for the interface regardless of how many traffic flows are present on the SAP/interface. It is a per-SAP/interface rate. The no form of the command sets overall-rate parameter back to the default value.
Default	max for cpu-protection-policy-id 1 to 253 6000 for cpu-protection-policy-id 254 (default access interface policy) max for cpu-protection-policy-id 255 (default network interface policy)
Parameters	<i>packet-rate-limit</i> — Specifies a packet arrival rate limit in packets per second. Values 1 to 65535, max (the max indicates no limit)

per-source-rate

Syntax	per-source-rate <i>packet-rate-limit</i> no per-source-rate
---------------	--

Context	config>sys>security>cpu-protection>policy
Description	<p>This command configures a per-source packet arrival rate limit. Use this command to apply a packet arrival rate limit on a per source basis. A source is defined as a unique combination of SAP and MAC source address (mac-monitoring) or SAP and source IP address (ip-src-monitoring). The CPU will receive no more than the configured packet rate from each source (only certain protocols are rate limited for ip-src-monitoring as configured under 'include-protocols' in the cpu protection policy). The measurement is cleared each second.</p> <p>This parameter is only applicable if the policy is assigned to an interface (some examples include saps, subscriber-interfaces, and spoke-sdps), and the mac-monitor or ip-src-monitor keyword is specified in the cpu-protection configuration of that interface.</p> <p>The ip-src-monitoring is useful in subscriber management architectures that have routers between the subscriber and the BNG (router). In layer-3 aggregation scenarios, all packets from all subscribers behind the same aggregation router will arrive with the same source MAC address and as such the mac-monitoring functionality can not differentiate traffic from different subscribers.</p>
Default	max, no limit
Parameters	<p><i>packet-rate-limit</i> — Specifies a per-source packet (per SAP/MAC source address or per SAP/IP source address) arrival rate limit in packets per second.</p> <p>Values 1 to 65535, max (max indicates no limit)</p>

port-overall-rate

Syntax	port-overall-rate <i>packet-rate-limit</i> [low-action-priority] no port-overall-rate
Context	config>sys>security>cpu-protection
Description	This command configures a per-port overall rate limit for CPU protection.
Parameters	<p><i>packet-rate-limit</i> — Specifies an overall per-port packet arrival rate limit in packets per second.</p> <p>Values 1 to 65535, max (indicates no limit)</p> <p>action-low-priority — Marks packets that exceed the rate as low-priority (for preferential discard later if there is congestion in the control plane) instead of discarding them immediately.</p> <p>Default max</p>

protocol-protection

Syntax **protocol-protection** [**allow-sham-links**] [**block-pim-tunneled**]

no protocol-protection

Context	config>sys>security>cpu-protection
Description	This command causes the network processor on the CPM to discard all packets received for protocols that are not configured on the particular interface. This helps mitigate DoS attacks by filtering invalid control traffic before it hits the CPU. For example, if an interface does not have IS-IS configured, then protocol protection will discard any IS-IS packets received on that interface.
Default	no protocol-protection
Parameters	<p>allow-sham-links — Allows sham links. As OSPF sham links form an adjacency over the MPLS-VP RN backbone network, when protocol-protection is enabled, the tunneled OSPF packets to be received over the backbone network must be explicitly allowed.</p> <p>block-pim-tunneled — - Blocks extraction and processing of PIM packets arriving at the SR-OS node inside a tunnel (for example, MPLS or GRE) on a network interface. With protocol-protection enabled and tunneled pim blocked, PIM in an mVPN on the egress DR will not switch traffic from the (*,G) to the (S,G) tree.</p>

cpu-protection

Syntax	cpu-protection <i>policy-id</i> no cpu-protection
Context	config>router>interface config>service>ies>interface config>service>ies>video-interface config>service>vpls>video-interface config>service>vprn>interface config>service>vprn>network-interface config>service>vprn>video-interface
Description	<p>Use this command to apply a specific CPU protection policy to the associated interface. For these interface types, the per-source rate limit is not applicable.</p> <p>If no CPU-protection policy is assigned to an interface, then the default policy is used to limit the overall-rate. The default policy is policy number 254 for access interfaces, 255 for network interfaces and no policy for video interfaces.</p> <p>The no form of the command reverts to the default values.</p>
Default	<p>cpu-protection 254 (for access interfaces)</p> <p>cpu-protection 255 (for network interfaces)</p> <p>no cpu-protection (for video interfaces)</p>

cpu-protection

Syntax	cpu-protection <i>policy-id</i> [mac-monitoring] [ip-src-monitoring] no cpu-protection
Context	config>subscriber-mgmt>msap-policy
Description	<p>Use this command to apply a specific CPU protection policy to the associated msap-policy. The specified cpu-protection policy will automatically be applied to any MSAPs that are create using the msap-policy.</p> <p>If no CPU-protection policy is assigned to a SAP, then a default policy is used to limit the overall-rate according to the default policy. The default policy is policy number 254 for access interfaces, 255 for network interfaces and no policy for video interfaces.</p> <p>The no form of the command reverts to the default values.</p>
Default	<p>cpu-protection 254 (for access interfaces)</p> <p>cpu-protection 255 (for network interfaces)</p> <p>The configuration of no cpu-protection returns the msap-policy to the default policies as shown above.</p>
Parameters	<p>mac-monitoring — Enables per SAP + source MAC address rate limiting using the per-source-rate from the associated cpu-protection policy.</p> <p>ip-src-monitoring — Enables per SAP + IP source address rate limiting for certain protocol packets using the per-source-rate and included-protocols from the associated cpu-protection policy. The ip-src-monitoring is useful in subscriber management architectures that have routers between the subscriber and the BNG (router). In layer-3 aggregation scenarios all packets from all subscribers behind the same aggregation router will arrive with the same source MAC address and as such the mac-monitoring functionality can not differentiate traffic from different subscribers.</p>

cpu-protection

Syntax	cpu-protection <i>policy-id</i> [mac-monitoring] [eth-cfm-monitoring [aggregate][car]] [ip-src-monitoring] no cpu-protection
Context	<p>config>service>ies>if>sap config>service>ies>if>spoke-sdp config>service>ies>sub-if>grp-if>sap config>service>vprn>if>sap config>service>vprn>if>spoke-sdp config>service>vprn>sub-if>grp-if>sap</p>

Description	<p>Use this command to apply a specific CPU protection policy to the associated msap-policy. The specified cpu-protection policy will automatically be applied to any MSAPs that are create using the msap-policy.</p> <p>If no CPU-protection policy is assigned to a SAP, then a default policy is used to limit the overall-rate according to the default policy. The default policy is policy number 254 for access interfaces, 255 for network interfaces and no policy for video interfaces.</p> <p>The no form of the command reverts to the default values.</p>
Default	<p>cpu-protection 254 (for access interfaces)</p> <p>cpu-protection 255 (for network interfaces)</p> <p>The configuration of no cpu-protection returns the msap-policy to the default policies as shown above.</p>
Parameters	<p>mac-monitoring — Enables per SAP + source MAC address rate limiting using the per-source-rate from the associated cpu-protection policy.</p> <p>ip-src-monitoring — Enables per SAP + IP source address rate limiting for certain protocol packets using the per-source-rate and include-protocols from the associated cpu-protection policy. The ip-src-monitoring is useful in subscriber management architectures that have routers between the subscriber and the BNG (router). In layer-3 aggregation scenarios all packets from all subscribers behind the same aggregation router will arrive with the same source MAC address and as such the mac-monitoring functionality can not differentiate traffic from different subscribers.</p> <p>eth-cfm-monitoring — Enables the Ethernet Connectivity Fault Management cpu-protection extensions on the associated SAP/SDP/template.</p> <p>aggregate — applies the rate limit to the sum of the per-peer packet rates.</p> <p>car — (Committed Access Rate) Ignores Eth-CFM packets when enforcing overall-rate.</p>

cpu-protection

Syntax	<p>cpu-protection <i>policy-id</i> [mac-monitoring] [eth-cfm-monitoring [aggregate][car]]</p> <p>no cpu-protection</p>
Context	<pre>config>service>epipe>sap config>service>epipe>spoke-sdp config>service>ipipe>sap config>service>template>vpls-sap-template config>service>vpls>mesh-sdp config>service>vpls>sap config>service>vpls>spoke-sdp</pre>
Description	<p>Use this command to apply a specific CPU protection policy to the associated SAP, SDP or template. If the mac-monitoring keyword is given then per MAC rate limiting should be performed, using the per-source-rate from the associated cpu-protection policy.</p>

If no CPU-protection policy is assigned to a SAP, then a default policy is used to limit the overall-rate according to the default policy. The default policy is policy number 254 for access interfaces, 255 for network interfaces and no policy for video interfaces.

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

Default cpu-protection 254 (for access interfaces)

cpu-protection 255 (for network interfaces)

The configuration of no cpu-protection returns the SAP/SDP/template to the default policies as shown above.

Parameters **mac-monitoring** — Enables per SAP + source MAC address rate limiting using the per-source-rate from the associated cpu-protection policy.

eth-cfm-monitoring — Enables the Ethernet Connectivity Fault Management cpu-protection extensions on the associated SAP/SDP/template.

aggregate — applies the rate limit to the sum of the per-peer packet rates.

car — (Committed Access Rate) Ignores Eth-CFM packets when enforcing overall-rate.

2.19.2.21 Distributed CPU Protection Commands

dist-cpu-protection

Syntax	dist-cpu-protection
Context	config>system>security
Description	This command enters the CLI context for configuration of the Distributed CPU Protection (DCP) feature.

policy

Syntax	[no] policy <i>policy-name</i>
Context	config>sys>security>dist-cpu-protection
Description	This command configures one of the maximum 16 Distributed CPU Protection policies. These policies can be applied to objects such as SAPs and network interfaces.
Parameters	<i>policy-name</i> — Name of the policy to be configured.

description

Syntax	[no] description <i>string</i>
Context	config>sys>security>dist-cpu-protection>policy
Description	This command allows you to set the description of the CPU Protection Policy.

rate

Syntax	rate kbps <i>kilobits-per-second</i> max [<i>mbs size</i>] [bytes kilobytes] rate packets { <i>ppi</i> max } within <i>seconds</i> [initial-delay <i>packets</i>] no rate
Context	config>sys>security>dist-cpu-protection>policy>static-policer config>sys>security>dist-cpu-protection>policy>local-monitoring-policer config>sys>security>dist-cpu-protection>policy>protocol>dynamic-parameters
Description	This command configures the rate and burst tolerance for the policer in either a packet rate or a bit rate.

The actual hardware may not be able to perfectly rate limit to the exact configured parameters. In this case, the configured parameters will be adapted to the closest supported rate. The actual (operational) parameters can be seen in CLI, for example, **show service id 33 sap 1/1/3:33 dist-cpu-protection detail**.

Default	rate packets max within 1
Parameters	<p>packets kbps — specifies that the rate is either in units of packets per interval or in units of kilobits per second. The packets option would typically be used for lower rates (for example, for per subscriber DHCP rate limiting) while the kbps option would typically be used for higher rates (for example, per interface BGP rate limiting).</p> <p>ppi — Specifies packets per interval. 0..255 or max (0 = all packets are non-conformant)</p> <ul style="list-style-type: none"> • rate of max = effectively disable the policer (always conformant) • rate of packets 0 = all packets considered non-conformant. <p>within seconds — Specifies the length of the ppi rate measurement interval.</p> <p>Values 1 to 32767</p> <p>initial-delay packets — The number of packets allowed (even at line rate) in an initial burst (or a burst after the policer bucket has drained to zero) in addition to the normal “ppi”. This would typically be set to a value that is equal to the number of received packets in several full handshakes/negotiations of the particular protocol.</p> <p>Values 1 to 255</p> <p>kbps kilobits-per-second — Specifies the kilobits per second.</p> <p>Values 1 to 20000000 max max = This effectively disables the policer (always conformant).</p> <p>mbs — The tolerance for the kbps rate</p> <p>Values 0 to 4194304. A configured mbs of 0 will cause all packets to be considered non-conformant.</p> <p>Default The default mbs sets the mbs to 10 ms of the kbps.</p> <p>bytes kilobytes — Specifies that the units of the mbs size parameter are either in bytes or kilobytes.</p>

detection-time

Syntax	detection-time <i>seconds</i>
Context	config>sys>security>dist-cpu-protection>policy>static-policer

Description	When a policer is declared as in an “exceed” state, it will remain as exceeding until a contiguous conformant period of detection-time passes. The detection-time only starts after the exceed-action hold-down is complete. If the policer detects another exceed during the detection count down then a hold-down is once again triggered before the policer re-enters the detection time (that is, the countdown timer starts again at the configured value). During the hold-down (and the detection-time), the policer is considered as in an “exceed” state.
Default	30
Parameters	<i>seconds</i> — Specifies in seconds.
Values	1 to 128000

dynamic-enforcement-policer-pool

Syntax	[no] dynamic-enforcement-policer-pool <i>number-of-policers</i>
Context	config>card>fp>dist-cpu-protection
Description	This command reserves a set of policers for use as dynamic enforcement policers for the Distributed CPU Protection (DCP) feature. Policers are allocated from this pool and instantiated as per-object-per-protocol dynamic enforcement policers after a local monitor is triggered for an object (such as a SAP or Network Interface). Any change to this configured value automatically clears the high water mark, timestamp, and failed allocation counts as seen under “show card x fp y dist-cpu-protection” and in the <code>tmnxFpDcpDynEnfrcPlcrStatTable</code> in the TIMETRA-CHASSIS-MIB. Decreasing this value to below the currently used/allocated number causes all dynamic policers to be returned to the free pool (and traffic returns to the local monitors).
Default	0
Parameters	<i>number-of-policers</i> — specifies the number of policers to be reserved.
Values	0, 1000 to 32k

exceed-action

Syntax	exceed-action { discard [hold-down <i>seconds</i>] low-priority [hold-down <i>seconds</i>] none }
Context	config>sys>security>dist-cpu-protection>policy>static-policer config>sys>security>dist-cpu-protection>policy>protocol>dynamic-parameters
Description	This command controls the action performed upon the extracted control packets when the configured policer rates are exceeded.
Default	none
Parameters	discard — Discards packets that are non-conformant.

low-priority — Marks packets that are non-conformant as low-priority (for example, discard eligible or out-profile). If there is congestion in the control plane of the SR OS then unmarked (for example, green, hi-prio or in-profile) control packets are given preferential treatment.

hold-down seconds — When this optional parameter is specified, it causes the following “hold-down” behavior.

When the SR OS software detects that an enforcement policer has marked or discarded one or more packets (software may detect this some time after the packets are actually discarded), and an optional **hold-down seconds** value has been specified for the **exceed-action**, then the policer will be set into a “mark-all” or “drop-all” mode that cause the following:

- the policer state to be updated as normal
- all packets to be marked (if the action is “low-priority”) or dropped (action = discard) regardless of the results of the policing decisions/actions/state.

The **hold-down** is cleared after approximately the configured time in seconds after it was set. The **hold-down seconds** option should be selected for protocols that receive more than one packet in a complete handshake/negotiation (for example, DHCP, PPP). **hold-down** is not applicable to a local monitoring policer. The “detection-time” will only start after any **hold-down** is complete. During the **hold-down** (and the detection-time), the policer is considered as in an “exceed” state. The policer may re-enter the hold-down state if an exceed packet is detected during the detection-time countdown.

Configuring the **indefinite** parameter value will cause hold down to remain in place until the operator clears it manually using a tools command (**tools perform security dist-cpu-protection release-hold-down**) or removes the dist-cpu-protection policy from the object.

Configuring the **none** parameter value will disable hold down.

Values 1 to 10080, indefinite, none

exceed-action

Syntax	exceed-action {discard low-priority none}
Context	config>sys>security>dist-cpu-protection>policy>local-monitoring-policer
Description	This command controls the action performed upon the extracted control packets when the configured policer rates are exceeded.
Default	none
Parameters	<p>discard — Discards packets that are non-conformant.</p> <p>low-priority — Marks packets that are non-conformant as low-priority (discard eligible or out-profile). If there is congestion in the control plane of the SR OS then unmarked (green, hi-prio or in-profile) control packets are given preferential treatment.</p> <p>none — no hold-down</p>

log-events

Syntax	[no] log-events [verbose]
Context	config>sys>security>dist-cpu-protection>policy>static-policer
Description	This command controls the creation of log events related to static-policer status and activity.
Default	log-events
Parameters	verbose — (Sends the same events as just “log-events” plus Hold Down Start and Hold Down End events. The optional “verbose” includes some events that are more likely used during debug/tuning/investigations.

local-monitoring-policer

Syntax	[no] local-monitoring-policer <i>policer-name</i> [create]
Context	config>sys>security>dist-cpu-protection>policy>local-monitoring-policer
Description	<p>This command configures a monitoring policer that is used to monitor the aggregate rate of several protocols arriving on an object (for example, SAP). When the local-monitoring-policer is determined to be in a non-conformant state (at the end of a minimum monitoring time of 60 seconds) then the system will attempt to allocate dynamic policers for the particular object for any protocols associated with the local monitor (for example, via the “protocol xyz enforcement” CLI command).</p> <p>If the system cannot allocate all the dynamic policers within 150 seconds, it will stop attempting to allocate dynamic policers, raise a LocMonExcdAllDynAlloc log event, and go back to using the local monitor. The local monitor may then detect exceeded packets again and make another attempt at allocating dynamic policers.</p> <p>Once this <i>policer-name</i> is referenced by a protocol then this policer will be instantiated for each “object” that is created and references this DDoS policy. If there is no policer free then the object will be blocked from being created.</p>
Parameters	<i>policy-name</i> — Specifies name of the policy.
Values	[32 chars max]

log-events

Syntax	[no] log-events [verbose]
Context	config>sys>security>dist-cpu-protection>policy>local-monitoring-policer
Description	This command controls the creation of log events related to local-monitoring-policer status and activity.

Default	log-events
Parameters	verbose — This parameter sends the same events as just “log-events” plus DcpLocMonExcd, DcpLocMonExcdAllDynAlloc, and DcpLocMonExcdAllDynFreed. The optional “verbose” includes some events that are more likely used during debug/tuning/investigations

protocol

Syntax	[no] protocol <i>name</i> [create]
Context	config>sys>security>dist-cpu-protection>policy
Description	This command creates the protocol for control in the policy.

Control packets that are both forwarded (which means they could be subject to normal QoS policy policing) and also copied for extraction are not subject to distributed cpu protection (including in the all-unspecified bucket). This includes traffic snooping (for example, PIM in VPLS) as well as control traffic that is flooded in an R-VPLS instance and also extracted to the CPM (ARP, ISIS and VRRP). Centralized per SAP/interface, cpu-protection can be employed to rate limit or mark this traffic if desired.

Explanatory notes for some of the protocols:

- bfd-cpm: includes all bfd handled on the CPM including cpm-np type, single hop and multi-hop, and MPLS-TP CC and CV bfd
- dhcp: includes dhcp for IPv4 and IPv6
- eth-cfm: 802.1ag and includes Y.1731. Eth-cfm packets on port and LAG based facility MEPs are not included (but packets on Tunnel MEPs are).
- icmp: includes IPv4 and IPv6 ICMP (including RS/RA/Redirect) except NS/NA Neighbor Discovery packets which are classified as a separate protocol 'ndis'
- isis: includes isis used for SPBM
- ldp: includes ldp and t-ldp
- mpls-ttl: MPLS packets that are extracted due to an expired mpls ttl field
- ndis: IPv6 NS/NA Neighbor Discovery (not including RS/RA/Redirect which are classified as part of the protocol 'icmp')
- ospf: includes all OSPFv2 and OSPFv3 packets.
- pppoe-pppoa: includes PADx, LCP, PAP/CHAP and NCPs

- **all-unspecified:** a special 'protocol'. When configured, this treats all extracted control packets that are not explicitly created in the dist-cpu-protection policy as a single aggregate flow (or "virtual protocol"). It lumps together "all the rest of the control traffic" to allow it to be rate limited as one flow. It includes all control traffic of all protocols that are extracted and sent to the CPM (even protocols that cannot be explicitly configured with the distributed cpu protection feature). Control packets that are both forwarded and copied for extraction are not included. If an operator later explicitly configures a protocol, then that protocol is suddenly no longer part of the "all-unspecified" flow. The "all-unspecified" protocol must be explicitly configured in order to operate.

"no protocol x" means packets of protocol x are not monitored and not enforced (although they do count in the fp protocol queue) on the objects to which this dist-cpu-protection policy is assigned, although the packets will be treated as part of the all-unspecified protocol if the all-unspecified protocol is created in the policy.

Default none

Parameters *names* — Signifies protocol name.

Values arp | dhcp | http-redirect | icmp | igmp | mld | ndis | pppoe-pppoa | all-unspecified | mpls-ttl | bfd-cpm | bgp | eth-cfm | isis | ldp | ospf | pim | rsvp.

enforcement

Syntax **enforcement** {**static** *policer-name* | **dynamic** {*mon-policer-name* | **local-mon-bypass**}}

Context config>sys>security>dist-cpu-protection>policy>protocols

Description This command configures the enforcement method for the protocol.

Default dynamic local-mon-bypass

Parameters **static** — the protocol is always enforced using a static-policer. Multiple protocols can reference the same static-policer. Packets of protocols that are statically enforced bypass any local monitors.

policer name — Specifies the name is a static-policer.

dynamic — A specific enforcement policer for this protocol for this SAP/object is instantiated when the associated local-monitoring-policer is determined to be in a non-conformant state (at the end of a minimum monitoring time of 60 seconds to reduce thrashing).

mon-policer-name — Specifies which local-monitoring-policer to use

local-mon-bypass — This parameter is used to not include packets from this protocol in the local monitoring function, and when the local-monitor "trips", do not instantiate a dynamic enforcement policer for this protocol.

detection-time

Syntax	detection-time <i>seconds</i>
Context	config>sys>security>dist-cpu-protection>policy>protocols>dynamic-parameters
Description	When a dynamic enforcing policer is instantiated, it will remain allocated until at least a contiguous conformant period of detection-time passes.

dynamic-parameters

Syntax	dynamic-parameters
Context	config>sys>security>dist-cpu-protection>policy>protocols
Description	The dynamic-parameters are used to instantiate a dynamic enforcement policer for the protocol when the associated local-monitoring-policer is considered as exceeding its rate parameters (at the end of a minimum monitoring time of 60 seconds).

log-events

Syntax	[no] log-events [verbose]
Context	config>sys>security>dist-cpu-protection>policy>protocols>dynamic-parameters
Description	This command controls the creation of log events related to dynamic enforcement policer status & activity
Default	log-events
Parameters	verbose — This parameter sends the send the same events as just “log-events” plus Hold Down Start, Hold Down End, DcpDynamicEnforceAlloc and DcpDynamicEnforceFreed events. This includes the allocation/de-allocation events (typically used for debug/tuning only – could be very noisy even when there is nothing much of concern).

static-policer

Syntax	[no] static-policer <i>policer-name</i> [create]
Context	config>sys>security>dist-cpu-protection>policy
Description	Configures a static enforcement policer that can be referenced by one or more protocols in the policy. Once this policer-name is referenced by a protocol, then this policer will be instantiated for each object (e.g. SAP or network interface) that is created and references this policy. If there is no policer resource available on the associated card/fp then the object will be blocked from being created. Multiple protocols can use the same static-policer.

Parameters *policy-name* — Specifies the name of the policy.
Values [32 chars max]

2.19.2.22 Extracted Protocol Traffic Priority Commands

init-extract-prio-mode

Syntax `init-extract-prio-mode {uniform | l3-classify}`

Context `config>card>fp`

Description This command determines the scheme used to select the initial drop priority of extracted control plane traffic. The initial drop priority of extracted packets can be either low or high priority. The drop priority of the extracted packets can be subsequently altered by mechanisms such as CPU protection. High-priority traffic receives preferential treatment in control plane congestion situations over low-priority traffic.

Default `uniform`

Parameters

uniform — Initializes the drop priority of all extracted control traffic as high priority. Drop priority can then be altered (marked low priority) by distributed CPU protection (DCP) or centralized CPU protection rate-limiting functions in order to achieve protocol and interface isolation.

l3-classify — Initializes the drop priority of Layer 3 extracted control traffic (BGP and OSPF) based on the QoS classification of the packets. This is useful in networks where the DSCP and EXP markings can be trusted as the primary method to distinguish, protect, and isolate good terminating protocol traffic from unknown or potentially harmful protocol traffic instead of using the rate-based DCP and centralized CPU protection traffic marking/coloring mechanisms (for example, **out-profile-rate** and **exceed-action low-priority**).

For network interfaces, the QoS classification profile result selects the drop priority (in = high priority, out = low priority) for extracted control traffic, and the default QoS classification maps different DSCP and EXP values to different in/out profile states. For access interfaces, the QoS classification priority result typically selects the drop priority for extracted control traffic. The default access QoS classification (**default-priority**) maps all traffic to **low**. If the queues in the access QoS policy are configured as **profile-mode** queues (rather than the default **priority-mode**) extracted traffic will use the QoS classification profile value configured against the associated FC (rather than the priority result) to select the drop priority.

Layer 2 extracted control traffic (ARP or ETH-CFM) and protocols that cannot always be QoS-classified, such as IS-IS, are initialized as low drop priority in order to protect Layer 2 protocol traffic on uniform interfaces (which would typically be subject to centralized CPU protection). Alternately, DCP can be used (by configuring a non-zero rate with **exceed-action** of **low-priority** for the **all-unspecified** protocol) to mark some of this traffic as high priority.

2.20 Security Show, Clear, Debug, Tools, and Admin Command Reference

2.20.1 Command Hierarchies

- [Show Commands](#)
- [Clear Commands](#)
- [Debug Commands](#)
- [Tools Commands](#)
- [Admin Commands](#)

2.20.1.1 Show Commands

2.20.1.1.1 Security

```
show
  — system
    — security
      — access-group [group-name]
      — authentication [statistics]
      — communities
      — cpm-filter
        — ip-filter [entry entry-id]
        — ipv6-filter [entry entry-id]
        — mac-filter [entry entry-id]
      — cpm-queue queue-id
      — cpu-protection
        — eth-cfm-monitoring [{service-id service-id sap-id sap-id} | {service-id
          service-id sdp-id sdp-id:vc-id}]
        — excessive-sources [service-id service-id sap-id sap-id]
        — policy [policy-id] association
        — protocol-protection
        — violators [port] [interface] [sap] [video] [sdp]
      — dist-cpu-protection
        — policy [policy-id] [association detail]
      — keychain keychain-name [detail]
      — management-access-filter
        — ip-filter [entry entry-id]
        — ipv6-filter [entry entry-id]
        — mac-filter [entry entry-id]
      — password-options
```

```

    — per-peer-queuing [detail]
    — per-peer-queuing
    — profile [user-profile-name]
    — source-address
    — ssh
    — user [user-name] [detail]
    — user [user-name] lockout
    — view [view-name] [detail]
  — certificate
    — ca-profile
    — ca-profile name [association]
    — ocsp-cache [entry-id]
    — statistics

show
  — card
    — fp
      — dist-cpu-protection

show
  — service
    — id
      — sap
        — dist-cpu-protection [detail]

show
  — router
    — interface
      — dist-cpu-protection [detail]

```

2.20.1.1.2 Login Control

```

show
  — users

```

2.20.1.2 Clear Commands

```

clear
  — router
    — authentication
      — statistics [interface ip-int-name | ip-address]
      — radius-proxy-server server-name statistics
    — cpm-filter
      — ip-filter [entry entry-id]
      — ipv6-filter [entry entry-id]
      — mac-filter [entry entry-id]
    — cpu-protection
      — excessive-sources

```

- **protocol-protection**
 - **violators** [port] [interface] [sap]
 - **cpm-queue** *queue-id*
- admin
- user
 - user
 - **clear lockout** {*name* | all}
 - **clear password-history** {*name* | all}

2.20.1.3 Debug Commands

- debug
- **radius** [detail] [hex]
 - **no radius**
 - [no] **ocsp**
 - [no] **ocsp** *profile-name*

2.20.1.4 Tools Commands

- tools
- dump
 - security
 - **dist-cpu-protection**
 - **violators enforcement** {sap | interface} card *slot-number* [fp *fp-number*]
 - **violators local-monitor** {sap | interface} card *slot-number* [fp *fp-number*]
 - perform
 - security
 - **dist-cpu-protection**
 - **release-hold-down interface** *interface-name* [protocol *protocol*] [static-policer *name*]
 - **release-hold-down sap** *sap-id* [protocol *protocol*] [static-policer *name*]

2.20.1.5 Admin Commands

- admin
- user
 - **clear lockout** {*user name* | all}
 - **clear password-history** {*user name* | all}

2.20.2 Command Descriptions

- [Show Commands](#)
- [Clear Commands](#)
- [Debug Commands](#)
- [Tools Commands](#)
- [Admin Commands](#)

2.20.2.1 Show Commands

The command outputs in the following section are examples only; actual displays may differ depending on supported functionality and user configuration.

2.20.2.1.1 Security Commands

access-group

- Syntax** `access-group [group-name]`
- Context** `show>system>security`
- Description** This command displays SNMP access group information.
- Parameters** *group-name* — This command displays information for the specified access group.
- Output** Security Access Group Output

[Table 20](#) describes security access group output fields..

Table 20 Show System Security Access Group Output Fields

Label	Description
Group name	The access group name.
Security model	The security model required to access the views configured in this node.
Security level	Specifies the required authentication and privacy levels to access the views configured in this node.
Read view	Specifies the variable of the view to read the MIB objects.

Table 20 Show System Security Access Group Output Fields (Continued)

Label	Description
Write view	Specifies the variable of the view to configure the contents of the agent.
Notify view	Specifies the variable of the view to send a trap about MIB objects.

Sample Output

```
A:ALA-4# show system security access-group
=====
Access Groups
=====
group name      security  security  read      write      notify
                 model    level    view      view      view
-----
snmp-ro         snmpv1   none     no-security          no-security
snmp-ro         snmpv2c  none     no-security          no-security
snmp-rw         snmpv1   none     no-security  no-security  no-security
snmp-rw         snmpv2c  none     no-security  no-security  no-security
snmp-rwa        snmpv1   none     iso           iso          iso
snmp-rwa        snmpv2c  none     iso           iso          iso
snmp-trap       snmpv1   none     snmpv1        iso          iso
snmp-trap       snmpv2c  none     snmpv2c       iso          iso
=====
A:ALA-7#
```

authentication

Syntax authentication [statistics]

Context show>system>security

Description This command displays system login authentication configuration and statistics.

Parameters **statistics** — Appends login and accounting statistics to the display.

Output Authentication Output

[Table 21](#) describes system security authentication output fields.

Table 21 Show System Security Authentication Output Fields

Label	Description
Sequence	The sequence in which authentication is processed.
Server address	The IP address of the RADIUS server.
Status	Current status of the RADIUS server.

Table 21 Show System Security Authentication Output Fields

Label	Description
Type	The authentication type.
Timeout (secs)	The number of seconds the router waits for a response from a RADIUS server.
Retry count	Displays the number of times the router attempts to contact the RADIUS server for authentication if there are problems communicating with the server.
Connection errors	Displays the number of times a user has attempted to login irrespective of whether the login succeeded or failed.
Accepted logins	The number of times the user has successfully logged in.
Rejected logins	The number of unsuccessful login attempts.
Sent packets	The number of packets sent.
Rejected packets	The number of packets rejected.

Sample Output

```
A:ALA-4# show system security authentication
=====
Authentication          sequence : radius tacplus local ldap exit-on-reject
=====
type                    status  timeout (secs)  retry count
  server address
  server name
-----
radius                  down    3                3
  192.170.0.30
  n/a
ldap                    up      3                3
  192.170.0.10(389)
  my_first_LDAP_server
ldap                    down    3                3
  0.0.0.0(389)
  n/a
-----
radius admin/oper status : up/down
ldap admin/oper status  : up/up
health check            : enabled (interval 30 secs)
-----
No. of Servers: 3
=====

A:ALA-4# show system security authentication statistics
=====
Authentication          sequence : radius tacplus ldap local
=====
type                    status  timeout (secs)  retry count
```

```

server address
server name
-----
ldap                               down                               3                               3
  135.243.194.179:10390
  n/a
-----

```

```

ldap admin/oper status      : down/down
health check                : enabled (interval 30 secs)
-----

```

No. of Servers: 1

=====
Login Statistics

```

server address                conn  accepted  rejected
                             errors logins   logins
-----
135.243.194.179              0     2         7
local                         n/a   10        8
=====

```

=====
Authorization Statistics (TACACS+)

```

server address                conn  sent      rejected
                             errors pkts    pkts
-----

```

=====
Accounting Statistics

```

server address                conn  sent      rejected
                             errors pkts    pkts
-----

```

A:ALA-4# show system security authentication

```

=====
Authentication                sequence : radius tacplus local ldap exit-on-reject
=====
type                            status  timeout (secs)  retry count
server address
server name
-----
radius                          up      5                5
  10.10.10.103
  n/a
radius                          up      5                5
  10.10.10.1
  n/a
radius                          up      5                5
  10.10.10.2
  n/a
radius                          up      5                5
  10.10.10.3
  n/a
-----

```

```

radius admin status : up
tacplus admin status : up
health check        : enabled (interval 30)
-----

```

```
No. of Servers: 4
=====
A:ALA-4#
```

```
A:ALA-7>show>system>security# authentication statistics
=====
Authentication                sequence : radius tacplus local
=====
type                            status  timeout (secs)  retry count
server address
-----
radius                          up      5                5
 10.10.10.103
radius                          up      5                5
 10.10.10.1
radius                          up      5                5
 10.10.10.2
radius                          up      5                5
 10.10.10.3
-----
radius admin status   : up
tacplus admin status : up
health check         : enabled (interval 30)
-----
```

```
No. of Servers: 4
=====
Login Statistics
=====
server address      connection errors  accepted logins  rejected logins
-----
10.10.10.103       0                  0                0
10.10.0.1          0                  0                0
10.10.0.2          0                  0                0
10.10.0.3          0                  0                0
local              n/a                1                0
=====
```

```
Authorization Statistics (TACACS+)
=====
server address      connection errors  sent packets     rejected packets
-----
```

```
Accounting Statistics
=====
server address      connection errors  sent packets     rejected packets
-----
10.10.10.103       0                  0                0
10.10.0.1          0                  0                0
10.10.0.2          0                  0                0
10.10.0.3          0                  0                0
=====
```

```
A:ALA-7#

*A:Dut-C# show system security authentication statistics

=====
Authentication                sequence : radius tacplus local
=====
```

```

type                status  timeout (secs)  retry count
server address
-----
radius              up      5                5
 10.10.10.103
radius              up      5                5
 10.10.10.1
radius              up      5                5
 10.10.10.2
radius              up      5                5
 10.10.10.3
-----
radius admin status : up
tacplus admin status : up
health check       : enabled (interval 30)
-----
No. of Servers: 4
=====

Login Statistics
=====
server address                conn  accepted  rejected
                              errors logins   logins
-----
local                          n/a    4          0
=====

Authorization Statistics (TACACS+)
=====
server address                conn  sent      rejected
                              errors pkts   pkts
-----

Accounting Statistics
=====
server address                conn  sent      rejected
                              errors pkts   pkts
=====

```

communities

- Syntax** **communities**
 - Context** show>system>security
 - Description** This command displays SNMP communities.
 - Output** Communities Output
- [Table 22](#) describes community output fields.

Table 22 Show Communities Output Fields

Label	Description
Community	The community string name for SNMPv1 and SNMPv2c access only.
Access	r — The community string allows read-only access.
	rw — The community string allows read-write access.
	rwa — The community string allows read-write access.
	mgmt — The unique SNMP community string assigned to the management router.
View	The view name.
Version	The SNMP version.
Group Name	The access group name.
No of Communities	The total number of configured community strings.

Sample Output

```
A:ALA-48# show system security communities
=====
Communities
=====
community      access  view          version  group name
-----
cli-readonly   r       iso           v2c     cli-readonly
cli-readwrite  rw      iso           v2c     cli-readwrite
public         r       no-security   v1 v2c  snmp-ro
-----
No. of Communities: 3
=====
A:ALA-48#
```

cpm-filter

- Syntax** **cpm-filter**
- Context** show>system>security
- Description** This command displays CPM filters.

ip-filter

- Syntax** `ip-filter [entry entry-id]`
- Context** `show>system>security>cpm-filter`
- Description** This command displays CPM IP filters.
- Parameters** *entry-id* — Identifies a CPM filter entry as configured on this system.
- Values** 1 to 6144
- Output** CPM Filter Output

[Table 23](#) describes CPM IP filter output fields.

Table 23 Show CPM IP Filter Output Fields

Label	Description
Entry-Id	Displays information about the specified management access filter entry
Dropped	Displays the number of dropped events.
Forwarded	Displays the number of forwarded events.
Description	Displays the CPM filter description.
Log ID	Displays the log ID where matched packets will be logged.
Src IP	Displays the source IP address(/netmask or prefix-list)
Dest. IP	Displays the destination IP address(/netmask).
Src Port	Displays the source port number (range).
Dest. Port	Displays the destination port number (range).
Protocol	Displays the Protocol field in the IP header.
Dscp	Displays the DSCP field in the IP header.
Fragment	Displays the 3-bit fragment flags or 13-bit fragment offset field.
ICMP Type	Displays the ICMP type field in the ICMP header.
ICMP Code	Displays the ICMP code field in the ICMP header.
TCP-syn	Displays the SYN flag in the TCP header.
TCP-ack	Displays the ACK flag in the TCP header
Match action	When the criteria matches, displays drop or forward packet.

Table 23 Show CPM IP Filter Output Fields (Continued)

Label	Description
Next Hop	In case match action is forward, indicates destination of the matched packet.
Dropped pkts	Indicates number of matched dropped packets
Forwarded pkts	Indicates number of matched forwarded packets.

Sample Output

```
A:ALA-35# show system security cpm-filter ip-filter
=====
CPM IP Filters
=====
Entry-Id  Dropped  Forwarded  Description
-----
101        25880     0          CPM-Filter 10.4.101.2 #101
102        25880     0          CPM-Filter 10.4.102.2 #102
103        25880     0          CPM-Filter 10.4.103.2 #103
104        25882     0          CPM-Filter 10.4.104.2 #104
105        25926     0          CPM-Filter 10.4.105.2 #105
106        25926     0          CPM-Filter 10.4.106.2 #106
107        25944     0          CPM-Filter 10.4.107.2 #107
108        25950     0          CPM-Filter 10.4.108.2 #108
109        25968     0          CPM-Filter 10.4.109.2 #109
110        25984     0          CPM-Filter 10.4.110.2 #110
111        26000     0          CPM-Filter 10.4.111.2 #111
112        26018     0          CPM-Filter 10.4.112.2 #112
113        26034     0          CPM-Filter 10.4.113.2 #113
114        26050     0          CPM-Filter 10.4.114.2 #114
115        26066     0          CPM-Filter 10.4.115.2 #115
116        26084     0          CPM-Filter 10.4.116.2 #116
=====
A:ALA-35#

A:ALA-35# show system security cpm-filter ip-filter entry 101
=====
CPM IP Filter Entry
=====
Entry Id      : 101
Description   : CPM-Filter 10.4.101.2 #101
-----
Filter Entry Match Criteria :
-----
Log Id        : n/a
Src. IP       : 10.4.101.2/32      Src. Port     : 0
Dest. IP      : 10.4.101.1/32      Dest. Port    : 0
Protocol      : 6                    Dscp          : ef
ICMP Type     : Undefined        ICMP Code     : Undefined
Fragment      : True             Option-present : Off
IP-Option     : 130/255          Multiple Option : True
TCP-syn       : Off             TCP-ack       : True
Match action   : Drop
=====
```


A:ALA-35#

ipv6-filter

- Syntax** `ip-filter [entry entry-id]`
- Context** `show>system>security>cpm-filter`
- Description** This command displays CPM IPv6 filters and only applies to the 7750 SR and 7950 XRS.
- Parameters** *entry-id* — Identifies a CPM IPv6 filter entry as configured on this system.
Values 1 to 6144
- Output** CPM Filter Output

[Table 24](#) describes CPM IPv6 filter output fields.

Table 24 Show CPM IPv6 Filter Output Fields

Label	Description
Entry-Id	Displays information about the specified management access filter entry
Dropped	Displays the number of dropped events.
Forwarded	Displays the number of forwarded events.
Description	Displays the CPM filter description.
Log ID	Log Id where matched packets will be logged.
Src IP	Displays Source IP address(/netmask)
Dest. IP	Displays Destination IP address(/netmask).
Src Port	Displays Source Port Number (range).
Dest. Port	Displays Destination Port Number (range).
next-header	Displays next-header field in the IPv6 header.
Dscp	Displays Traffic Class field in the IPv6 header.
ICMP Type	Displays ICMP type field in the icmp header.
ICMP Code	Displays ICMP code field in the icmp header.
TCP-syn	Displays the SYN flag in the TCP header.
TCP-ack	Displays the ACK flag in the TCP header
Match action	When criteria matches, displays drop or forward packet.

Table 24 Show CPM IPv6 Filter Output Fields (Continued)

Label	Description
Next Hop	In case match action is forward, indicates destination of the matched packet.
Dropped pkts	Indicating number of matched dropped packets
Forwarded pkts	Indicating number of matched forwarded packets.

The following is an output example on the 7750 SR:

```
A:ALA-35# show system security cpm-filter ipv6-filter
=====
CPM IPv6 Filters
=====
Entry-Id Dropped Forwarded Description
-----
101      25880    0      CPM-Filter 11::101:2 #101
102      25880    0      CPM-Filter 11::102:2 #102
103      25880    0      CPM-Filter 11::103:2 #103
104      25880    0      CPM-Filter 11::104:2 #104
105      25880    0      CPM-Filter 11::105:2 #105
106      25880    0      CPM-Filter 11::106:2 #106
107      25880    0      CPM-Filter 11::107:2 #107
108      25880    0      CPM-Filter 11::108:2 #108
109      25880    0      CPM-Filter 11::109:2 #109
=====
A:ALA-35#
```

```
A:ALA-35# show system security cpm-filter ipv6-filter entry 101
=====
CPM IPv6 Filter Entry
=====
Entry Id : 1
Description : CPM-Filter 11::101:2 #101
-----
Filter Entry Match Criteria :
-----
Log Id : n/a
Src. IP : 11::101:2      Src. Port : 0
Dest. IP : 11::101:1    Dest. Port : 0
next-header : none      Dscp : Undefined
ICMP Type : Undefined   ICMP Code : Undefined
TCP-syn : Off          TCP-ack : Off
Match action : Drop
Dropped pkts : 25880    Forwarded pkts : 0
=====
A:ALA-35#
```

cpm-queue

- Syntax** `cpm-queue queue-id`
- Context** `show>system>security`
- Description** This command displays CPM queues.
- Parameters** *queue-id* — Specifies an integer value that identifies a CPM queue.
- Values** 0, 33 to 2000
- Output** CPM queue Output

[Table 25](#) describes CPM queue output fields..

Table 25 Show CPM IPv6 Filter Output Fields

Label	Description
PIR	Displays the administrative Peak Information Rate (PIR) for the queue.
CIR	Displays the amount of bandwidth committed to the queue.
CBS	Displays the amount of buffer drawn from the reserved buffer portion of the queue's buffer pool.
MBS	Displays the maximum queue depth to which a queue can grow.

Sample Output

```
A:ALA-35# show system security cpm-queue 1001
=====
CPM Queue Entry
=====
Queue Id          : 1001
-----
Queue Parameters :
-----
PIR                : 10000000          CIR                : 1000000
CBS                : 4096             MBS                : 8192
=====
A:ALA-35#
```

cpu-protection

- Syntax** `cpu-protection`
- Context** `show>system>security`
- Description** This command enables the context to display CPU protection information.

Output The following output is an example of ETH CFM monitoring

Sample Output

```

show system security cpu-protection eth-cfm-monitoring
=====
SAP's where the protection policy Eth-CFM rate limit is exceeded
=====
SAP-Id                               Service-Id  Plcy
-----
1/1/1                                 3           100
-----
1 SAP('s) found
=====
SDP's where the protection policy Eth-CFM rate limit is exceeded
=====
SDP-Id          Service-Id  Plcy
-----
1:3             3           100
-----
1 SDP('s) found
=====

show system security cpu-protection eth-cfm-monitoring service-id 3 sap-id 1/1/1
=====
Flows exceeding the Eth-CFM monitoring rate limit
=====
Service-Id : 3
SAP-Id      : 1/1/1
Plcy        : 100
-----
Limit  MAC-Address          Level  OpCode
  First-Time              Last-Time              Violation-Periods
-----
0      8c:8c:8c:8c:8c:8c    1      18
      03/21/2009 23:32:29  03/21/2009 23:34:39  4000000019
61234  8d:8d:8d:8d:8d:8d    2      19
      03/21/2009 23:32:39  03/21/2009 23:34:59  4000000020
61234  Aggregated          3      20
      03/21/2009 23:32:49  03/21/2009 23:35:19  4000000021
61234  8f:8f:8f:8f:8f:8f    4      21
      03/21/2009 23:32:59  03/21/2009 23:35:39  4000000022
61234  90:90:90:90:90:90    5      22
      03/21/2009 23:33:09  03/21/2009 23:35:59  4000000023
61234  91:91:91:91:91:91    6      23
      03/21/2009 23:33:19  03/21/2009 23:36:19  4000000024
61234  92:92:92:92:92:92    7      24
      03/21/2009 23:33:29  03/21/2009 23:36:39  4000000025
max    Aggregated          0      25
      03/21/2009 23:33:39  03/21/2009 23:36:59  4000000026
0      94:94:94:94:94:94    1      26
      03/21/2009 23:33:49  03/21/2009 23:37:19  4000000027
-----
9 flows(s) found
=====

```

```
show system security cpu-protection eth-cfm-monitoring service-id 3 sdp-id 1:3
=====
Flows exceeding the Eth-CFM monitoring rate limit
=====
Service-Id : 3
SDP-Id      : 1:3
Plcy        : 100
-----
Limit  MAC-Address      Level  OpCode
  First-Time          Last-Time          Violation-Periods
-----
0      8c:8c:8c:8c:8c:8c  1      18
      03/21/2009 23:32:29  03/21/2009 23:34:39  3000000019
61234  8d:8d:8d:8d:8d:8d  2      19
      03/21/2009 23:32:39  03/21/2009 23:34:59  3000000020
61234  Aggregated          3      20
      03/21/2009 23:32:49  03/21/2009 23:35:19  3000000021
61234  8f:8f:8f:8f:8f:8f  4      21
      03/21/2009 23:32:59  03/21/2009 23:35:39  3000000022
61234  90:90:90:90:90:90  5      22
      03/21/2009 23:33:09  03/21/2009 23:35:59  3000000023
61234  91:91:91:91:91:91  6      23
      03/21/2009 23:33:19  03/21/2009 23:36:19  3000000024
61234  92:92:92:92:92:92  7      24
      03/21/2009 23:33:29  03/21/2009 23:36:39  3000000025
max    Aggregated          0      25
      03/21/2009 23:33:39  03/21/2009 23:36:59  3000000026
0      94:94:94:94:94:94  1      26
      03/21/2009 23:33:49  03/21/2009 23:37:19  3000000027
-----
9 flow(s) found
=====
```

```
show system security cpu-protection excessive-sources service-id 3 sdp-id 1:3
=====
Sources exceeding the per-source rate limit
=====
Service-Id : 3
SDP-Id      : 1:3
Plcy        : 100
Limit       : 65534
-----
MAC-Address      First-Time          Last-Time          Violation-Periods
-----
00:00:00:00:00:01 03/22/2009 00:41:59 03/22/2009 01:53:39 3000000043
00:00:00:00:00:02 03/22/2009 00:43:39 03/22/2009 01:56:59 3000000044
00:00:00:00:00:03 03/22/2009 00:45:19 03/22/2009 02:00:19 3000000045
00:00:00:00:00:04 03/22/2009 00:46:59 03/22/2009 02:03:39 3000000046
00:00:00:00:00:05 03/22/2009 00:48:39 03/22/2009 02:06:59 3000000047
-----
5 source(s) found
=====
```

```
show system security cpu-protection violators sdp
=====
SDP's where the protection policy overall rate limit is violated
=====
```

SDP-Id	Service-Id	Plcy	Limit	First-Time	Last-Time	Violation-Periods
1:1	3					
100	61234	05/01/2010	01:43:53	06/27/2010	22:37:20	3000000007
1:2	3					
255	max	05/01/2010	01:43:55	06/27/2010	22:37:23	3000000008
1:3	3					
100	61234	05/01/2010	01:43:57	06/27/2010	22:37:26	3000000009
1:4	3					
255	max	05/01/2010	01:43:59	06/27/2010	22:37:29	3000000010
1:5	3					
100	61234	05/01/2010	01:44:01	06/27/2010	22:37:32	3000000011

5 SDP('s) found

show system security cpu-protection excessive-sources

SAP's where the protection policy per-source rate limit is exceeded

SAP-Id	Service-Id	Plcy	Limit
1/1/1	3		
100	65534		

1 SAP('s) found

SDP's where the protection policy per-source rate limit is exceeded

SDP-Id	Service-Id	Plcy	Limit
1:3	3	100	65534
1:4	3	255	max
1:5	3	100	65534

3 SDP('s) found

show system security cpu-protection policy association

Associations for CPU Protection policy 100

Description : (Not Specified)

SAP associations

Service Id	Type
SAP 1/1/1	mac-monitoring
SAP 1/1/2	eth-cfm-monitoring aggr car
SAP 1/1/3	eth-cfm-monitoring
SAP 1/1/4	

Number of SAP's : 4

SDP associations

Service Id	Type
3	VPLS

```
SDP 1:1          eth-cfm-monitoring aggr car
SDP 1:3          eth-cfm-monitoring aggr
SDP 1:5          mac-monitoring
SDP 17407:4123456789 eth-cfm-monitoring car
-----
Number of SDP's : 4
Interface associations
-----
None
Managed SAP associations
-----
None
Video-Interface associations
-----
None
=====
Associations for CPU Protection policy 254
=====
Description : Default (Modifiable) CPU-Protection Policy assigned to Access
              Interfaces
SAP associations
-----
None
SDP associations
-----
None
Interface associations
-----
Router-Name : Base
              ies6If
Router-Name : vprn7
              vprn7If
-----
Number of interfaces : 2
Managed SAP associations
-----
None
Video-Interface associations
-----
None
=====
Associations for CPU Protection policy 255
=====
Description : Default (Modifiable) CPU-Protection Policy assigned to Network
              Interfaces
SAP associations
-----
None
SDP associations
-----
Service Id : 3          Type : VPLS
  SDP 1:2
  SDP 1:4          eth-cfm-monitoring
Service Id : 6          Type : IES
  SDP 1:6
Service Id : 7          Type : VPRN
  SDP 1:7
Service Id : 9          Type : Epipe
```

```

SDP 1:9
Service Id : 300                               Type : VPLS
SDP 1:300
-----
Number of SDP's : 6
Interface associations
-----
Router-Name : Base
              system
-----
Number of interfaces : 1
Managed SAP associations
-----
None
Video-Interface associations
-----
None
=====

```

```

show system security cpu-protection policy 100 association
=====
Associations for CPU Protection policy 100
=====
Description : (Not Specified)

```

```

SAP associations
-----
Service Id : 3                               Type : VPLS
SAP 1/1/1                                     mac-monitoring
SAP 1/1/2                                     eth-cfm-monitoring aggr car
SAP 1/1/3                                     eth-cfm-monitoring
SAP 1/1/4
-----

```

```

Number of SAP's : 4
SDP associations
-----
Service Id : 3                               Type : VPLS
SDP 1:1                                     eth-cfm-monitoring aggr car
SDP 1:3                                     eth-cfm-monitoring aggr
SDP 1:5                                     mac-monitoring
SDP 17407:4123456789 eth-cfm-monitoring car
-----

```

```

Number of SDP's : 4
Interface associations
-----
None
Managed SAP associations
-----
None
Video-Interface associations
-----
None
=====

```

A:bksim130#

```

show system security cpu-protection violators
=====

```



```

Ports where a rate limit is violated
=====
Port-Id
  Type Limit First-Time          Last-Time          Violation-Periods
-----
No ports found
=====
Interfaces where the protection policy overall rate limit is violated
=====
Interface-Name          Router-Name
  Plcy Limit First-Time          Last-Time          Violation-Periods
-----
No interfaces found
=====
SAP's where the protection policy overall rate limit is violated
=====
SAP-Id          Service-Id
  Plcy Limit First-Time          Last-Time          Violation-Periods
-----
1/1/1
  100 61234 05/01/2010 01:43:41 06/27/2010 22:37:02 3000000001
-----
1 SAP('s) found
=====
SDP's where the protection policy overall rate limit is violated
=====
SDP-Id          Service-Id
  Plcy Limit First-Time          Last-Time          Violation-Periods
-----
1:1
  100 61234 05/01/2010 01:43:41 06/27/2010 22:37:02 3000000001
1:2
  255 max 05/01/2010 01:43:43 06/27/2010 22:37:05 3000000002
1:3
  100 61234 05/01/2010 01:43:45 06/27/2010 22:37:08 3000000003
1:4
  255 max 05/01/2010 01:43:47 06/27/2010 22:37:11 3000000004
1:5
  100 61234 05/01/2010 01:43:49 06/27/2010 22:37:14 3000000005
-----
5 SDP('s) found
=====
Video clients where the protection policy per-source rate limit is violated
=====
Client IP Address  Video-Interface          Service-Id
  Plcy Limit First-Time          Last-Time          Violation-Periods
-----
No clients found
=====

```

eth-cfm-monitoring

Syntax `eth-cfm-monitoring [{service-id service-id sap-id sap-id} | {service-id service-id sdp-id sdp-id:vc-id}]`

Context `show>system>security>cpu-protection`

Description This command displays sources exceeding their eth-cfm-monitoring rate limit.

dist-cpu-protection

Syntax `dist-cpu-protection`

Context `show>card>fp`

Description This command displays Distributed CPU Protection parameters and status at the per card and forwarding plane level.

Output [Table 26](#) describes Distributed CPU Protection output fields.

Table 26 Show Distributed CPU Protection Output Fields

Label	Description
Card	The card identifier
Forwarding Plane(FP)	Identifies the instance of the FP (FastPath) chipset. Some cards have a single FP (for example, an IOM3-XP) and some cards can contain multiple FPs (for example, an XCM can house two FPs via its two XMA's).
Dynamic Enforcement Policer Pool	The configured size of the dynamic-enforcement-policer-pool for this card/FP.
Dynamic-Policers Currently In Use	The number of policers from the dynamic enforcement policer pool that are currently in use. The policers are allocated from the pool and instantiated as per-object-per-protocol dynamic enforcement policers after a local monitor triggered for an object (such as a SAP or Network Interface).
Hi-WaterMark Hit Count	The maximum Currently In Use value since it was last cleared (clear card x fp y dist-cpu-protection)
Hi-WaterMark Hit Time	The time at which the current Hi-WaterMark Hit Count was first recorded.
Dynamic-Policers Allocation Fail Count	Indicates how many times the system attempted to allocate dynamic enforcement policers but could not get enough the fill the request.

Sample Output

```
*A:nodeA# show card 1 fp 1 dist-cpu-protection
=====
Card : 1 Forwarding Plane(FP) : 1
=====
Dynamic Enforcement Policer Pool : 2000
-----
```

```

-----
Statistics Information
-----
Dynamic-Policers Currently In Use      : 48
Hi-WaterMark Hit Count                : 72
Hi-WaterMark Hit Time                 : 01/03/2013 15:08:42 UTC
Dynamic-Policers Allocation Fail Count : 0
-----
=====

```

dist-cpu-protection

- Syntax** `dist-cpu-protection [detail]`
- Context** `show>service>id>sap`
- Description** This command displays Distributed CPU Protection parameters and status at the per SAP level.
- Parameters** *detail* — Include the adapted operational rate parameters in the CLI output. The adapted Oper. parameters are only applicable if the policer is instantiated (for example, if the associated forwarding plane is operational, or for an interface if there is a physical port configured for the interface, or if the dynamic policers are allocated), otherwise values of 0 kb/s, etc are displayed.
- Output** Distributed CPU Protection Policer Output
- [Table 27](#) describes Distributed CPU Protection Policer Output output fields.

Table 27 Show Distributed CPU Protection Policer Output Fields

Label	Description
Distributed CPU Protection Policy	The DCP policy assigned to the object.
Policer-Name	The configured name of the static policer
Card/FP	The card and FP identifier. FP identifies the instance of the FP (FastPath) chipset. Some cards have a single FP (for example, IOM3-XP) and some cards can contain multiple FPs (for example, an XCM can house two FPs via its two XMA).

Table 27 Show Distributed CPU Protection Policer Output Fields

Label	Description
Policer-State	The state of the policer with the following potential values:
	<i>Exceed</i> - The policer has been detected as non-conformant to the associated DCP policy parameters (e.g. packets exceeded the configured rate and the DCP polling process identified this occurrence)
	<i>Conform</i> - The policer has been detected as conformant to the associated DCP policy parameters (rate)
	<i>not-applicable</i> - Newly created policers or policers that are not currently instantiated. This includes policers configured on linecards that are not in service.
Protocols Mapped	A list of protocols that are configured to map to the particular policer.
Oper. xyz fields	The actual hardware may not be able to perfectly rate limit to the exact configured rate parameters in a DCP policy. In this case the configured rate parameters will be adapted to the closest supported rate. These adapted operational values are displayed in CLI when the “detail” keyword is included in the show command. The adapted Oper. parameters are only applicable if the policer is instantiated (e.g. if the associated forwarding plane is operational, or for an interface if there is a physical port configured for the interface, or if the dynamic policers are allocated), otherwise values of 0 kb/s, etc are displayed.
	<i>Oper. Kbps</i> - The adapted ‘kilobits-per-second’ value for DCP ‘kbps’ rates
	<i>Oper. MBS</i> - The adapted ‘mbs size’ value for DCP ‘kbps’ rates
	<i>Oper. Depth</i> - The calculated policer bucket depth in packets (for DCP ‘packets’ rates) or in bytes (for DCP ‘kbps’ rates)
	<i>Oper. Packets</i> - The adapted ‘ppi’ value for DCP ‘packets’ rates
	<i>Oper. Within</i> - The adapted ‘within seconds’ value for DCP ‘packets’ rates
	<i>Oper. Init. Delay</i> - The adapted ‘initial-delay packets’ value for DCP ‘packets’ rates

Table 27 Show Distributed CPU Protection Policer Output Fields

Label	Description
Exceed-Count	The count of packets exceeding the policing parameters since the given policer was previously declared as conformant or newly instantiated. This counter has the same behavior as the exceed counter in the DCP the log events – they are baselined (reset) when the policer transitions to conformant.
Detec. Time Remain	The remaining time in the detection-time countdown during which a policer in the exceed state is being monitored to see if it is once again conformant.
Hold-Down Remain	The remaining time in the hold-down countdown during which a policer is treating all packets as exceeding.
All Dyn-Plcr Alloc.	Indicates that all the dynamic enforcement policers have been allocated and instantiated for a given local-monitor.
Dyn-Policer Alloc.	Indicates that a dynamic policer has been instantiated.

Sample Output

```
*A:nodeA# show service id 33 sap 1/1/3:33 dist-cpu-protection detail
=====
Service Access Points(SAP) 1/1/3:33
=====
Distributed CPU Protection Policy : test1
-----
Statistics/Policer-State Information
=====
-----
Static Policer
-----
Policer-Name      : arp
Card/FP           : 1/1
Protocols Mapped  : arp
Exceed-Count      : 0
Detec. Time Remain : 0 seconds
Operational (adapted) rate parameters:
  Oper. Packets   : 5 ppi
  Oper. Initial Delay: 6 packets
  Oper. Depth     : 0 packets
Policer-State      : Conform

Policer-Name      : dhcp
Card/FP           : 1/1
Protocols Mapped  : dhcp
Exceed-Count      : 0
Detec. Time Remain : 0 seconds
Operational (adapted) rate parameters:
  Oper. Kbps      : 2343 kbps
  Oper. MBS       : 240 kilobytes
  Oper. Depth     : 0 bytes
Policer-State      : Conform

... (snip)
```

```

*A:nodaA# show service id 33 sap 1/1/3:34 dist-cpu-protection detail
=====
Service Access Points(SAP) 1/1/3:34
=====
Distributed CPU Protection Policy : test2
-----
Statistics/Policer-State Information
=====
Static Policer
-----
No entries found
-----
Local-Monitoring Policer
-----
Policer-Name       : my-local-mon1
Card/FP            : 1/1                Policer-State      : conform
Protocols Mapped   : arp, pppoe-pppoa
Exceed-Count       : 0
All Dyn-Plcr Alloc. : False
Operational (adapted) rate parameters:
  Oper. Packets    : 10 ppi              Oper. Within       : 8 seconds
  Oper. Initial Delay: 8 packets
  Oper. Depth      : 0 packets
-----
Dynamic-Policer (Protocol)
-----
Protocol (Dyn-Plcr) : arp
Card/FP             : 1/1                Protocol-State     : not-applicable
Exceed-Count        : 0
Detec. Time Remain  : 0 seconds          Hold-Down Remain. : none
Dyn-Policer Alloc. : False
Operational (adapted) rate parameters: unknown

Protocol (Dyn-Plcr) : pppoe-pppoa
Card/FP             : 1/1                Protocol-State     : not-applicable
Exceed-Count        : 0
Detec. Time Remain  : 0 seconds          Hold-Down Remain. : none
Dyn-Policer Alloc. : False
Operational (adapted) rate parameters: unknown
-----

```

dist-cpu-protection

Syntax	dist-cpu-protection [detail]
Context	show>router>interface
Description	This command displays Distributed CPU Protection parameters and status at the router Interface level.

Parameters **detail** — Specifies to include the adapted operational rate parameters in the CLI output. The adapted Oper. parameters are only applicable if the policer is instantiated (for example, if the associated forwarding plane is operational, or for an interface if there is a physical port configured for the interface, or if the dynamic policers are allocated), otherwise values of 0 kb/s, and so on, are displayed.

Output Distributed CPU Protection Policer Output

[Table 28](#) describes Distributed CPU Protection Policer Output output fields.

Table 28 Show Distributed CPU Protection Policer Output Fields

Label	Description
Distributed CPU Protection Policy	Displays the DCP policy assigned to the object.
Policer-Name	Displays the configured name of the static policer
Card/FP	Displays the card and FP identifier. FP identifies the instance of the FP (FastPath) chipset. Some cards have a single FP (for example, IOM3-XP) and some cards can contain multiple FPs (for example, an XCM can house two FPs via its two XMA).
Policer-State	Displays the state of the policer with the following potential values:
	<i>Exceed</i> - The policer has been detected as non-conformant to the associated DCP policy parameters (packets exceeded the configured rate and the DCP polling process identified this occurrence)
	<i>Conform</i> - The policer has been detected as conformant to the associated DCP policy parameters (rate)
Protocols Mapped	<i>not-applicable</i> - Newly created policers or policers that are not currently instantiated. This includes policers configured on linecards that are not in service.
	Displays a list of protocols that are configured to map to the particular policer.

Table 28 Show Distributed CPU Protection Policer Output Fields

Label	Description
Oper. xyz fields	<p>The actual hardware may not be able to perfectly rate limit to the exact configured rate parameters in a DCP policy. In this case the configured rate parameters will be adapted to the closest supported rate. These adapted operational values are displayed in CLI when the detail keyword is included in the show command. The adapted Oper. parameters are only applicable if the policer is instantiated (for example, if the associated forwarding plane is operational, or for an interface if there is a physical port configured for the interface, or if the dynamic policers are allocated), otherwise values of 0 kb/s, etc are displayed.</p> <p><i>Oper. Kbps</i> - Displays the adapted 'kilobits-per-second' value for DCP 'kbps' rates</p> <p><i>Oper. MBS</i> - Displays the adapted 'mbs size' value for DCP 'kbps' rates</p> <p><i>Oper. Depth</i> - Displays the calculated policer bucket depth in packets (for DCP 'packets' rates) or in bytes (for DCP 'kbps' rates)</p> <p><i>Oper. Packets</i> - Displays the adapted 'ppi' value for DCP 'packets' rates</p> <p><i>Oper. Within</i> - Displays the adapted 'within seconds' value for DCP 'packets' rates</p> <p><i>Oper. Init. Delay</i> - Displays the adapted 'initial-delay packets' value for DCP 'packets' rates</p>
Exceed-Count	Displays the count of packets exceeding the policing parameters since the given policer was previously declared as conformant or newly instantiated. This counter has the same behavior as the exceed counter in the DCP the log events – they are baselined (reset) when the policer transitions to conformant.
Detec. Time Remain	Displays the remaining time in the detection-time countdown during which a policer in the exceed state is being monitored to see if it is once again conformant.
Hold-Down Remain	Displays the remaining time in the hold-down countdown during which a policer is treating all packets as exceeding.
All Dyn-Plcr Alloc.	Indicates that all the dynamic enforcement policers have been allocated and instantiated for a given local-monitor.
Dyn-Policer Alloc.	Indicates that a dynamic policer has been instantiated.

Sample Output


```
*A:Dut-A# show router interface "test" dist-cpu-protection detail
=====
Interface "test" (Router: Base)
=====
Distributed CPU Protection Policy : dcpuPol
-----
Statistics/Policer-State Information
=====
Static Policer
-----
Policer-Name      : staticArpPolicer
Card/FP           : 4/1                Policer-State      : Exceed
Protocols Mapped  : arp
Exceed-Count      : 10275218
Detec. Time Remain : 29 seconds        Hold-Down Remain.  : none
Operational (adapted) Rate Parameters:
  Oper. Packets   : 100 ppi            Oper. Within       : 1 seconds
  Oper. Initial Delay: none
  Oper. Depth     : 100 packets
-----
Local-Monitoring Policer
-----
Policer-Name      : localMonitor
Card/FP           : 4/1                Policer-State      : Exceed
Protocols Mapped  : icmp, ospf
Exceed-Count      : 8019857
All Dyn-Plcr Alloc. : True
Operational (adapted) Rate Parameters:
  Oper. Packets   : 200 ppi            Oper. Within       : 1 seconds
  Oper. Initial Delay: none
  Oper. Depth     : 0 packets
-----
Dynamic-Policer (Protocol)
-----
Protocol(Dyn-Plcr) : icmp
Card/FP           : 4/1                Protocol-State     : Exceed
Exceed-Count      : 1948137
Detec. Time Remain : 29 seconds        Hold-Down Remain. : none
Dyn-Policer Alloc. : True
Operational (adapted) Rate Parameters:
  Oper. Kbps      : 25 kbps            Oper. MBS         : 256 bytes
  Oper. Depth     : 274 bytes
-----
Protocol(Dyn-Plcr) : ospf
Card/FP           : 4/1                Protocol-State     : Exceed
Exceed-Count      : 1487737
Detec. Time Remain : 29 seconds        Hold-Down Remain. : none
Dyn-Policer Alloc. : True
Operational (adapted) Rate Parameters:
  Oper. Kbps      : 25 kbps            Oper. MBS         : 256 bytes
  Oper. Depth     : 284 bytes
-----
=====
```

excessive-sources

- Syntax** **excessive-sources** [**service-id** *service-id* **sap-id** *sap-id*]
- Context** show>system>security>cpu-protection
- Description** This command displays sources exceeding their per-source rate limit.
- Parameters** *service-id* — Displays information for services exceeding their per-source rate limit.
sap-id — Displays information for SAPs exceeding their per-source rate limit.

policy

- Syntax** **policy** [*policy-id*] **association**
- Context** show>system>security>cpu-protection
show>system>security>dist-cpu-protection
- Description** This command displays CPU protection policy information.
- Parameters** *policy-id* — Displays CPU protection policy information for the specified policy ID>
association — This keyword displays policy-id associations.

protocol-protection

- Syntax** **protocol-protection**
- Context** show>system>security>cpu-protection
- Description** This command display all interfaces with non-zero drop counters.

violators

- Syntax** **violators** [**port**] [**interface**] [**sap**] [**video**] [**sdp**]
- Context** show>system>security>cpu-protection
- Description** This command displays all interfaces, ports or SAPs with CPU protection policy violators. It also includes objects (saps, interfaces) that exceed the out-profile-rate and have the log-events keyword enabled for the out-profile-rate in the cpu-protection policy associated with the object.
- Parameters** **port** — Displays violators associated with the port.
interface — Displays violators associated with the interface.
sap — Displays violators associated with the SAP.

video — Displays violators associated with the video entity.

sdp — Displays violators associated with the SDP.

Output The following is an output example of CPU protection violators.

Sample Output

```
*A:SecuritySR7>config>sys>security>cpu-protection>policy# show system security
cpu-protection violators
=====
Ports where a rate limit is violated
=====
Port-Id
  Type Limit First-Time          Last-Time          Violation-Periods
-----
No ports found
=====
Interfaces where the protection policy overall rate limit is violated
=====
Interface-Name          Router-Name
  Plcy Limit First-Time          Last-Time          Violation-Periods
-----
toIxia
  255 1000 10/02/2012 18:38:23 10/02/2012 18:39:31 70
-----
1 interface(s) found
=====
SAP's where the protection policy overall rate limit is violated
=====
SAP-Id          Service-Id
  Plcy Limit First-Time          Last-Time          Violation-Periods
-----
No SAP's found
=====
SDP's where the protection policy overall rate limit is violated
=====
SDP-Id          Service-Id
  Plcy Limit First-Time          Last-Time          Violation-Periods
-----
No SDP's found
=====
Video clients where the protection policy per-source rate limit is violated
=====
Client IP Address  Video-Interface          Service-Id
  Plcy Limit First-Time          Last-Time          Violation-Periods
-----
No clients found
=====
```

mac-filter

Syntax **mac-filter** [**entry** *entry-id*]

- Context** show>system>security>cpm-filter
- Description** This command displays CPM MAC filters.
- Parameters** *entry-id* — Displays information about the specified entry.
Values 1 to 2048

Output**Sample Output**

```
*B:bksim67# show system security cpm-filter mac-filter
=====
CPM Mac Filter (applied)
=====
Entry-Id  Dropped   Forwarded Description
-----
1          23002     47094
-----
Num CPM Mac filter entries: 1
=====
*B:bksim67#
```

mac-filter

- Syntax** **mac-filter** [**entry** *entry-id*]
- Context** show>system>security>management-access-filter
- Description** This command displays management access MAC filters.
- Parameters** *entry-id* — Displays information about the specified entry.
Values 1 to 9999

Output**Sample Output**

```
*B:bksim67# show system security management-access-filter mac-filter
=====
Mac Management Access Filter
=====
filter type      : mac
Def. Action      : permit
Admin Status     : enabled (no shutdown)
-----
Entry           : 1           Action           : deny
FrameType       : ethernet_II  Svc-Id           : Undefined
Src Mac         : Undefined
Dest Mac        : Undefined
Dot1p           : Undefined   Ethertype        : Disabled
DSAP            : Undefined   SSAP             : Undefined
-----
```

```
Snap-pid          : Undefined          ESnap-oui-zero    : Undefined
cfm-opcode        : Undefined
Log               : disabled          Matches           : 0
=====
*B:bksim67#
```

keychain

- Syntax** `keychain [key-chain] [detail]`
- Context** `show>system>security`
- Description** This command displays keychain information.
- Parameters** *key-chain* — Specifies the keychain name to display.
detail — Displays detailed keychain information.
- Output**

Sample Output

```
*A:ALA-A# show system security keychain test
=====
Key chain:test
=====
TCP-Option number send      : 254          Admin state   : Up
TCP-Option number receive   : 254          Oper state    : Up
=====
*A:ALA-A#
*A:ALA-A# show system security keychain test detail
=====
Key chain:test
=====
TCP-Option number send      : 254          Admin state   : Up
TCP-Option number receive   : 254          Oper state    : Up
=====
Key entries for key chain: test
=====
Id          : 0
Direction   : send-receive      Algorithm     : hmac-sha-1-96
Admin State : Up                  Valid         : Yes
Active      : Yes                Tolerance    : 300
Begin Time  : 2007/02/15 18:28:37 Begin Time (UTC) : 2007/02/15 17:28:37
End Time    : N/A                End Time (UTC)  : N/A
=====
Id          : 1
Direction   : send-receive      Algorithm     : aes-128-cmac-96
Admin State : Up                  Valid         : Yes
Active      : No                 Tolerance    : 300
Begin Time  : 2007/02/15 18:27:57 Begin Time (UTC) : 2007/02/15 17:27:57
End Time    : 2007/02/15 18:28:13 End Time (UTC)   : 2007/02/15 17:28:13
=====
Id          : 2
Direction   : send-receive      Algorithm     : aes-128-cmac-96
```

```

Admin State      : Up                Valid           : Yes
Active          : No                Tolerance      : 500
Begin Time      : 2007/02/15 18:28:13 Begin Time (UTC) : 2007/02/15 17:28:13
End Time        : 2007/02/15 18:28:37 End Time (UTC)   : 2007/02/15 17:28:37
=====
*A:ALA-A#
    
```

management-access-filter

- Syntax** **management-access-filter**
- Context** show>system>security
- Description** This command displays management access filter information for IP and MAC filters.

ip-filter

- Syntax** **ip-filter [entry entry-id]**
- Context** show>system>security>mgmt-access-filter
- Description** This command displays management-access IP filters.
- Parameters** *entry-id* — Displays information for the specified entry.
 Values 1 to 9999
- Output** Management Access Filter Output

[Table 29](#) describes management access filter output fields.

Table 29 Show Management Access Filter Output Fields

Label	Description
Def. action	Permit Specifies that packets not matching the configured selection criteria in any of the filter entries are permitted.
	Deny Specifies that packets not matching the configured selection criteria in any of the filter entries are denied and that a ICMP host unreachable message will be issued.
	Deny-host-unreachable Specifies that packets not matching the configured selection criteria in the filter entries are denied.
Entry	The entry ID in a policy or filter table.

Table 29 Show Management Access Filter Output Fields (Continued)

Label	Description
Description	A text string describing the filter.
Src IP	The source IP address used for management access filter match criteria.
Src interface	The interface name for the next hop to which the packet should be forwarded if it hits this filter entry.
Dest port	The destination port.
Matches	The number of times a management packet has matched this filter entry.
Protocol	The IP protocol to match.
Action	The action to take for packets that match this filter entry.

Sample Output

```
*A:Dut-F# show system security management-access-filter ip-filter
=====
IPv4 Management Access Filter
=====
filter type:      : ip
Def. Action      : permit
Admin Status     : enabled (no shutdown)
-----
Entry            : 1
Src IP           : 192.168.0.0/16
Src interface    : undefined
Dest port       : undefined
Protocol        : undefined
Router          : undefined
Action          : none
Log             : disabled
Matches         : 0
=====
*A:Dut-F#
```

ipv6-filter

- Syntax** **ipv6-filter [entry entry-id]**
- Context** show>system>security>mgmt-access-filter
- Description** This command displays management-access IPv6 filters and only applies to the 7750 SR and 7950 XRS.

Parameters *entry-id* — Specifies the IPv6 filter entry ID to display.

Values 1 to 9999

Output

Sample Output

```
*A:Dut-C# show system security management-access-filter ipv6-filter entry 1
=====
IPv6 Management Access Filter
=====
filter type      : ipv6
Def. Action      : permit
Admin Status     : enabled (no shutdown)
-----
Entry           : 1
Src IP          : 2001::1/128
Flow label      : undefined
Src interface    : undefined
Dest port       : undefined
Next-header     : undefined
Router          : undefined
Action          : permit
Log             : enabled
Matches         : 0
=====
*A:Dut-C# s
```

password-options

Syntax password-options

Context show>system>security

Description This command displays configured password options.

Output Password Options Output

[Table 30](#) describes password options output fields.

Table 30 Show Password Options Output Fields

Label	Description
Password aging in days	Displays the number of days a user password is valid before the user must change their password.
Time required between password changes	Displays the time interval between changed passwords.

Table 30 Show Password Options Output Fields (Continued)

Label	Description
Number of invalid attempts permitted per login	Displays the number of unsuccessful login attempts allowed for the specified time .
Time in minutes per login attempt	Displays the period of time, in minutes, that a specified number of unsuccessful attempts can be made before the user is locked out.
Lockout period (when threshold breached)	Displays the number of minutes that the user is locked out if the threshold of unsuccessful login attempts has been exceeded.
Authentication order	Displays the sequence in which password authentication is attempted among RADIUS, TACACS+, and local passwords.
User password history length	Displays the size of the password history file to be stored.
Accepted password length	Displays the minimum length required for local passwords.
Credits for each character type	Displays the credit for each character type. A credit is obtained for a particular character type; for example, uppercase, lowercase, numeric, or special character. Credits per character type are configurable. Credits can be used towards the minimum length of the password, so a trade-off can be made between a very long, simple password and a short, complex one.
Required character types	Displays the character types that are required in a password; for example, uppercase, lowercase, numeric, or special character.
Minimum number different character types	Displays the minimum number of each different character types in a password.
Required distance with previous password	Displays the minimum Levenshtein distance between a new password and the old password.
Allow consecutively repeating a character	Displays the number of times the same character is allowed to be repeated consecutively.
Allow passwords containing username	Displays whether the user name is allowed as part of the password.
Palindrome allowed	Displays whether palindromes are allowed as part of the password.

Sample Output

```
A:ALA-7# show system security password-options
=====
Password Options
=====
Password aging in days                : none
Time required between password changes : 0d 00:10:00

Number of invalid attempts permitted per login : 3
Time in minutes per login attempt            : 5
Lockout period (when threshold breached)     : 10
Authentication order                       : radius tacplus local
User password history length                : disabled
Accepted password length                   : 6..56 characters
Credits for each character type             : none
Required character types                    : none
Minimum number different character types    : 0
Required distance with previous password    : 5
Allow consecutively repeating a character   : always
Allow passwords containing username         : yes
Palindrome allowed                         : no
=====
A:ALA-7#
```

per-peer-queuing

- Syntax** **per-peer-queuing**
- Context** show>system>security
- Description** This command enables or disables CPMCFM hardware queuing per peer. TTL security only operates when per-peer-queuing is enabled.
- Output** Per-Peer-Queuing Output

[Table 31](#) describes per-peer-queuing output fields.

Table 31 Show Per-Peer-Queuing Output Fields

Label	Description
Per Peer Queuing	Displays the status (enabled or disabled) of CPM hardware queuing per peer.
Total Num of Queues	Displays the total number of hardware queues.
Num of Queues In Use	Displays the total number of hardware queues in use.

Sample Output

```
A:ALA-48# show system security per-peer-queuing
=====
CPM Hardware Queuing
```

```

=====
Per Peer Queuing      : Enabled
Total Num of Queues   : 8192
Num of Queues In Use  : 2
=====
A:ALA-48# configure

```

profile

- Syntax** `profile [user-profile-name]`
- Context** `show>system>security`
- Description** This command displays user profile information.
If the *profile-name* is not specified, then information for all profiles are displayed.
- Parameters** *user-profile-name* — Displays information for the specified user profile.
- Output** User Profile Output
[Table 32](#) describes user profile output fields.

Table 32 Show User Profile Output Fields

Label	Description
User Profile	Displays the profile name used to deny or permit user console access to a hierarchical branch or to specific commands.
Def. action	Permit all — Permits access to all commands. Deny — Denies access to all commands. None — No action is taken.
Entry	The entry ID in a policy or filter table.
Description	Displays the text string describing the entry.
Match Command	Displays the command or subtree commands in subordinate command levels.
Action	Permit all — Commands matching the entry command match criteria are permitted. Deny — Commands not matching the entry command match criteria are not permitted.
No. of profiles	The total number of profiles listed.

Sample Output

```
A:ALA-7# show system security profile administrative
```

```

=====
User Profile
=====
User Profile : administrative
Def. Action  : permit-all
-----
Entry       : 10
Description :
Match Command: configure system security
Action      : permit
-----
Entry       : 20
Description :
Match Command: show system security
Action      : permit
-----
No. of profiles:
=====
A:ALA-7#
    
```

source-address

- Syntax** **source-address**
- Context** show>system>security
- Description** This command displays source-address configured for applications.
- Output** Source Address Output

[Table 33](#) describes source address output fields.

Table 33 Show Source Address Output Fields

Label	Description
Application	Displays the source-address application.
IP address Interface Name	Displays the source address IP address or interface name.
Oper status	Up: The source address is operationally up. Down: The source address is operationally down.

Sample Output

```

A:SR-7# show system security source-address
=====
Source-Address applications
=====
Application          IP address/Interface Name          Oper status
-----
telnet                10.20.1.7                          Up
    
```

```
radius                loopback1                Up
=====
A:SR-7#
```

ssh

- Syntax** **ssh**
 - Context** show>system>security
 - Description** This command displays all the SSH sessions as well as the SSH status and fingerprint. The type of SSH application (CLI, SCP, SFTP or NETCONF) is indicated for each SSH connection.
 - Output** SSH Options Output
- [Table 34](#) describes SSH output fields .

Table 34 Show System Security SSH Options Output Fields

Label	Description
SSH status	SSH is enabled: Displays that SSH server is enabled. SSH is disabled: Displays that SSH server is disabled.
SSH Preserve Key	Enabled: Displays that preserve-key is enabled. Disabled: Displays that preserve-key is disabled.
SSH protocol version 1	Enabled: Displays that SSH1 is enabled. Disabled: Displays that SSH1 is disabled.
SSH protocol version 2	Enabled: Displays that SSH2 is enabled. Disabled: Displays that SSH2 is disabled.
Key fingerprint	The key fingerprint is the server's identity. Clients trying to connect to the server verify the server's fingerprint. If the server fingerprint is not known, the client may not continue with the SSH session since the server might be spoofed.
Connection	The IP address of the connected router(s) (remote client).
Encryption	des: Data encryption using a private (secret) key. 3des: An encryption method that allows proprietary information to be transmitted over untrusted networks.
Username	The name of the user.
Version	The SSH version number.
Server Name	The type of SSH application (CLI, SCP, SFTP or NETCONF)

Table 34 Show System Security SSH Options Output Fields (Continued)

Label	Description (Continued)
Number of SSH sessions	The total number of SSH sessions.

Sample output

```
*A:ALA-49# show system security ssh
=====
SSH Server
=====
Administrative State      : Enabled
Operational State       : Up
Preserve Key             : Enabled

SSH Protocol Version 1   : Disabled

SSH Protocol Version 2   : Enabled
DSA Host Key Fingerprint : 88:41:1c:7e:97:64:df:a0:e4:54:c2:cc:3d:dd:c7:70
RSA Host Key Fingerprint : 63:b8:c4:8a:17:b7:1c:95:35:91:c9:08:75:cc:31:a3
-----
Connection      Username      Version  ServerName  Status
-----
138.120.214.254  admin        2         netconf     connected
138.120.140.148  admin        2         cli         connected
-----
Number of SSH sessions : 2
=====
```

user

- Syntax** **user** [*user-id*] [**detail**]
 user [*user-id*] **lockout**
- Context** show>system>security
- Description** This command displays user registration information.

If no command line options are specified, summary information for all users displays.
- Parameters** *user-id* — Displays information for the specified user.

 Default All users

 detail — Displays detailed user information to the summary output.

 lockout — Displays information about any users who are currently locked out.
- Output** User Output

 [Table 35](#) describes user output fields.

Table 35 Show System Security User Output Fields

Label	Description
User ID	The name of a system user
Users	
New Pwd	y — The user must change their password at the next login n — The user does not need to change their password at the next login
User Permissions	console y — The user is authorized for console access n — The user is not authorized for console access
	ftp y — The user is authorized for FTP access n — The user is not authorized for FTP access
	li y — The user is authorized for LI access n — The user is not authorized for LI access
	snmp y — The user is authorized for SNMP access n — The user is not authorized for SNMP access
	netconf y — The user is authorized for NETCONF access n — The user is not authorized for NETCONF access
	grpc y — The user is authorized for gRPC access n — The user is not authorized for gRPC access
Password Expires	The number of days after which the user must change their password
Login Attempt	The number of times that the user has attempted to log in, irrespective of whether the login succeeded or failed
Failed Logins	The number of unsuccessful login attempts
Local Conf	y — Password authentication is based on the local password database n — Password authentication is not based on the local password database
Number of users	The total number of listed users

Table 35 Show System Security User Output Fields (Continued)

Label	Description (Continued)
User Configuration Detail	
new pw required	yes — The user must change their password at the next login no — The user does not need to change their password at the next login
cannot change pw	yes — The user does not have the ability to change their password no — The user has the ability to change their password
home directory	The local home directory for the user for both console and FTP access
restricted to home	yes — The user is not allowed to navigate to a directory higher in the directory tree on the home directory device no — The user is allowed to navigate to a directory higher in the directory tree on the home directory device
login exec file	The user's login exec file which executes whenever the user successfully logs in to a console session
profile	The security profiles associated with the user
locked-out	Whether the user is currently locked out, and, if they are locked out, how much time remains before the user can attempt to log into the node again
Currently Failed Login Attempts	
Remaining Login Attempts	The number of login attempts remaining before the user is locked out
Remaining Lockout Time (min:sec)	The number of minutes and seconds remaining until the lockout expires and the user can attempt to log in again

The following are output examples for the 7450 ESS and 7750 SR:

```

show system security user
=====
Users
=====
user id      need   user permissions      password attempted failed local
            new pwd console ftp snmp      expires  logins  logins  conf
-----
admin       n     y     n   n                never    21     0     y
=====

show system security user detail
=====
Users
=====

```



```
=====
user id      need      user permissions  password  attempted  failed  local
             new pwd  console ftp snmp  expires   logins   logins   conf
-----
admin        n         y         n  n         never     21      0        y
=====
```

User Configuration Detail

```
=====
user id      : admin
-----
```

console parameters

```
-----
new pw required :                no cannot change pw : no
home directory  : cf3:\
restricted to home : no
login exec file :
profile         : administrative
-----
```

snmp parameters

```
=====
show system security user detail
=====
```

Users

```
=====
User ID      New User Permissions          Password Login  Failed Local
             Pwd console ftp li snmp netconf grpc Expires Attempt Logins Conf
-----
admin        n  y         y  n  y  y         n  never  9      0        y
-----
```

Number of users : 1

```
=====
User Configuration Detail
=====
```

```
=====
user id      : admin
-----
```

console parameters

```
-----
new pw required : no                cannot change pw : no
home directory  :
restricted to home : no
login exec file :
profile         : default
profile         : administrative
locked-out      : no
-----
```

snmp parameters

```
=====
show system security user lockout
=====
```

Currently Failed Login Attempts

```
=====
User ID Remaining Login attempts Remaining Lockout Time (min:sec)
-----
```

jason123 N/A 9:56

```
-----
Number of users : 1
-----
```

With the introduction of the PKI on an SR (SSH Server) the authentication process can be done via PKI or password. SSH client usually authenticate via PKI and password if PKI is configured on the client. In this case PKI takes precedence over password in most clients.

All client authentications are logged and display in the **show>system>security>user detail**. [Table 36](#) shows the rules where pass and fail attempts are logged.

Table 36 Pass/Fail Login Attempts

Authentication Order	Client (i.e., putty)	Server (i.e., SR)		CLI Show System Security Attempts (SR)	
	Private Key Programmed	Public Key Configured	Password Configured	Logins Attempts	Failed Logins
1. Public Key	Yes	Yes	N/A	Increment	
2. Password	Yes	Yes (No match between client and server. Go to password.)	Yes	Increment	
	Yes	No	Yes	Increment	
	No	N/A	Yes	Increment	
	No	N/A	No		Increment
1. Public Key (only)	Yes	Yes	N/A	Increment	
	Yes	Yes (No match between client and server. Go go password.)			Increment
	Yes		N/A		Increment
	No		N/A		Increment

view

Syntax **view** [*view-name*] [**detail**]

Context show>system>security

Description This command displays the SNMP MIB views.

Parameters *view-name* — Specify the name of the view to display output. If no view name is specified, the complete list of views displays.

detail — Displays detailed view information.

Output View Output

[Table 37](#) describes show view output fields.

Table 37 Show View Output Fields

Label	Description
view name	The name of the view. Views control the accessibility of a MIB object within the configured MIB view and subtree.
oid tree	The object identifier of the ASN.1 subtree.
mask	The bit mask that defines a family of view subtrees.
permission	Indicates whether each view is included or excluded
No. of Views	Displays the total number of views.

Sample Output

```
A:ALA-48# show system security view
=====
Views
=====
view name      oid tree      mask          permission
-----
iso            1             1             included
read1         1.1.1.1      11111111     included
write1        2.2.2.2      11111111     included
testview      1             11111111     included
testview      1.3.6.1.2    11111111     excluded
mgmt-view     1.3.6.1.2.1.2  included
mgmt-view     1.3.6.1.2.1.4  included
mgmt-view     1.3.6.1.2.1.5  included
mgmt-view     1.3.6.1.2.1.6  included
mgmt-view     1.3.6.1.2.1.7  included
mgmt-view     1.3.6.1.2.1.31 included
mgmt-view     1.3.6.1.2.1.77 included
mgmt-view     1.3.6.1.4.1.6527.3.1.2.3.7 included
mgmt-view     1.3.6.1.4.1.6527.3.1.2.3.11 included
vprn-view     1.3.6.1.2.1.2  included
vprn-view     1.3.6.1.2.1.4  included
vprn-view     1.3.6.1.2.1.5  included
vprn-view     1.3.6.1.2.1.6  included
vprn-view     1.3.6.1.2.1.7  included
vprn-view     1.3.6.1.2.1.15 included
vprn-view     1.3.6.1.2.1.23 included
vprn-view     1.3.6.1.2.1.31 included
vprn-view     1.3.6.1.2.1.68 included
vprn-view     1.3.6.1.2.1.77 included
```

vprn-view	1.3.6.1.4.1.6527.3.1.2.3.7		included
vprn-view	1.3.6.1.4.1.6527.3.1.2.3.11		included
vprn-view	1.3.6.1.4.1.6527.3.1.2.20.1		included
no-security	1		included
no-security	1.3.6.1.6.3		excluded
no-security	1.3.6.1.6.3.10.2.1		included
no-security	1.3.6.1.6.3.11.2.1		included
no-security	1.3.6.1.6.3.15.1.1		included
on-security	2	00000000	included

No. of Views: 33

=====

A:ALA-48#

certificate

- Syntax** **certificate**
- Context** show
- Description** This command displays certificate information.

ca-profile

- Syntax** **ca-profile**
ca-profile *name* [**association**]
- Context** show>certificate
- Description** This command shows certificate-authority profile information.
- Parameters** *name* — Specifies the name of the Certificate Authority (CA) profile.
association — Displays associated CA profiles.

ocsp-cache

- Syntax** **ocsp-cache** [*entry-id*]
- Context** show>certificate
- Description** This command displays the current cached OCSP results. The output includes the following information:

Certificate issuer

Certificate serial number

OCSP result

Cache entry expire time

Parameters *entry-id* — Specifies the local cache entry identifier of the certificate that was validated by the OCSP responder.

statistics

Syntax **statistics**

Context show>certificate

Description This command shows certificate related statistics.

2.20.2.1.2 Login Control

users

Syntax **users**

Context show

Description Displays console user login and connection information.

Output Users Output

[Table 38](#) describes show users output fields.

Table 38 Show Users Output Fields

Label	Description
User	The user name.
Type	The user is authorized this access type.
From	The originating IP address.
Login time	The time the user logged in.
Idle time	The amount of idle time for a specific login.
Number of users	Displays the total number of users logged in.

Sample Console Users Output

```
A:ALA-7# show users
=====
```

```

User           Type      From           Login time           Idle time
=====
testuser       Console   --            21FEB2007 04:58:55  0d 00:00:00  A
=====
Number of users : 1
'A' indicates user is in admin mode
=====
A:ALA-7#

```

2.20.2.2 Clear Commands

statistics

- Syntax** `statistics [interface ip-int-name | ip-address]`
- Context** `clear>router>authentication`
- Description** This command clears authentication statistics.
- Parameters** *ip-int-name* — Clears the authentication statistics for the specified interface name. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes
- ip-address* — Clears the authentication statistics for the specified IP address.

ip-filter

- Syntax** `ip-filter [entry entry-id]`
- Context** `clear>cpm-filter`
- Description** This command clears IP filter statistics.
- Parameters** *entry-id* — Specifies a particular CPM IP filter entry.
- Values** 1 to 2048

ipv6-filter

- Syntax** `ipv6-filter [entry entry-id]`
- Context** `clear>cpm-filter`
- Description** This command clears IPv6 filter statistics.

Parameters *entry-id* — Specifies a particular CPM IP filter entry.
Values 1 to 2048

mac-filter

Syntax **mac-filter** [**entry** *entry-id*]
Context clear>cpm-filter
Description This command clears MAC filter statistics.
Parameters *entry-id* — Specifies a particular CPM MAC filter entry.
Values 1 to 2048

ipv6-filter

Syntax **ipv6-filter** [**entry** *entry-id*]
Context clear>cpm-filter
Description This command clears IPv6 filter information and only applies to the 7750 SR and 7950 XRS.
Parameters *entry-id* — Specifies a particular CPM IPv6 filter entry.
Values 1 to 2048

2.20.2.2.1 CPU Protection Commands

cpu-protection

Syntax **cpu-protection**
Context clear
Description This command enables the context to clear CPU protection data.

excessive-sources

Syntax **excessive-sources**
Context clear>cpu-protection
Description This command clears the records of sources exceeding their per-source rate limit.

protocol-protection

Syntax	protocol-protection
Context	clear>cpu-protection
Description	This command clears the interface counts of packets dropped by protocol protection.

violators

Syntax	violators [port][interface][sap]
Context	clear>cpu-protection
Description	This command clears the rate limit violator record.
Parameters	port — Clears entries for ports. interface — Clears entries for interfaces. sap — Clears entries for SAPs.

cpm-queue

Syntax	cpm-queue <i>queue-id</i>
Context	clear
Description	This command clears CPM queue information.
Parameters	<i>queue-id</i> — Specifies the CPM queue ID. Values 33 to 2000

radius-proxy-server

Syntax	radius-proxy-server <i>server-name</i> statistics
Context	clear>router
Description	This command clears RADIUS proxy server data.
Parameters	<i>server-name</i> — Specifies the proxy server name. statistics — Clears statistics for the specified server.

2.20.2.3 Debug Commands

radius

Syntax	radius [detail] [hex] no radius
Context	debug
Description	This command enables debugging for RADIUS connections. The no form of the command disables the debug output.
Parameters	detail — Displays detailed output. hex — Displays the packet dump in hex format.

OCSP

Syntax	[no] ocsp
Context	debug
Description	This command enables debug output of OCSP protocol for the CA profile. The no form of the command disables the debug output.

ca-profile

Syntax	[no] ca-profile <i>profile-name</i>
Context	debug>ocsp
Description	This command enables debug output of a specific CA profile. The no form of the command disables the debug output.

2.20.2.4 Tools Commands

dist-cpu-protection

Syntax	dist-cpu-protection
Context	tools>perform>security tools>dump>security
Description	This command displays to release Distributed CPU Protection parameters and status at the per card and forwarding plane level.

release-hold-down

Syntax	release-hold-down interface <i>interface-name</i> [protocol <i>protocol</i>] [static-policer <i>name</i>] release-hold-down sap <i>sap-id</i> [protocol <i>protocol</i>] [static-policer <i>name</i>]
Context	tools>perform>security>dist-cpu-protection
Description	This command is used to release a Distributed CPU Protection (DCP) policer from a hold-down countdown (or indefinite hold-down if configured as such).
Parameters	interface <i>interface-name</i> — Specifies Router interface name. sap <i>sap-id</i> — Specifies sap identifier. protocol <i>protocol</i> — Specifies DCP protocol name (for example, arp, dhcp) static-policer <i>name</i> — Specifies DCP static policer name as defined in the DCP policy.

violators

Syntax	violators enforcement { sap interface } card <i>slot-number</i> [fp <i>fp-number</i>] violators local-monitor { sap interface } card <i>slot-number</i> [fp <i>fp-number</i>]
Context	tools>dump>security>dist-cput protection
Description	This command shows the non-conformant enforcement policers and local monitors.
Parameters	sap — -Indicates to display the violators associated with SAPs interface — - Indicates to display the violators associated with router interfaces. enforcement — Shows exceed and hold-down for Static and Dynamic Policers. local-monitor — Shows state of dynamic policer allocation for Local Monitoring Policers. card <i>slot-number</i> — The physical slot number for the card. Values 1 to n (n is platform dependent)

fp *fp-number* — Identifies the instance of the FP (FastPath) chipset. Some cards have a single FP (for example, an IOM3-XP) and some cards can contain multiple FPs (for example, an XCM can house two FPs via its two XMAAs).

Values 1 to 2

Output Users Output

Table 39 describes show users output fields.

Table 39 Output Parameters

Label	Description
Interface	The name of the router interface
Policer/Protocol	The configured name of the static policer (indicated with an [S]) or the DCP protocol name for a dynamic policer (indicated with a [D]).
[S] / [D]	indicates a static vs dynamic policer
Hld Rem	The remaining time in the hold-down countdown during which a policer is treating all packets as exceeding.

Sample Output

```
*A:Dut-A# tools dump security dist-cpu-protection violators enforcement interface
card 4 fp 1
=====
Distributed Cpu Protection Current Interface Enforcer Policer Violators
=====
Interface                Policer/Protocol          Hld Rem
-----
Violators on Slot-4 Fp-1
-----
test                      staticArpPolicer         [S] none
test                      icmp                     [D] none
test                      ospf                     [D] none
-----
[S]-Static [D]-Dynamic [M]-Monitor
=====
```

2.20.2.5 Admin Commands

clear lockout

- Syntax** `clear lockout {user name | all}`
- Context** admin>user
- Description** This command is used to clear any lockouts for a specific user, or for all users.
- Parameters** *name* — Specifies locked username.

clear password-history

- Syntax** `clear password-history {user name | all}`
- Context** admin>user
- Description** This command is used to clear old passwords used by a specific user, or for all users.
- Parameters** *name* — Specifies username.

3 SNMP

3.1 In This Chapter

This chapter provides information to configure SNMP.

Topics in this chapter include:

- [SNMP Overview](#)
 - [SNMP Architecture](#)
 - [Management Information Base](#)
 - [SNMP Protocol Operations](#)
 - [SNMP Versions](#)
 - [Management Information Access Control](#)
 - [User-Based Security Model Community Strings](#)
 - [Views](#)
 - [Access Groups](#)
 - [Users](#)
 - [Per-VPRN Logs and SNMP Access](#)
 - [Per-SNMP Community Source IP Address Validation](#)
- [Which SNMP Version to Use?](#)
- [Configuration Notes](#)

3.2 SNMP Overview

This section provides an overview of the Simple Network Management Protocol (SNMP).

3.2.1 SNMP Architecture

The Service Assurance Manager (SAM) is comprised of two elements: managers and agents. The manager is the entity through which network management tasks are facilitated. Agents interface managed objects. Managed devices, such as bridges, hubs, routers, and network servers can contain managed objects. A managed object can be a configuration attribute, performance statistic, or control action that is directly related to the operation of a device.

Managed devices collect and store management information and use Simple Network Management Protocol (SNMP). SNMP is an application-layer protocol that provides a message format to facilitate communication between SNMP managers and agents. SNMP provides a standard framework to monitor and manage devices in a network from a central location.

An SNMP manager controls and monitors the activities of network hosts which use SNMP. An SNMP manager can obtain (get) a value from an SNMP agent or store (set) a value in the agent. The manager uses definitions in the management information base (MIB) to perform operations on the managed device such as retrieving values from variables or blocks of data, replying to requests, and processing traps.

Between the SNMP agent and the SNMP manager the following actions can occur:

- The manager can get information from the agent.
- The manager can set the value of a MIB object that is controlled by an agent.
- The agent can send traps to notify the manager of significant events that occur on the router.

3.2.2 Management Information Base

A MIB is a formal specifications document with definitions of management information used to remotely monitor, configure, and control a managed device or network system. The agent's management information consists of a set of network objects that can be managed with SNMP. Object identifiers are unique object names that are organized in a hierarchical tree structure. The main branches are defined by the Internet Engineering Task Force (IETF). When requested, the Internet Assigned Numbers Authority (IANA) assigns a unique branch for use by a private organization or company. The branch assigned to Nokia (TiMetra) is 1.3.6.1.4.1.6527.

The SNMP agent provides management information to support a collection of IETF specified MIBs and a number of MIBs defined to manage device parameters and network data unique to Nokia's router.

3.2.3 SNMP Protocol Operations

Between the SNMP agent and the SNMP manager the following actions can occur:

- The manager can get information from the agent.
- The manager can set the value of a MIB object that is controlled by an agent.
- The agent notifies the manager of significant events that occur on the router.

3.2.4 SNMP Versions

The agent supports multiple versions of the SNMP protocol.

- SNMP Version 1 (SNMPv1) is the original Internet-standard network management framework.
SNMPv1 uses a community string match for authentication.
- The OS implementation uses SNMPv2c, the community-based administrative framework for SNMPv2. SNMPv2c uses a community string match for authentication.
- In SNMP Version 3 (SNMPv3), USM defines the user authentication and encryption features. View Access Control MIB (VACM) defines the user access control features. The SNMP-COMMUNITY-MIB is used to associate SNMPv1/ SNMPv2c community strings with SNMPv3 VACM access control.
SNMPv3 uses a username match for authentication.

3.2.5 Management Information Access Control

By default, the OS implementation of SNMP uses SNMPv3. SNMPv3 incorporates security model and security level features. A security model is the authentication type for the group and the security level is the permitted level of security within a security model. The combination of the security level and security model determines which security mechanism handles an SNMP packet.

To implement SNMPv1 and SNMPv2c configurations, several access groups are predefined. These access groups provide standard read-only, read-write, and read-write-all access groups and views that can simply be assigned community strings. In order to implement SNMP with security features, security models, security levels, and USM communities must be explicitly configured. Optionally, additional views which specify more specific OIDs (MIB objects in the subtree) can be configured.

Access to the management information in as SNMPv1/SNMPv2c agent is controlled by the inclusion of a community name string in the SNMP request. The community defines the sub-set of the agent's managed objects can be accessed by the requester. It also defines what type of access is allowed: read-only or read-write.

The use of community strings provide minimal security and context checking for both agents and managers that receive requests and initiate trap operations. A community string is a text string that acts like a password to permit access to the agent on the router.

Nokia's implementation of SNMP has defined three levels of community-named access:

- Read-Only permission — Grants only read access to objects in the MIB, except security objects.
- Read-Write permission — Grants read and write access to all objects in the MIB, except security objects.
- Read-Write-All permission — Grants read and write access to all objects in the MIB, including security objects.

3.2.6 User-Based Security Model Community Strings

User-based security model (USM) community strings associates a community string with an SNMPv3 access group and its view. The access granted with a community string is restricted to the scope of the configured group.

3.2.7 Views

Views control the access to a managed object. The total MIB of a router can be viewed as a hierarchical tree. When a view is created, either the entire tree or a portion of the tree can be specified and made available to a user to manage the objects contained in the subtree. Object identifiers (OIDs) uniquely identify managed objects. A view defines the type of operations for the view such as read, write, or notify.

OIDs are organized in a hierarchical tree with specific values assigned to different organizations. A view defines a subset of the agent's managed objects controlled by the access rules associated with that view.

The following system-provisioned views are available through the **config>system>security>snmp# view** context, which are particularly useful when configuring SNMPv1 and SNMPv2c:

- “iso” view—intended for administrative-type access to the entire supported object tree (except Lawful Interception)
- “no-security” view—similar to “iso” view, but removes access to several security areas of the object tree (such as SNMP communities, user and profile configuration, SNMP engine ID, etc). The “no-security” view is generally recommended over the “iso” view to reduce access to security objects.
- “li-view” view—provides access to a small set of Lawful Interception related objects
- “mgmt-view” view—provides access to IF-MIB and a few other basics
- “vprn-view” view—used to limit access to objects associated with a specific VPRN (for example, the Per-VPRN Logs and SNMP Access feature)

The Nokia SNMP agent associates SNMPv1 and SNMPv2c community strings with a SNMPv3 view.

3.2.8 Access Groups

Access groups associate a user group and a security model to the views the group can access. An access group is defined by a unique combination of a group name, security model (SNMPv1, SNMPv2c, or SNMPv3), and security level (no-authorization-no-privacy, authorization-no-privacy, or privacy).

An access group, in essence, is a template which defines a combination of access privileges and views. A group can be associated to one or more network users to control their access privileges and views.

When configuring access groups, the “no-security” view is generally recommended over the “iso” view in order to restrict access to security objects.

A set of system-provisioned access groups and system-created communities are available in SR OS. The system-provisioned groups and communities that begin with “cli-” are only used for internal CLI management purposes and are not exposed to external SNMP access.

Additional access parameters must be explicitly configured if the preconfigured access groups and views for SNMPv1 and SNMPv2c do not meet your security requirements.

3.2.9 Users

By default, authentication and encryption parameters are not configured. Authentication parameters which a user must use in order to be validated by the router can be modified. SNMP authentication allows the device to validate the managing node that issued the SNMP message and determine if the message has been tampered with.

User access and authentication privileges must be explicitly configured. In a user configuration, a user is associated with an access group, which is a collection of users who have common access privileges and views (see [Access Groups](#)).

3.2.10 Per-VPRN Logs and SNMP Access

Configuration of VPRN-specific logs (with VPRN-specific syslog destinations, SNMP trap/notification groups, etc) is supported in addition to the global logs configured under “config log”. The event streams for vprn logs contain only events that are associated with the particular vprn.

Each VPRN service can be configured with a set of SNMP v1/v2c community strings. These communities are mapped to the default “snmp-vprn” and “snmp-vprn-ro” views, which limit SNMP access to objects associated with a specific VPRN. For example, walking the ifTable (IF-MIB) using the community configured for VPRN 5 will return counters and status for VPRN 5. See the “vprn <x> snmp community” command description for more details.

3.2.11 Per-SNMP Community Source IP Address Validation

SNMPv1 and SNMPv2c requests can be validated against per-snmp-community whitelists (**src-access-list**) of configured source IPv4 and IPv6 addresses. Source IP address lists can be configured and then associated with an SNMP community.

SNMPv1 and SNMPv2c requests that fail the source IP address and community validation checks are discarded and are logged as SNMP event 2003 authenticationFailure (suppressed by default under “event-control”).

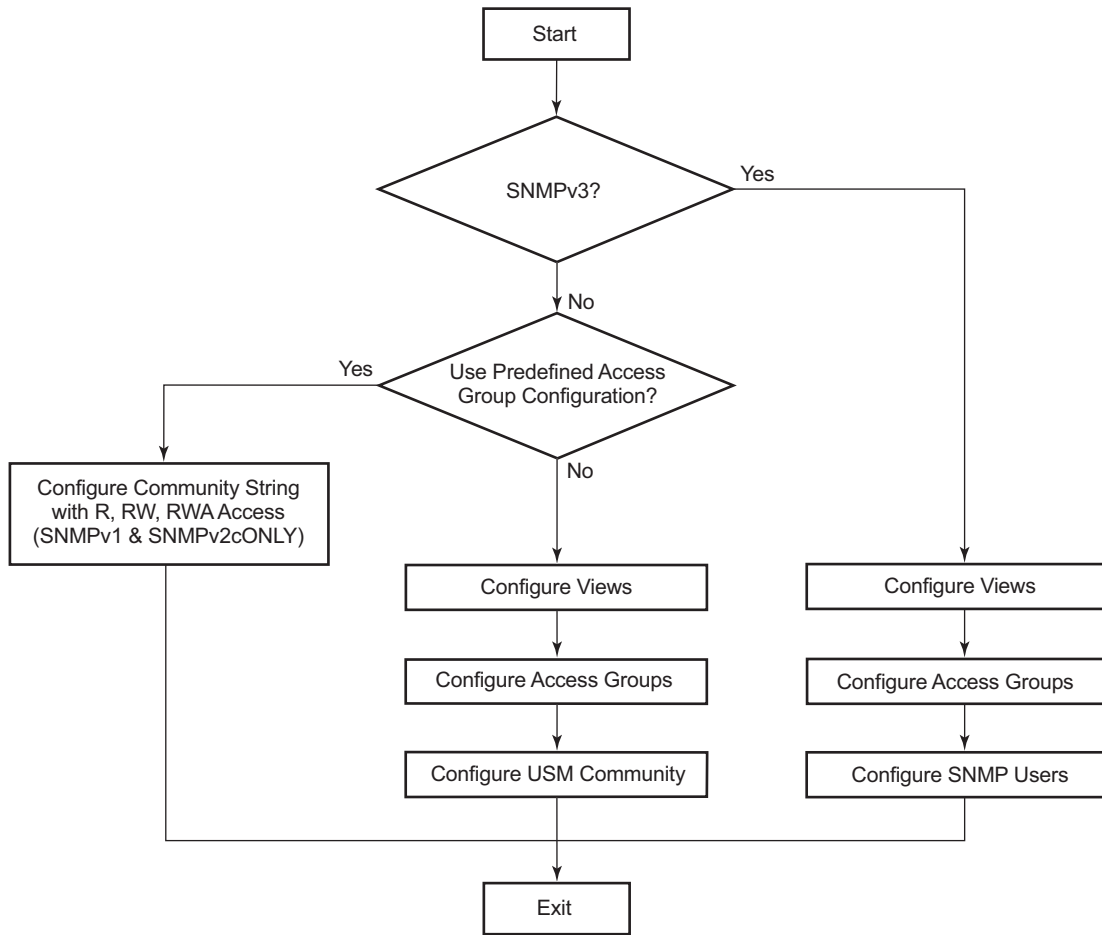
3.3 Which SNMP Version to Use?

SNMPv1 and SNMPv2c do not provide security, authentication, or encryption. Without authentication, a non authorized user could perform SNMP network management functions and eavesdrop on management information as it passes from system to system. Many SNMPv1 and SNMPv2c implementations are restricted read-only access, which, in turn, reduces the effectiveness of a network monitor in which network control applications cannot be supported.

To implement SNMPv3, an authentication and encryption method must be assigned to a user in order to be validated by the router. SNMP authentication allows the router to validate the managing node that issued the SNMP message and determine if the message was tampered with.

[Figure 14](#) depicts the configuration requirements to implement SNMPv1/SNMPv2c, and SNMPv3.

Figure 14 SNMPv1 and SNMPv2c Configuration and Implementation Flow



al_0203

3.4 Configuration Notes

This section describes SNMP configuration caveats.

3.4.1 General

- To avoid management systems attempting to manage a partially booted system, SNMP will remain in a shut down state if the configuration file fails to complete during system startup. While shutdown, SNMP gets and sets are not processed. However, notifications are issued if an SNMP trap group has been configured.

In order to enable SNMP, the portions of the configuration that failed to load must be initialized properly. Start SNMP with the **config>system>snmp>no shutdown** CLI command.

- Use caution when changing the SNMP engine ID. If the SNMP engine ID is changed in the **config>system>snmp> engineID** *engine-id* context, the current configuration must be saved and a reboot must be executed. If not, the previously configured SNMP communities and logger trap-target notify communities will not be valid for the new engine ID.

3.5 Configuring SNMP with CLI

This section provides information about configuring SNMP with CLI.

Topics in this section include:

- [SNMP Configuration Overview](#)
- [Basic SNMP Security Configuration](#)
- [Configuring SNMP Components](#)

3.6 SNMP Configuration Overview

This section describes how to configure SNMP components which apply to SNMPv1 and SNMPv2c, and SNMPv3 on the router.

- [Configuring SNMPv1 and SNMPv2c](#)
- [Configuring SNMPv3](#)

3.6.1 Configuring SNMPv1 and SNMPv2c

Nokia routers are based on SNMPv3. To use the routers with SNMPv1 and/or SNMPv2c, SNMP community strings must be configured. Three pre-defined access methods are available when SNMPv1 or SNMPv2c access is required. Each access method (**r**, **rw**, or **rwa**) is associated with an SNMPv3 access group that determines the access privileges and the scope of managed objects available. The **community** command is used to associate a community string with a specific access method and the required SNMP version (SNMPv1 or SNMPv2c). The access methods are:

- Read-Only — Grants read only access to the entire management structure with the exception of the security area.
- Read-Write — Grants read and write access to the entire management structure with the exception of the security area.
- Read-Write-All — Grants read and write access to the entire management structure, including security.

If the predefined access groups do not meet your access requirements, then additional access groups and views can be configured. The **usm-community** command is used to associate an access group with an SNMPv1 or SNMPv2c community string.

SNMP trap destinations are configured in the **config>log>snmp-trap-group** context.

3.6.2 Configuring SNMPv3

The OS implements SNMPv3. If security features other than the default views are required, then the following parameters must be configured:

- Configure views
- Configure access groups
- Configure SNMP users

3.7 Basic SNMP Security Configuration

This section provides information to configure SNMP parameters and provides examples of common configuration tasks. The minimal SNMP parameters are:

For SNMPv1 and SNMPv2c:

- Configure community string parameters.

For SNMPv3:

- Configure view parameters
- Configure SNMP group
- Configure access parameters
- Configure user with SNMP parameters

The following displays SNMP default views, access groups, and attempts parameters.

```
A:ALA-1>config>system>security>snmp# info detail
-----
view iso subtree 1
    mask ff type included
exit
view no-security subtree 1
    mask ff type included
exit
view no-security subtree 1.3.6.1.6.3
    mask ff type excluded
exit
view no-security subtree 1.3.6.1.6.3.10.2.1
```



```

        mask ff type included
    exit
    view no-security subtree 1.3.6.1.6.3.11.2.1
        mask ff type included
    exit
    view no-security subtree 1.3.6.1.6.3.15.1.1
        mask ff type included
    exit
    access group snmp-ro security-model snmpv1 security-level no-auth-
no-
privacy read no-security notify no-security
        access group snmp-ro security-model snmpv2c security-level no-auth-
no-
privacy read no-security notify no-security
        access group snmp-rw security-model snmpv1 security-level no-auth-
no-
privacy read no-security write no-security notify no-security
        access group snmp-rw security-model snmpv2c security-level no-auth-
no-
privacy read no-security write no-security notify no-security
        access group snmp-rwa security-model snmpv1 security-level no-auth-
no-
privacy read iso write iso notify iso
        access group snmp-rwa security-model snmpv2c security-level no-auth-
no-
privacy read iso write iso notify iso
        access group snmp-trap security-model snmpv1 security-level no-auth-
no-
privacy notify iso
        access group snmp-trap security-model snmpv2c security-level no-
auth-
no-privacy notify iso
        attempts 20 time 5 lockout 10

```

3.8 Configuring SNMP Components

Use the CLI syntax displayed below to configure the following SNMP scenarios:

- [Configuring a Community String](#)
- [Configuring View Options](#)
- [Configuring Access Options](#)
- [Configuring USM Community Options](#)
- [Configuring Other SNMP Parameters](#)

CLI Syntax:

```

config>system>security>snmp
attempts [count] [time minutes1] [lockout minutes2]
community community-string access-permissions [version
SNMP version]
usm-community community-string group group-name
view view-name subtree oid-value

```

```

        mask mask-value [type {included|excluded}]
access group group-name security-model security-model
security-level security-level [context context-name
[prefix-match]] [read view-name-1] [write view-name-2]
[notify view-name-3]

```

3.8.1 Configuring a Community String

SNMPv1 and SNMPv2c community strings are used to define the relationship between an SNMP manager and agent. The community string acts like a password to permit access to the agent. The access granted with a community string is restricted to the scope of the configured group.

One or more of these characteristics associated with the string can be specified:

- Read-only, read-write, and read-write-all permission for the MIB objects accessible to the community.
- The SNMP version, SNMPv1 or SNMPv2c.

Default access features are pre-configured by the agent for SNMPv1/SNMPv2c.

Use the following CLI syntax to configure community options:

```

config>system>security>snmp
community community-string access-permissions [version
SNMP version]

```

The following displays an SNMP community configuration example:

```

*A:cses-A13>config>system>security>snmp# info
-----
community "uTdc9j48PBRkxn5DcSjchk" hash2 rwa version both
community "Lla.RtAyRW2" hash2 r version v2c
community "r0a159kIOfg" hash2 r version both
-----
*A:cses-A13>config>system>security>snmp#

```

3.8.2 Configuring View Options

Use the following CLI syntax to configure view options:

```

CLI Syntax: config>system>security>snmp
view view-name subtree oid-value
mask mask-value [type {included|excluded}]

```

The following displays a view configuration example:

```
*A:cses-A13>config>system>security>snmp# info
-----
      view "testview" subtree "1"
        mask ff
      exit
      view "testview" subtree "1.3.6.1.2"
        mask ff type excluded
      exit
      community "uTdc9j48PBRkxn5DcSjchk" hash2 rwa version both
      community "Lla.RtAyRW2" hash2 r version v2c
      community "r0a159kIOfg" hash2 r version both
-----
*A:cses-A13>config>system>security>snmp#
```

3.8.3 Configuring Access Options

The **access** command creates an association between a user group, a security model and the views that the user group can access. Access must be configured unless security is limited to the preconfigured access groups and views for SNMPv1 and SNMPv2. An access group is defined by a unique combination of the group name, security model and security level.

Use the following CLI syntax to configure access features:

```
CLI Syntax:  config>system>security>snmp
                access group group-name security-model security-model
                security-level security-level [context context-name
                [prefix-match]] [read view-name-1] [write view-name-2]
                [notify view-name-3]
```

The following displays an access configuration with the view configurations.

```
*A:cses-A13>config>system>security>snmp# info
-----
      view "testview" subtree "1"
        mask ff
      exit
      view "testview" subtree "1.3.6.1.2"
        mask ff type excluded
      exit
      access group "test" security-model usm security-level auth-no-pr
      ivacy read "testview" write "testview" notify "testview"
      community "uTdc9j48PBRkxn5DcSjchk" hash2 rwa version both
      community "Lla.RtAyRW2" hash2 r version v2c
      community "r0a159kIOfg" hash2 r version both
-----
*A:cses-A13>config>system>security>snmp#
```

Use the following CLI syntax to configure user group and authentication parameters:

```
CLI Syntax:  config>system>security# user user-name
                access [ftp] [snmp] [console]
                snmp
                    authentication [none] | [[hash]{md5 key | sha key }
                    privacy {none | des-key | aes-128-cfb-key key}]
                group group-name
```

The following displays a user's SNMP configuration example.

```
A:ALA-1>config>system>security# info
-----
user "testuser"
access snmp
snmp
authentication hash md5 e14672e71d3e96e7a1e19472527ee969 privacy none
group testgroup
exit
exit
...
-----
A:ALA-1>config>system>security#
```

3.8.4 Configuring USM Community Options

User-based security model (USM) community strings associate a community string with an SNMPv3 access group and its view. The access granted with a community string is restricted to the scope of the configured group.

By default, the OS implementation of SNMP uses SNMPv3. However, to implement SNMPv1 and SNMPv2c, USM community strings must be explicitly configured.

Use the following CLI syntax to configure USM community options:

```
CLI Syntax:  config>system>security>snmp
                usm-community community-string group group-name
```

The following displays a SNMP community configuration example:

```
A:ALA-1>config>system>security>snmp# info
-----
view "testview" subtree "1"
    mask ff
    exit
view "testview" subtree "1.3.6.1.2"
    mask ff type excluded
    exit
access group "test" security-model usm security-level auth-no-pr
```

```
ivacy read "testview" write "testview" notify "testview"  
    community "uTdc9j48PBRkxn5DcSjchk" hash2 rwa version both  
    community "Lla.RtAyRW2" hash2 r version v2c  
    community "r0a159kIOfg" hash2 r version both  
-----  
A:ALA-1>config>system>security>snmp#
```

The group **grouptest** was configured in the **config>system>security>snmp>access** CLI context.

3.8.5 Configuring Other SNMP Parameters

Use the following CLI syntax to modify the system SNMP options:

CLI Syntax:

```
config>system>snmp  
engineID engine-id  
general-port port  
packet-size bytes  
no shutdown
```

The following example displays the system SNMP default values:

```
A:ALA-104>config>system>snmp# info detail  
-----  
    shutdown  
    engineID "0000xxxx000000000xxxxx00"  
    packet-size 1500  
    general-port 161  
-----  
A:ALA-104>config>system>snmp#
```


3.9 SNMP Configuration Command Reference

3.9.1 Command Hierarchies

- [SNMP System Commands](#)
- [SNMP Security Commands](#)

3.9.1.1 SNMP System Commands

```

config
  — system
    — snmp
      — engineID engine-id
      — no engineID
      — general-port port
      — no general-port
      — packet-size bytes
      — no packet-size
      — streaming
        — [no] shutdown
      — [no] shutdown
  
```

3.9.1.2 SNMP Security Commands

Refer to the SR OS Services Guide for information about configuring SNMP in a VPRN service.

```

config
  — system
    — security
      — snmp
        — access group group-name security-model security-model security-
          level security-level [context context-name [prefix-match]] [read
          view-name-1] [write view-name-2] [notify view-name-3]
        — no access group group-name [security-model security-model]
          [security-level security-level] [context context-name [prefix-
          match]] [read view-name-1] [write view-name-2] [notify view-name-
          3]
        — attempts [count] [ time minutes1] [ lockout minutes2]
        — no attempts
        — community community-string [hash | hash2] access-permissions
          [version SNMP-version] [src-access-list list-name]
  
```

- **no community** *community-string* [*hash* | *hash2*]
- **usm-community** *community-string* *group* *group-name* [*src-access-list* *list-name*]
- **no usm-community** *community-string*
- **[no] src-access-list** *list-name*
 - **src-host** *host-name* *address* *ip-address*
 - **no src-host** *host-name*
- **view** *view-name* *subtree* *oid-value*
- **no view** *view-name* [*subtree* *oid-value*]
 - **mask** *mask-value* [*type* {*included* | *excluded*}]
 - **no mask**

The following commands configure user-specific SNMP features. Refer to the **Security** section for CLI syntax and command descriptions.

```

config
  — system
    — security
      — [no] user user-name
        — [no] snmp
          — authentication {[none] | [[hash] {md5 key-1 | sha key-1}
            privacy {none|des-key|aes-128-cfb-key key-2}]
          — authentication group-name
          — [no] group group-name

```

3.9.2 Command Descriptions

- [SNMP System Commands](#)
- [SNMP Security Commands](#)

3.9.2.1 SNMP System Commands

engineID

Syntax	[no] engineID <i>engine-id</i>
Context	config>system>snmp
Description	This command sets the SNMP engineID to uniquely identify the SNMPv3 node. By default, the engineID is generated using information from the system backplane.

If SNMP engine ID is changed in the **config>system>snmp> engineID** *engine-id* context, the current configuration must be saved and a reboot must be executed. If not, the previously configured SNMP communities and logger trap-target notify communities will not be valid for the new engine ID.



Note: In conformance with IETF standard RFC 2274, *User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3)*, hashing algorithms which generate SNMPv3 MD5 or SHA security digest keys use the engineID. Changing the SNMP engineID invalidates all SNMPv3 MD5 and SHA security digest keys and may render the node unmanageable.

When a chassis is replaced, use the engine ID of the first system and configure it in the new system to preserve SNMPv3 security keys. This allows management stations to use their existing authentication keys for the new system.

Ensure that the engine IDs are not used on multiple systems. A management domain can only have one instance of each engineID.

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

Default The engine ID is system generated.

Parameters *engine-id* — An identifier from 10 to 64 hexadecimal digits (5 to 32 octet number), uniquely identifying this SNMPv3 node. This string is used to access this node from a remote host with SNMPv3.

general-port

Syntax **general-port** *port-number*
no general-port

Context config>system>snmp

Description This command configures the port number used by this node to receive SNMP request messages and to send replies. SNMP notifications generated by the agent are sent from the port specified in the **config>log>snmp-trap-group>trap-target** CLI command.

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

Default general-port 161

Parameters *port-number* — The port number used to send SNMP traffic other than traps.

Values 1 to 65535

packet-size

Syntax	packet-size <i>bytes</i> no packet-size
Context	config>system>snmp
Description	This command configures the maximum SNMP packet size generated by this node. The no form of this command restores the default value.
Default	packet-size 1500
Parameters	<i>bytes</i> — The SNMP packet size in bytes. Values 484 to 9216

snmp

Syntax	snmp
Context	config>system
Description	This command creates the context to configure SNMP parameters.

streaming

Syntax	streaming
Context	config>system>snmp
Description	This command enables the proprietary SNMP request/response bundling and TCP-based transport mechanism for optimizing network management of the router nodes. In higher latency networks, synchronizing router MIBs from network management via streaming takes less time than synchronizing via classic SNMP UDP requests. Streaming operates on TCP port 1491 and runs over IPv4 or IPv6.

shutdown

Syntax	[no] shutdown
Context	config>system>snmp>streaming
Description	This command administratively disables proprietary SNMP request/response bundling and TCP-based transport mechanism for optimizing network management of the router nodes.. The no form of the command administratively re-enables SNMP request/response bundling and TCP-based transport mechanism.

Default shutdown

shutdown

Syntax [no] shutdown

Context config>system>snmp

Description This command administratively disables SNMP agent operations. System management can then only be performed using the command line interface (CLI). Shutting down SNMP does not remove or change configuration parameters other than the administrative state. This command does not prevent the agent from sending SNMP notifications to any configured SNMP trap destinations. SNMP trap destinations are configured under the **config>log>snmp-trap-group** context.

This command is automatically invoked in the event of a reboot when the processing of the configuration file fails to complete or when an SNMP persistent index file fails while the **bof persist on** command is enabled.

The **no** form of the command administratively enables SNMP which is the default state.

Default no shutdown

3.9.2.2 SNMP Security Commands

access group

Syntax [no] access group *group-name* security-model *security-model* security-level *security-level* [**context** *context-name* [**prefix-match**]] [**read** *view-name-1*] [**write** *view-name-2*] [**notify** *view-name-3*]

Context config>system>security>snmp

Description This command creates an association between a user group, a security model, and the views that the user group can access. Access parameters must be configured unless security is limited to the preconfigured access groups and views for SNMPv1 and SNMPv2. An access group is defined by a unique combination of the group name, security model and security level.

Access groups are used by the `usm-community` command.

Access must be configured unless security is limited to SNMPv1/SNMPv2c with community strings (see the [community](#)).

Default access group configurations cannot be modified or deleted.

To remove the user group with associated, security model(s), and security level(s), use:

no access group *group-name*

To remove a security model and security level combination from a group, use:

no access group *group-name* **security-model** {**snmpv1** | **snmpv2c** | **usm**} **security-level** {**no-auth-no-privacy** | **auth-no-privacy** | **privacy**}

Parameters

group-name — Specify a unique group name up to 32 characters.

security-model {**snmpv1** | **snmpv2c** | **usm**} — Specifies the security model required to access the views configured in this node. A group can have multiple security models. For example, one view may only require SNMPv1/ SNMPv2c access while another view may require USM (SNMPv3) access rights.

security-level {**no-auth-no-priv** | **auth-no-priv** | **privacy**} — Specifies the required authentication and privacy levels to access the views configured in this node.

security-level no-auth-no-privacy — Specifies that no authentication and no privacy (encryption) is required. When configuring the user's authentication, select the **none** option.

security-level auth-no-privacy — Specifies that authentication is required but privacy (encryption) is not required. When this option is configured, both the **group** and the **user** must be configured for authentication.

security-level privacy — Specifies that both authentication and privacy (encryption) is required. When this option is configured, both the **group** and the user must be configured for **authentication**. The user must also be configured for **privacy**.

context *context-name* — Specifies a set of SNMP objects that are associated with the context-name.

The *context-name* is treated as either a full context-name string or a context name prefix depending on the keyword specified (**exact** or **prefix**).

prefix-match — Specifies the context name **prefix-match** keywords, **exact** or **prefix**. This parameter applies only to the 7750 SR.

The VPRN context names begin with a **vprn** prefix. The numerical value is associated with the service ID that the VPRN was created with and identifies the service in the service domain. For example, when a new VPRN service is created such as **config>service>vprn 2345 customer 1**, a VPRN with context name **vprn2345** is created.

The **exact** keyword specifies that an exact match between the context name and the prefix value is required. For example, when **context vprn2345 exact** is entered, matches for only **vprn2345** are considered.

The **prefix** keyword specifies that only a match between the prefix and the starting portion of context name is required. If only the **prefix** keyword is specified, simple wildcard processing is used. For example, when **context vprn prefix** is entered, all **vprn** contexts are matched.

Default **exact**

read *view-name* — Specifies the keyword and variable of the view to read the MIB objects.
This command must be configured for each view to which the group has read access.

Default none

write *view-name* — Specifies the keyword and variable of the view to configure the contents of the agent.

This command must be configured for each view to which the group has write access.

Values Up to 32 characters

notify *view-name* — specifies keyword and variable of the view to send a trap about MIB objects.

This command must be configured for each view to which the group has notify access.

Values none

attempts

Syntax **attempts** [*count*] [**time** *minutes1*] [**lockout** *minutes2*]
no attempts

Context config>system>security>snmp

Description This command configures a threshold value of unsuccessful SNMP connection attempts allowed in a specified time frame. The command parameters are used to counter denial of service (DOS) attacks through SNMP.

If the threshold is exceeded, the host is locked out for the lockout time period.

If multiple **attempts** commands are entered, each command overwrites the previously entered command.

The **no** form of the command restores the default values, in which 20 failed SNMP attempts are allowed in a 5 minute period with a 10 minute lockout for the host if exceeded.

Default attempts 20 time 5 lockout 10

Parameters *count* — The number unsuccessful SNMP attempts allowed for the specified **time**.

Default 20

Values 1 to 64

time *minutes1* — The period of time, in minutes, that a specified number of unsuccessful attempts can be made before the host is locked out.

Default 5

Values 0 to 60

lockout *minutes2* — The lockout period in minutes where the host is not allowed to login. When the host exceeds the attempted count times in the specified time, then that host is locked out from any further login attempts for the configured time period.

Default 10

Values 0 to 1440

community

Syntax **community** *community-string* [**hash** | **hash2**] *access-permissions* [**version** *SNMP-version*]
[**src-access-list** *list-name*]
no community *community-string* [**hash** | **hash2**]

Context config>system>security>snmp

Description This command creates SNMP community strings for SNMPv1 and SNMPv2c access. This command is used in combination with the predefined access groups and views. To create custom access groups and views and associate them with SNMPv1 or SNMPv2c access use the [usm-community](#) command.

When configured, community implies a security model for SNMPv1 and SNMPv2c only.

For SNMPv3 security, the [access group](#) command must be configured.

The **no** form of the command removes the specified community string.

Parameters *community-string* — Configure the SNMPv1 and/or SNMPv2c community string.

Values *community-string* — 32 characters maximum

hash-key — 33 characters maximum

hash2-key — 96 characters maximum

hash — Specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified

hash2 — Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

access-permissions — Configures the access permissions for objects in the MIB.

r — Grants only read access to objects in the MIB, except security objects, using the internal "snmp-ro" access group and the "no-security" snmp view.

rw — Grants read and write access to all objects in the MIB, using the internal "snmp-rw" access group and the "no-security" snmp view.

rwa — Grants read and write access to all objects in the MIB, including security, using the internal "snmp-rwa" access group and the "iso" snmp view.

mgmt — Assigns a unique SMMP community string for SNMP access via the "management" router instance. This community uses the internal "snmp-mgmt" access group and the "mgmt" snmp view.

vpls-mgmt — Assigns a unique SNMP community string for SNMP access via the "vpls-management" router instance. This community uses the internal "snmp-vpls-mgmt" access group and "mgmt-view" snmp view.

version {v1 | v2c | both} — Configures the scope of the community string to be for SNMPv1, SNMPv2c, or both SNMPv1 and SNMPv2c access.

Default both

list-name — Configures the **community** to reference a specific [src-access-list](#), which will be used to validate the source IP address of all received SNMP requests that use this **community**. Multiple **community**, **usm-community**, or **vprn snmp community** instances can reference the same **src-access-list**.

mask

Syntax **mask** *mask-value* [**type** {**included** | **excluded**}]
no mask

Context config>system>security>snmp>view view-name

Description The mask value and the mask type, along with the *oid-value* configured in the **view** command, determines the access of each sub-identifier of an object identifier (MIB subtree) in the view.

Each bit in the mask corresponds to a sub-identifier position. For example, the most significant bit for the first sub-identifier, the next most significant bit for the second sub-identifier, and so on. If the bit position on the sub-identifier is available, it can be included or excluded.

For example, the MIB subtree that represents MIB-II is 1.3.6.1.2.1. The mask that catches all MIB-II would be 0xfc or 0b11111100.

Only a single mask may be configured per view and OID value combination. If more than one entry is configured, each subsequent entry overwrites the previous entry.

Per RFC 2575, *View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP)*, each MIB view is defined by two sets of view subtrees, the included view subtrees, and the excluded view subtrees. Every such view subtree, both the included and the excluded ones, are defined in this table. To determine if a particular object instance is in a particular MIB view, compare the object instance's object identifier (OID) with each of the MIB view's active entries in this table. If none match, then the object instance is not in the MIB view. If one or more match, then the object instance is included in, or excluded from, the MIB view according to the value of `vacmViewTreeFamilyType` in the entry whose value of `vacmViewTreeFamilySubtree` has the most sub-identifiers.

The **no** form of this command removes the mask from the configuration.

- Parameters** *mask-value* — The mask value associated with the OID value determines whether the sub-identifiers are included or excluded from the view. (Default: all 1s)
- The mask can be entered either:
- In hex. For example, 0xfc.
 - In binary. For example, 0b11111100.



Note: If the number of bits in the bit mask is less than the number of sub-identifiers in the MIB subtree, then the mask is extended with ones until the mask length matches the number of sub-identifiers in the MIB subtree.

type {included | excluded} — Specifies whether to include or exclude MIB subtree objects. *included* - All MIB subtree objects that are identified with a 1 in the mask are available in the view. (Default: *included*).

excluded - All MIB subtree objects that are identified with a 1 in the mask are denied access in the view. (Default: *included*).

Default **included**

snmp

- Syntax** **snmp**
- Context** config>system>security
- Description** This command creates the context to configure SNMPv1, SNMPv2, and SNMPv3 parameters.

src-access-list

- Syntax** **src-access-list** *list-name*
no src-access-list *list-name*
- Context** config>system>security>snmp
- Description** This command is used to identify a list of source IP addresses that can be used to validate SNMPv1 and SNMPv2c requests once the list is associated with one or more SNMPv1 and SNMPv2c communities.

An src-address-list referenced by one or more **community** instances is used to verify the source IP addresses of an SNMP request using the **community** regardless of which VPRN/VRF interface (or 'Base' interface) the request arrived on. For example, if an SNMP request arrives on an interface in vprn 100 but the request is referencing a **community**, then the source IP address in the packet would be validated against the src-address-list configured for

the **community**. This occurs regardless of whether the request is destined to a VPRN interface address and the VPRN has SNMP access enabled, or the request is destined to the base system address via GRT leaking. If the request's source IP address does not match the *ip-address* of any of the **src-hosts** contained in the list, then the request will be discarded and logged as an SNMP authentication failure.

Using `src-access-list` validation can have an impact on the time it takes for an SR OS node to reply to an SNMP request. It is recommended to keep the lists short, including only the addresses that are needed, and to place SNMP managers that send the highest volume of requests, such as the NSP NFM-P, at the top of the list.

You can configure a maximum of 16 **src-access-lists**. Each **src-access-list** can contain a maximum of 16 **src-hosts**.

The **no** form of this command removes the named `src-access-list`. You cannot remove an **src-access-list** that is referenced by one or more **community** instances.

Parameters *list-name* — Configures the name or key of the **src-access-list**. The *list-name* parameter must begin with a letter (a-z or A-Z).

src-host

Syntax **src-host** *host-name* **address** *ip-address*
no src-host *host-name*

Context `config>system>security>snmp>src-access-list`

Description This command is used to configure a source IP address entry that can be used to validate SNMPv1 and SNMPv2c requests.

The **no** form of this command removes the specified entry.

Parameters *host-name* — Configures the name of the **src-host** entry.
ip-address — Configures an allowed source address for SNMP requests. This can be an IPv4 or IPv6 address.

Values `ipv4-address: a.b.c.d`
`ipv6-address: x:x:x:x:x:x:x`
`x:x:x:x:x:d.d.d.d`
`x: [0..FFFF]H`
`d: [0..255]D`

usm-community

Syntax **usm-community** *community-string* **group** *group-name* [*src-access-list list-name*]
no usm-community *community-string*

Context	config>system>security>snmp
Description	<p>This command is used to associate a community string with an SNMPv3 access group and its view. The access granted with a community string is restricted to the scope of the configured group.</p> <p>Nokia's SR OS implementation of SNMP uses SNMPv3. In order to implement SNMPv1 and SNMPv2c configurations, several access groups are predefined. In order to implement SNMP with security features (Version 3), security models, security levels, and USM communities must be explicitly configured. Optionally, additional views which specify more specific OIDs (MIB objects in the subtree) can be configured.</p> <p>The no form of this command removes a community string.</p>
Parameters	<p><i>community-string</i> — Specifies the SNMPv1/SNMPv2c community string to determine the SNMPv3 access permissions to be used.</p> <p><i>group</i> — Specifies the group that governs the access rights of this community string. This group must be configured first in the config system security snmp access group context.</p> <p><i>list-name</i> — Specifies the usm-community to reference a specific src-access-list that will be used to validate the source IP address of all received SNMP requests that use this usm-community. Multiple community, usm-community, or vprn snmp community instances can reference the same src-access-list.</p>

view

Syntax	<p>view <i>view-name</i> subtree <i>oid-value</i></p> <p>no view <i>view-name</i> [subtree <i>oid-value</i>]</p>
Context	config>system>security>snmp
Description	<p>This command configures a view. Views control the accessibility of a MIB object within the configured MIB view and subtree. Object identifiers (OIDs) uniquely identify MIB objects in the subtree. OIDs are organized hierarchically with specific values assigned by different organizations.</p> <p>Once the subtree (OID) is identified, a mask can be created to select the portions of the subtree to be included or excluded for access using this particular view. See the mask command. The view(s) configured with this command can subsequently be used in read, write, and notify commands which are used to assign specific access group permissions to created views and assigned to particular access groups.</p> <p>Multiple subtrees can be added or removed from a view name to tailor a view to the requirements of the user access group.</p> <p>The no view <i>view-name</i> command removes a view and all subtrees.</p> <p>The no view <i>view-name</i> subtree <i>oid-value</i> removes a sub-tree from the view name.</p>

Default No views are defined.

Parameters *view-name* — Enter a 1 to 32 character view name. (Default: *none*)

oid-value — The object identifier (OID) value for the *view-name*. This value, for example, 1.3.6.1.6.3.11.2.1, combined with the mask and include and exclude statements, configures the access available in the view.

It is possible to have a view with different subtrees with their own masks and include and exclude statements. This allows for customizing visibility and write capabilities to specific user requirements.

3.10 SNMP Show Command Reference

3.10.1 Command Hierarchies

3.10.1.1 Show Commands

```
show
  — snmp
    — counters
    — streaming
      — counters
  — system
    — information
    — security
      — access-group [group-name]
      — authentication [statistics]
      — password-options
      — per-peer-queuing
      — profile [profile-name]
      — snmp
        — community [community-string]
        — src-access-list [list-name]
    — ssh
    — user [user-id] [detail]
    — view [view-name] [detail]
```

3.10.2 Command Descriptions

The command outputs in the following section are examples only; actual displays may differ depending on supported functionality and user configuration.

3.10.2.1 Show Commands

counters

Syntax **counters**

Context show>snmp

Description This command displays SNMP counters information. SNMP counters will continue to increase even when SNMP is shut down. Some internal modules communicate using SNMP packets.

Output Counters Output

[Table 40](#) describes the SNMP counters output fields.

Table 40 Show Counters Output Fields

Label	Description
in packets	Displays the total number of messages delivered to SNMP from the transport service.
in gets	Displays the number of SNMP get request PDUs accepted and processed by SNMP.
in getnexts	Displays the number of SNMP get next PDUs accepted and processed by SNMP.
in sets	Displays the number of SNMP set request PDUs accepted and processed by SNMP.
out packets	Displays the total number of SNMP messages passed from SNMP to the transport service.
out get responses	Displays the number of SNMP get response PDUs generated by SNMP.
out traps	Displays the number of SNMP Trap PDUs generated by SNMP.
variables requested	Displays the number of MIB objects requested by SNMP.
variables set	Displays the number of MIB objects set by SNMP as the result of receiving valid SNMP set request PDUs.

Sample Output

```
A:ALA-1# show snmp counters
=====
SNMP counters:
=====
  in packets : 463
-----
  in gets    : 93
  in getnexts : 0
  in sets    : 370
  out packets: 463
-----
  out get responses : 463
  out traps        : 0
  variables requested: 33
  variables set     : 497
```

```
=====
A:ALA-1#
```

counters

- Syntax** **counters**
- Context** show>snmp>streaming
- Description** This command displays counters information for the proprietary SNMP streaming protocol. Output: Counters Output - The following table describes SNMP streaming counters output fields.
- Output** Counters Output

[Table 41](#) describes the SNMP streaming counters output fields.

Table 41 Show Streaming Counters Output Fields

Label	Description
in getTables	Displays the number of GetTable request packets received.
in getManys	Displays the number of GetMany request packets received.
out responses	Displays the number of response packets sent.

Sample Output

```
*A:Dut-B# show snmp streaming counters
=====
STREAMING counters:
=====
      in getTables   : 772
      in getManys   : 26
-----
      out responses  : 848
=====
```

information

- Syntax** **information**
- Context** show>system
- Description** This command lists the SNMP configuration and statistics.
- Output** System Information Output Fields

Table 42 describes system information output fields.

Table 42 Show System Information Output Fields

Label	Description
System Name	The name configured for the device.
System Type	Indicates the SR OS platform type (for example, 7750 SR-12).
Chassis Topology	Indicates the inter-chassis topology mode in which the system is operating. Standalone indicates that the system is comprised of a single physical router chassis. Extended (XRS-40) on a 7950 XRS-based system indicates that two router chassis are connected together in a back-to-back topology with no additional switch fabric chassis. An extended chassis topology is comprised of two XRS-20 chassis and is also known as an XRS-40 system.
System Contact	The text string that identifies the contact name for the device.
System Location	The text string that identifies the location of the device.
System Coordinates	The text string that identifies the system coordinates for the device location. For example, "37.390 -122.0550" is read as latitude 37.390 north and longitude 122.0550 west.
System Up Time	The time since the last reboot.
SNMP Port	The port which SNMP sends responses to management requests.
SNMP Engine ID	The ID for either the local or remote SNMP engine to uniquely identify the SNMPv3 node.
SNMP Max Message Size	The maximum size SNMP packet generated by this node.
SNMP Admin State	Enabled — SNMP is administratively enabled.
	Disabled — SNMP is administratively disabled.
SNMP Oper State	Enabled — SNMP is operationally enabled.
	Disabled — SNMP is operationally disabled.

Table 42 Show System Information Output Fields (Continued)

Label	Description
SNMP Index Boot Status	Persistent — Persistent indexes at the last system reboot was enabled.
	Disabled — Persistent indexes at the last system reboot was disabled.
SNMP Sync State	The state when the synchronization of configuration files between the primary and secondary CPMs finish.
Telnet/SSH/FTP Admin	Displays the administrative state of the Telnet, SSH, and FTP sessions.
Telnet/SSH/FTP Oper	Displays the operational state of the Telnet, SSH, and FTP sessions.
BOF Source	The boot location of the BOF.
Image Source	primary — Specifies whether the image was loaded from the primary location specified in the BOF.
	secondary — Specifies whether the image was loaded from the secondary location specified in the BOF.
	tertiary — Specifies whether the image was loaded from the tertiary location specified in the BOF.
Config Source	primary — Specifies whether the configuration was loaded from the primary location specified in the BOF.
	secondary — Specifies whether the configuration was loaded from the secondary location specified in the BOF.
	tertiary — Specifies whether the configuration was loaded from the tertiary location specified in the BOF.
Last Booted Config File	Displays the URL and filename of the configuration file used for the most recent boot.
Last Boot Cfg Version	Displays the version of the configuration file used for the most recent boot.
Last Boot Config Header	Displays header information of the configuration file used for the most recent boot.
Last Boot Index Version	Displays the index version used in the most recent boot.
Last Boot Index Header	Displays the header information of the index used in the most recent boot.
Last Saved Config	Displays the filename of the last saved configuration.

Table 42 Show System Information Output Fields (Continued)

Label	Description
Time Last Saved	Displays the time the configuration was most recently saved.
Changes Since Last Save	Yes — The configuration changed since the last save.
	No — The configuration has not changed since the last save.
Time Last Modified	Displays the time of the last modification.
Max Cfg/BOF Backup Rev	The maximum number of backup revisions maintained for a configuration file. This value also applies to the number of revisions maintained for the BOF file.
Cfg-OK Script	URL — The location and name of the CLI script file executed following successful completion of the boot-up configuration file execution.
	N/A — No CLI script file is executed.
Cfg-OK Script Status	Successful/Failed — The results from the execution of the CLI script file specified in the Cfg-OK Script location.
	Not used — No CLI script file was executed.
Cfg-Fail Script	URL — The location and name of the CLI script file executed following a failed boot-up configuration file execution.
	Not used — No CLI script file was executed.
Cfg-Fail Script Status	Successful/Failed — The results from the execution of the CLI script file specified in the Cfg-Fail Script location.
	Not used — No CLI script file was executed.
Management IP address	The Management IP address of the node.
DNS Server	The DNS address of the node.
DNS Domain	The DNS domain name of the node.
BOF Static Routes	To — The static route destination.
	Next Hop — The next hop IP address used to reach the destination.
	Metric — Displays the priority of this static route versus other static routes.
	None — No static routes are configured.

Sample Output

The following is an output example of the 7950 XRS:

```
*A:7950 XRS-20# show system information
=====
System Information
=====
System Name           : 7950 XRS-20
System Type           : 7950 XRS-20
Chassis Topology     : Standalone
System Version       : C-10.0.B1-103
System Contact       :
System Location      :
System Coordinates   :
System Active Slot   : A
System Up Time       : 19 days, 18:43:59.66 (hr:min:sec)

SNMP Port            : 161
SNMP Engine ID       : 0000197f0000ac9fff000000
SNMP Engine Boots    : 1
SNMP Max Message Size : 1500
SNMP Admin State     : Disabled
SNMP Oper State      : Disabled
SNMP Index Boot Status : Not Persistent
SNMP Sync State      : N/A

Tel/Tel6/SSH/FTP Admin : Enabled/Disabled/Enabled/Disabled
Tel/Tel6/SSH/FTP Oper  : Up/Down/Up/Down

BOF Source           : cf3:
Image Source         : primary
Config Source        : primary
Last Booted Config File: ftp://*:*@kandhcp214/tftpboot/bksimgrp31/images/bksim3
                        106/bksim3106.cfg
Last Boot Cfg Version : WED MAY 23 11:58:26 2012 UTC
Last Boot Config Header: # TiMOS-C-14.0.B1-217 cpm/
x86_64 Nokia 7950 XRS Copyright (c)
                        2000-2016 Nokia. # All rights
                        reserved. All use subject to applicable license
                        agreements. # Built on Wed Jul 13 19:09:32 PDT 2016
                        by builder in /rel14.0/b1/B1-217/panos/main

Last Boot Index Version: N/A
Last Boot Index Header : # TiMOS-C-0.0.I3339 cpm/i386 Nokia 7950 XRS
                        Copyright (c) 2000-2016 Nokia. # All rights
                        reserved. All use subject to applicable license
                        agreements. # Built on Tue May 22 18:46:56 PDT 2016
                        by builder in /rel14.0/I3339/panos/main # Generated
                        WED MAY 23 11:58:26 2016 UTC

Last Saved Config     : ftp://*:*@kandhcp214/tftpboot/bksimgrp31/images/bksim3
                        106/bksim3106.cfg
Time Last Saved       : 2012/05/28 10:38:31
Changes Since Last Save: Yes
User Last Modified    : admin
Time Last Modified    : 2012/06/06 17:06:15
Max Cfg/BOF Backup Rev : 5
Cfg-OK Script         : N/A
Cfg-OK Script Status  : not used
Cfg-Fail Script       : N/A
Cfg-Fail Script Status : not used
```

```

Management IP Addr      : 138.120.214.159/24
Primary DNS Server      : 138.120.252.56
Secondary DNS Server    : 138.120.252.48
Tertiary DNS Server     : 138.120.252.49
DNS Domain              : labs.ca.nokia.com
DNS Resolve Preference : ipv4-only
BOF Static Routes      :
  To                    Next Hop
  135.244.0.0/16       138.120.214.1
  138.120.0.0/16      138.120.214.1

ICMP Vendor Enhancement: Disabled
=====

```

access-group

- Syntax** `access-group group-name`
- Context** `show>system>security`
- Description** This command displays access-group information.
- Output** System Information Output

[Table 43](#) describes the access-group output fields.

Table 43 Show System Security Access-Group Output Fields

Label	Description
Group name	The access group name.
Security model	The security model required to access the views configured in this node.
Security level	Specifies the required authentication and privacy levels to access the views configured in this node.
Read view	Specifies the view to read the MIB objects.
Write view	Specifies the view to configure the contents of the agent.
Notify view	Specifies the view to send a trap about MIB objects.
No. of access groups	The total number of configured access groups.

Sample Output

```

A:ALA-1# show system security access-group
=====

```

```

Access Groups
=====
group name      security      security      read           write          notify
                model         level         view           view           view
-----
snmp-ro         snmpv1       none          no-security    no-security    no-security
snmp-ro         snmpv2c     none          no-security    no-security    no-security
snmp-rw         snmpv1       none          no-security    no-security    no-security
snmp-rw         snmpv2c     none          no-security    no-security    no-security
snmp-rwa        snmpv1       none          iso            iso            iso
snmp-rwa        snmpv2c     none          iso            iso            iso
snmp-trap       snmpv1       none          no-security    no-security    iso
snmp-trap       snmpv2c     none          no-security    no-security    iso
-----
No. of Access Groups: 8
=====
A:ALA-1#

A:ALA-1# show system security access-group detail
=====
Access Groups
=====
group name      security      security      read           write          notify
                model         level         view           view           view
-----
snmp-ro         snmpv1       none          no-security    no-security    no-security
-----
No. of Access Groups:
...
=====
A:ALA-1#

```

authentication

- Syntax** authentication [statistics]
- Context** show>system>security
- Description** This command displays authentication information.
- Output** Authentication Output

[Table 44](#) describes the authentication output fields.

Table 44 Show Authentication Output Fields

Label	Description
sequence	The authentication order in which password authentication, authorization, and accounting is attempted among RADIUS, TACACS+, and local passwords.
server address	The address of the RADIUS, TACACS+, or local server.

Table 44 Show Authentication Output Fields (Continued)

Label	Description (Continued)
status	The status of the server.
type	The type of server.
timeout (secs)	Number of seconds the server will wait before timing out.
retry count	The number of attempts to retry contacting the server.
radius admin status	The administrative status of the RADIUS protocol operation.
tacplus admin status	The administrative status of the TACACS+ protocol operation.
health check	Specifies whether the RADIUS and TACACS+ servers will be periodically monitored. Each server will be contacted every 30 seconds. If in this process a server is found to be unreachable, or a previously unreachable server starts responding, based on the type of the server, a trap will be sent.
No. of Servers	The total number of servers configured.

Sample Output

```

A:ALA-49>show>system>security# authentication
=====
Authentication                sequence : radius tacplus local
=====
type                          status  timeout (secs)  retry count
server address
-----
radius                         up      5                5
 10.10.10.103
radius                         up      5                5
 10.10.10.1
radius                         up      5                5
 10.10.10.2
radius                         up      5                5
 10.10.10.3
-----
radius admin status : up
tacplus admin status : up
health check       : enabled (interval 30)
-----
No. of Servers: 4
=====
A:ALA-49>show>system>security#

```

password-options

Syntax password-options

- Context** show>system>security
- Description** This command displays password options.
- Output** Password-Options Output

Table 45 describes password-options output fields.

Table 45 Show Password-Options Output Fields

Label	Description
Password aging in days	Number of days a user password is valid before the user must change his password.
Number of invalid attempts permitted per login	Displays the maximum number of unsuccessful login attempts allowed for a user.
Time in minutes per login attempt	Displays the time in minutes that user is to be locked out.
Lockout period (when threshold breached)	Displays the number of minutes the user is locked out if the threshold of unsuccessful login attempts has exceeded.
Authentication order	Displays the most preferred method to authenticate and authorize a user.
Configured complexity options	Displays the complexity requirements of locally administered passwords, HMAC-MD5-96, HMAC-SHA-96 and DES-keys configured in the authentication section.
Minimum password length	Displays the minimum number of characters required in the password.

Sample Output

```
A:ALA-48>show>system>security# password-options
=====
Password Options
=====
Password aging in days                : 365
Number of invalid attempts permitted per login : 5
Time in minutes per login attempt      : 5
Lockout period (when threshold breached) : 20
Authentication order                  : radius tacplus local
Configured complexity options          :
Minimum password length                : 8
=====
A:ALA-48>show>system>security#
```

per-peer-queuing

- Syntax** `per-peer-queuing`
- Context** `show>system>security`
- Description** This command displays displays the number of queues in use by the Qchip, which in turn is used by PPQ, CPM filter, SAP, etc.
- Output** Per-Peer_Queueing Output

[Table 46](#) describes the per-peer-queuing output fields.

Table 46 Show per-peer-queuing Output Fields

Label	Description
Per Peer Queuing	Displays whether per-peer-queuing is enabled or disabled. When enabled, a peering session is established and the router will automatically allocate a separate CPM hardware queue for that peer. When disabled, no hardware queuing per peer occurs.
Total Num of Queues	Displays the total number of CPM hardware queues.
Num of Queues In Use	Displays the number of CPM hardware queues that are in use.

Sample Output

```
A:ALA-48>show>system>security# per-peer-queuing
=====
CPM Hardware Queuing
=====
Per Peer Queuing      : Enabled
Total Num of Queues   : 8192
Num of Queues In Use  : 0
=====
A:ALA-48>show>system>security#
```

profile

- Syntax** `profile [profile-name]`
- Context** `show>system>security`
- Description** This command displays user profiles for CLI command tree permissions.
- Parameters** *profile-name* — Specify the profile name to display information about a single user profile. If no profile name is displayed, the entire list of profile names are listed.
- Output** Profile Output

Table 47 describes the profile output fields.

Table 47 Show Profile Output Fields

Label	Description
User Profile	default — The action to be given to the user profile if none of the entries match the command.
	administrative — specifies the administrative state for this profile.
Def. Action	none — No action is given to the user profile when none of the entries match the command.
	permit-all — The action to be taken when an entry matches the command.
Entry	10 - 80 — Each entry represents the configuration for a system user.
Description	A text string describing the entry.
Match Command	administrative — Enables the user to execute all commands.
	configure system security — Enables the user to execute the config system security command.
	enable-admin — Enables the user to enter a special administrative mode by entering the enable-admin command.
	exec — Enables the user to execute (exec) the contents of a text file as if they were CLI commands entered at the console.
	exit — Enables the user to execute the exit command.
	help — Enables the user to execute the help command.
	logout — Enables the user to execute the logout command.
	password — Enables the user to execute the password command.
	show config — Enables the user to execute the show config command.
	show — Enables the user to execute the show command.
Action	show system security — Enables the user to execute the show system security command.
	permit — Enables the user access to all commands.
	deny-all — Denies the user access to all commands.

Sample Output

```
A:ALA-48>config>system>snmp# show system security profile
=====
User Profile
=====
User Profile : test
Def. Action  : none
-----
Entry       : 1
Description :
Match Command:
Action      : unknown
=====
User Profile : default
Def. Action  : none
-----
Entry       : 10
Description :
Match Command: exec
Action      : permit
-----
Entry       : 20
Description :
Match Command: exit
Action      : permit
-----
Entry       : 30
Description :
Match Command: help
Action      : permit
-----
...
-----
Entry       : 80
Description :
Match Command: enable-admin
Action      : permit
=====

User Profile : administrative
Def. Action  : permit-all
-----
Entry       : 10
Description :
Match Command: configure system security
Action      : permit
-----
Entry       : 20
Description :
Match Command: show system security
Action      : permit
-----
No. of profiles: 3
=====
A:ALA-48>config>system>snmp#
```

snmp

- Syntax** `snmp`
- Context** `show>system>security`
- Description** This command enables the context to show SNMP information.

community

- Syntax** `community [community-string]`
- Context** `show>system>security>snmp`
- Description** This command lists SNMP communities and characteristics. Including the *community-name* parameter modifies the output to include all details for the specified community, including the source IP address list and validation failure counters.
- Output** [Table 48](#) describes the community output fields.

Sample Output

Table 48 Show Community Output Fields

Label	Description
Community	The community string name for SNMPv1 and SNMPv2c access only.
Access	r — The community string allows read-only access.
	rw — The community string allows read-write access.
	rwa — The community string allows read-write access.
	mgmt — The unique SNMP community string assigned to the management router.
	vpls-mgmt — The unique SNMP community string assigned for vpls management
View	The view name.
Version	The SNMP version.
Group Name	The access group name.
src-access-list	The name of the list of source IP addresses that are allowed to use the community, as configured using the community configuration command.

Table 48 Show Community Output Fields (Continued)

Label	Description
authFailures	The number of SNMP requests that have failed validation using this community .
No of Communities	The total number of configured community strings.



Note: The system-created communities that begin with “cli-” are only used for internal CLI management purposes and are not exposed to external SNMP access.

```
A:ALA-1# show system security snmp community

=====
Communities
=====
community      access  view          version  group name
-----
cli-li-readwrite  n/a    li-view       v2c     cli-li-readwrite
cli-readonly     r      iso           v2c     cli-readonly
cli-readwrite    rw     iso           v2c     cli-readwrite
my-privatel      rw     iso           v1 v2c  snmp-rwa
my-public2       r      no-security   v1 v2c  snmp-ro
test-123         rwa   n/a           v2c     snmp-trap
-----
No. of Communities: 6
=====
A:ALA-1#

A:ALA-1# show system security snmp community "my-public2"

=====
Communities
=====
community      access  view          version  group name
-----
my-public2     r      no-security   v1 v2c  snmp-ro
              my-list1
              5
=====
A:ALA-1#
```

src-access-list

Syntax `src-access-list [list-name]`

Context `show>system>security>snmp`

Description This command displays source access lists and the hosts for each. Including the *list-name* parameter modifies the output show only the specified **src-access-list**.

Output Source Access List Output

Table 49 describes the source access list output fields.

Sample Output

Table 49 Show Source Access List Output Fields

Label	Description
List Name	The name of the src-access-list .
Host Name	The name of the src-host .
Host Address	The IP address of the src-host .
Total Access Lists	The total number of source access lists displayed.

```
A:ALA-1# show system security snmp src-access-list
=====
Source Access Lists
=====
List Name
  HostName                Host Address
-----
L1
  H1                      100.100.100.1
  H2                      100.100.100.2
L2
  HA                      100.100.101.1
  HB                      100.100.101.2
-----
Total Access Lists: 2
=====
A:ALA-1#
```

```
A:ALA-1# show system security snmp src-access-list L1
=====
Source Access Lists
=====
List Name
  HostName                Host Address
-----
L1
  H1                      100.100.100.1
  H2                      100.100.100.2
-----
Total Access Lists: 1
=====
A:ALA-1#
```

ssh

Syntax	ssh
Context	show>system>security
Description	This command displays all the SSH sessions as well as the SSH status and fingerprint.
Output	SSH Options Output

[Table 50](#) describes SSH output fields.

Table 50 Show SSH Output Fields

Label	Description
SSH status	SSH is enabled — Displays that SSH server is enabled.
	SSH is disabled — Displays that SSH server is disabled.
Key fingerprint	The key fingerprint is the server's identity. Clients trying to connect to the server verify the server's fingerprint. If the server fingerprint is not known, the client may not continue with the SSH session since the server might be spoofed.
Connection	The IP address of the connected router(s) (remote client).
Encryption	des — Data encryption using a private (secret) key.
	3des — An encryption method that allows proprietary information to be transmitted over untrusted networks.
Username	The name of the user.
Number of SSH sessions	The total number of SSH sessions.

Sample output

```
A:ALA-7# show system security ssh
SSH is enabled
Key fingerprint: 34:00:f4:97:05:71:aa:b1:63:99:dc:17:11:73:43:83
=====
Connection Encryption Username
=====
192.168.5.218 3des admin
-----
Number of SSH sessions : 1
=====
A:ALA-7#

A:ALA-49>config>system>security# show system security ssh
SSH is disabled
A:ALA-49>config>system>security#
```

user

- Syntax** `users [user-id] [detail]`
- Context** `show>system>security`
- Description** This command displays user information.
- Output** User Output

[Table 51](#) describes user information output fields.

Table 51 Show User Output Fields

Label	Description
User ID	The name of a system user.
Need New PWD	Yes — The user must change his password at the next login.
	No — The user is not forced to change his password at the next login.
User Permission	Console — Specifies whether the user is permitted console/Telnet access.
	FTP — Specifies whether the user is permitted FTP access.
	SNMP — Specifies whether the user is permitted SNMP access.
Password expires	The date on which the current password expires.
Attempted logins	The number of times the user has attempted to login irrespective of whether the login succeeded or failed.
Failed logins	The number of unsuccessful login attempts.
Local Conf.	Y — Password authentication is based on the local password database.
	N — Password authentication is not based on the local password database.

Sample Output

```
A:ALA-1# show system security user
=====
Users
=====
user id          need   user permissions  password   attempted  failed  local
                  new pwd console ftp snmp  expires   logins   logins  conf
-----
admin            n      y      n  n      never     2        0       y
testuser        n      n      n  y      never     0        0       y
```

Number of users : 2

view

- Syntax** view [view-name] [detail]
- Context** show>system>security
- Description** This command lists one or all views and permissions in the MIB-OID tree.
- Output** System Security View Output

Table 52 describes system security view output fields.

Table 52 Show System Security View Output Fields

Label	Description
View name	The name of the view. Views control the accessibility of a MIB object within the configured MIB view and subtree.
OID tree	The Object Identifier (OID) value. OIDs uniquely identify MIB objects in the subtree.
Mask	The mask value and the mask type, along with the <i>oid-value</i> configured in the view command, determines the access of each sub-identifier of an object identifier (MIB subtree) in the view.
Permission	Included — Specifies to include MIB subtree objects.
	Excluded — Specifies to exclude MIB subtree objects.
No. of Views	The total number of configured views.
Group name	The access group name.

Sample Output

```
A:ALA-1# show system security view
=====
Views
=====
view name          oid tree          mask          permission
-----
iso                1                included
no-security        1                included
no-security        1.3.6.1.6.3       excluded
no-security        1.3.6.1.6.3.10.2.1 included
no-security        1.3.6.1.6.3.11.2.1 included
no-security        1.3.6.1.6.3.15.1.1 included
-----
No. of Views: 6
```



```
=====
A:ALA-1#

A:ALA-1# show system security view no-security detail
=====
Views
=====
view name          oid tree          mask          permission
-----
no-security        1                 included
no-security        1.3.6.1.6.3       excluded
no-security        1.3.6.1.6.3.10.2.1 included
no-security        1.3.6.1.6.3.11.2.1 included
no-security        1.3.6.1.6.3.15.1.1 included
-----
No. of Views: 5
=====
no-security used in
=====
group name
-----
snmp-ro
snmp-rw
=====
A:ALA-1#
```


4 NETCONF

4.1 In This Chapter

This chapter provides information to configure NETCONF.

Topics in this chapter include:


- [NETCONF Overview](#)
- [NETCONF in SR OS](#)
- [Establishing a NETCONF Session](#)
- [XML Content Layer](#)
- [XML Content Layer Examples](#)
- [CLI Content Layer](#)
- [CLI Content Layer Examples](#)

4.2 NETCONF Overview

NETCONF is a standardized IETF configuration management protocol published in RFC 6241. It is secure, connection-oriented, and runs on top of the SSHv2 transport protocol as specified in RFC 6242. NETCONF can be used as an alternative to CLI or SNMP for managing an SR OS.

NETCONF is an XML-based protocol used to configure network devices. It uses RPC messaging for communication between a NETCONF client and the NETCONF server running on the SR OS. An RPC message and configuration data is encapsulated within an XML document. These XML documents are exchanged between a NETCONF client and a NETCONF server in a request/response type of interaction. The SR OS NETCONF interface supports both configuration and retrieval of operational information. [Figure 15](#) shows a NETCONF RPC request.

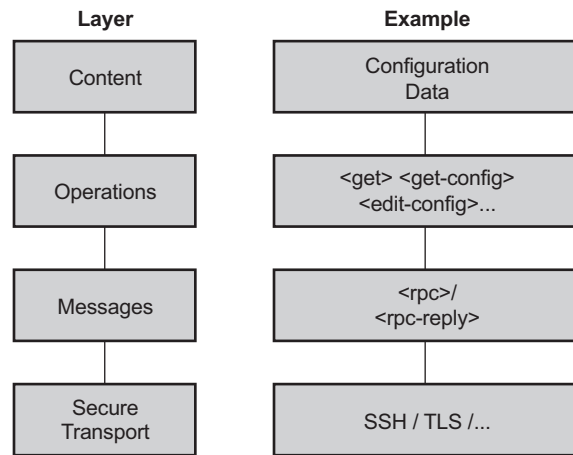
Figure 15 NETCONF RPC Request



al_0446

NETCONF can be conceptually partitioned into four layers as described in RFC 6241. Figure 16 shows the NETCONF layers.

Figure 16 NETCONF Layers (RFC 6241)



al_0447

4.3 NETCONF in SR OS

NETCONF can be used on an SR OS router to perform router management operations including:

- Changing the configuration of the router (<edit-config> operation)
- Reading the configuration of the router (<get-config> operation, equivalent to the **info** command in the CLI)
- Reading operational status and data (and associated configuration information) (<get> operation, equivalent to the **show** commands in the CLI)

NETCONF is not used for notifications on an SR OS router; for example, log events, syslog, or SNMP notifications (traps).

The equivalent of some admin commands are available via the SR OS NETCONF interface:

- **admin save** can be done using the <copy-config> operation
- **admin rollback** commands are supported using a CLI content layer <cli-action> RPC

The **bof**, **debug**, **tools**, **clear**, and other general CLI operational commands (for example, **telnet** or **ping**) are not supported via NETCONF on an SR OS.

The SR OS NETCONF server advertises base capability 1.1 (in addition to 1.0).

SR OS supports both a CLI content layer and an XML-based content layer for NETCONF.

4.3.1 YANG Data Models

The SR OS NETCONF XML content layer supports two similar proprietary configuration data models. Each configuration data model is described in a set of YANG modules. A unique set of XML namespaces is used for each of the two data models.

The YANG modules for the first configuration data model (Alcatel-Lucent Base-R13 SR OS YANG modules) have the following attributes:

- The names of the modules and sub-modules are **alu-conf-*-r13** (for example, **alu-conf-log-r13**). Note the **-r13** suffix at the end of the names.

- The Alcatel-Lucent Base-R13 models consist of a set of modules with groupings that are all used by a single top-level configuration module called `alu-conf-r13`. All configuration data in the Alcatel-Lucent Base-R13 models sits in the `urn:alcatel-lucent.com:sros:ns:yang:conf-r13` XML namespace.
- The modules cannot be used with the `<candidate>` datastore.
- Although the Base-R13 modules were first introduced in SR OS Release 13.0, they do not just contain objects from Release 13.0. For example, features from Release 14.0.R1 are also configurable using the versions of the Base-R13 modules that are distributed with SR OS Release 14.0.R1.

The YANG modules for the second configuration data model (Nokia SR OS YANG modules) have the following attributes:

- The names of the modules are **nokia-conf** (for example, `nokia-conf-log`). They have no `-r13` suffix in the names.
- The Nokia SR OS YANG models are divided into a single top-level configuration module (`nokia-conf`), a single top-level state module (`nokia-state`), a set of submodules (for example, `nokia-conf-system`), and a set of **nokia-types-*** modules. All configuration data in the Nokia SR OS YANG models sit in the `urn:nokia.com:sros:ns:yang:sr:conf` XML namespace. All state data in the Nokia SR OS YANG models sits in the `urn:nokia.com:sros:ns:yang:sr:state` XML namespace.
- The modules can be used with the `<candidate>` datastore.

The two configuration data models are not interchangeable. An XML request based on the Alcatel-Lucent Base-R13 YANG modules will not work if applied to a router using the `urn:nokia.com:sros:ns:yang:sr:conf` namespace (and vice versa).

All configuration modules and **types** modules are advertised in the SR OS NETCONF server `<hello>`. Submodules are not advertised in the `<hello>`.

The proprietary configuration YANG data models both closely align to the SR OS CLI configuration tree structure and commands.

The **bof**, **admin**, **tools**, **debug**, or **clear** branches of the CLI do not have equivalent YANG data models.

4.3.2 Transport and Sessions

SSH transport for NETCONF is supported on TCP port 830 with IPv4 or IPv6 in the Base routing instance. NETCONF SSH sessions (such as CLI, SCP and sFTP sessions) are subject to any configurable and non-configurable session limits; for example, inbound-max-sessions. Both the SSH server and NETCONF protocol must be enabled in the router configuration in order to use NETCONF. NETCONF sessions can be disconnected using the **admin disconnect** command. See the CLI section for details.

NETCONF sessions do not time out automatically and are not subject to the CLI session timeout. Operators can disconnect sessions manually if they need to.

A client establishing a NETCONF session must log into the router so user accounts must exist for NETCONF on the SR OS. An access type 'netconf' is provided. For access to the Base-R13 SR OS YANG data model, both **console** and **netconf** access must be configured for the user. For access to the Nokia SR OS YANG data model, only **netconf** access is necessary.

Only authentication via the local user database is supported for NETCONF users and sessions (no RADIUS or TACACS+ authentication).

Command authorization is not supported for the Nokia SR OS YANG data model. Once a NETCONF session is established and the user is authenticated then all configuration data is available via the Nokia SR OS YANG data model.

Command authorization is supported for the Alcatel-Lucent Base-R13 SR OS YANG modules. Also, access to various CLI config and show commands (via the CLI content layer) is controlled through the profile assigned to the user that is used to authenticate the underlying SSH session.

Access to LI commands using the Alcatel-Lucent Base-R13 SR OS YANG modules is based on the **access li** configuration setting for the user.

If a NETCONF request attempts to execute a CLI command which is outside the scope of its access profile, an error response will be sent.

Example - A user request, with **show** command, that is not in the scope of the user's access profile.

```
<?xml version="1.0" encoding="UTF-8"?>
<rpc message-id="101" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <get>
    <filter>
      <oper-data-format-cli-block>
        <cli-show>system security profile </cli-show>
      </oper-data-format-cli-block>
    </filter>
```

```

    </get>
  </rpc>
]]>]]>

```

Reply:

```

<?xml version="1.0" encoding="UTF-8"?>
<rpc-reply message-id="101" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <rpc-error>
    <error-type>application</error-type>
    <error-tag>operation-failed</error-tag>
    <error-severity>error</error-severity>
    <error-info>
      <err-element>cli-show</err-element>
    </error-info>
    <error-message>
      command failed - 'show system security profile'
      MINOR: CLI Command not allowed for this user.
    </error-message>
  </rpc-error>
</rpc-reply>
]]>]]>

```

4.3.3 Datastores and URLs

SR OS supports the <running> datastore, the <candidate> datastore, the <startup> datastore, and <url> tags.



Note: <url> is not a datastore in itself.

Support for the <candidate> datastore capability is advertised via the SR OS NETCONF server <hello> using the urn:ietf:params:netconf:capability:candidate:1.0 capability string.

All configuration changes (using <edit-config>) made to the <running> datastore via NETCONF take immediate operational effect. Configuration changes to the <candidate> datastore take effect after a successful <commit> operation.

The <startup> datastore and <url> tags can only be used with <copy-config> and <delete-config> and are not supported with any other operations (including <edit-config>, <get-config>, <get>, <validate>, etc).

The :startup capability is advertised in the SR OS NETCONF server <hello>:

```
<capability>urn:ietf:params:netconf:capability:startup:1.0</capability>
```


The <url> tag supports the same options as CLI <file-url>: local urls (CF) and remote urls (ftp and tftp).

The :url capability is advertised in the SR OS NETCONF server <hello>:

```
<capability>urn:ietf:params:netconf:capability:url:1.0?scheme=ftp,tftp,file</capability>
```

The following examples show the format of each URL scheme:

- <target><url>ftp://name:passwd@a.b.c.d/usr/myfiles/myfile.cfg</url></target>
- <target><url>tftp://name:passwd@a.b.c.d/usr/myfiles/myfile.cfg</url></target>
- <target><url>file:///cf3:/myfiles/myfile.cfg</url></target>
- <target><url>cf3:/myfiles/myfile.cfg</url></target>



Note: The examples use “///” for the file URL. Also, the file://localhost/... format is not supported.

The <startup> datastore is identified by following the bof primary-config/secondary-config/tertiary-config paths as configured by the operator. The <startup> datastore is effectively an alias for a URL (a special URL used for system startup) with some extra resiliency (primary/secondary/tertiary).

The BOF is not considered part of any configuration datastore.

Debug configuration (such as debug mirrors, or anything saved with **admin debug-save**) is not considered part of any configuration datastore.

Lawful Interception configuration information is contained in the <running> datastore but is not saved in the <startup> datastore. The equivalent of the CLI **li save** command is available in an <edit-config> using the Alcatel-Lucent Base-R13 SR OS YANG modules.

Configuration changes done via NETCONF are subject to CLI rollback (**revert**, **save**, and so on) and are included in the configuration when the operator performs an **admin save** in the CLI.

Only the data model described by Nokia SR OS YANG modules can be used with the <candidate> datastore. The data model described by the Alcatel-Lucent Base-R13 SR OS YANG modules is not applicable to the <candidate> datastore but does work with the <running> datastore. All <edit-config> requests to the candidate datastore must use the urn:nokia.com:sros:ns:yang:sr:conf namespace.

The candidate datastore supports the XML content layer only. Requests/replies to/from the candidate datastore cannot contain the CLI content layer.

4.3.4 NETCONF Operations and Capabilities

The following base protocol operations are supported:

- `<get>`
- `<get-config>`
- `<edit-config>`
- `<copy-config>` and `<delete-config>`
- `<lock>`
- `<unlock>`
- `<close-session>`
- `<kill-session>`

The following optional capabilities from RFC 6241 are supported:

- Writable-Running Capability
- Candidate Configuration Capability
 - `<commit>` operation
 - `<discard-changes>` operation
- Validate Capability
 - `<validate>` operation
- Distinct Startup Capability
- URL Capability

The following capability from RFC 6243 is supported:

- With-defaults Capability

The `<edit-config>` operation's `<error-option>` is not supported. SR OS implements the stop-on-error behavior by default. The continue-on-error and rollback-on-error are not supported.

One rpc request can only contain one operation.

[Table 53](#) shows supported NETCONF operations.

Table 53 NETCONF Operations

Operation	Arguments
get-config	source [filter]
edit-config	target [default-operation][test-option][error-option] config
copy-config	target source
delete-config	target
lock	target
unlock	target
get	[filter]
close-session	n/a
kill-session	session-id
discard-changes	n/a
validate	source
commit	n/a



Note: Bracketed arguments are optional.

4.3.4.1 <get>

The CLI content layer <get> operation is supported with both configuration and state data returned in a <get> reply. An XML content layer <get> operation, supported with both configurations and state data, being returned in a <get> reply as per the NOKIA SR OS YANG data model only.

A <get> request is first analyzed for syntax errors before any execution starts. If a syntax error is found then a single global <rpc-error> for the entire request is sent in the reply.

Responses are provided for each item in the request until the first item with an error is found. The item with an error has a <response> tag containing some error information, followed by an <rpc-error> tag (and sub-tags). The reply is then returned and subsequent items are not executed.

The <rpc-error> for an individual item (i.e. for a non-syntax error) is after the </response> information and not inside the <response>.

Example — <get> request with a non-syntax error in the 2nd item:

```
<?xml version="1.0" encoding="UTF-8"?>
<rpc message-id="101" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <get>
    <filter>
      <oper-data-format-cli-block>
        <cli-show>router interface "system"</cli-show>
        <cli-show>router mpls lsp</cli-show>
        <cli-show>system security ssh</cli-show>
      </oper-data-format-cli-block>
    </filter>
  </get>
</rpc>
]]>]]>
```

Reply:

```
<?xml version="1.0" encoding="UTF-8"?>
<rpc-reply message-id="101" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <data xmlns="urn:alcatel-lucent.com:sros:ns:yang:cli-content-layer-r13">
    <oper-data-format-cli-block>
      <item>
        <cli-show>router interface "system"</cli-show>
        <response>
```

```
=====
Interface Table (Router: Base)
=====
Interface-Name      Adm      Opr (v4/v6)  Mode      Port/SapId
  IP-Address                               PfxState
-----
system              Up        Up/Down      Network   system
  144.23.63.5/32                               n/a
-----
Interfaces : 1
=====
```

```
      </response>
    </item>
    <item>
      <cli-show>router mpls lsp</cli-show>
      <response>
        MINOR: CLI MPLS is not configured.
      </response>
    <rpc-error>
      <error-type>application</error-type>
      <error-tag>operation-failed</error-tag>
      <error-severity>error</error-severity>
      <error-info>
        <err-element>cli-show</err-element>
      </error-info>
      <error-message>
        command failed - 'show router mpls lsp'
      </error-message>
```

```

        </rpc-error>
      </item>
    </oper-data-format-cli-block>
  </data>
</rpc-reply>
]]>]]>

```

4.3.4.2 <get-config>

The <get-config> operation returns non-default configuration by default for the Alcatel-Lucent Base-R13 SR OS YANG modules (the 'trim' mode as per RFC 6243).

The <get-config> operation returns data nodes that were set by a client to their default values for the NOKIA SR OS modules (the 'explicit' mode as per RFC 6243).

4.3.4.3 <edit-config>

The following values for the <test-option> parameter under <edit-config> are supported:

- test-then-set
- set
- test-only

4.3.4.4 <copy-config> and <delete-config>

The <copy-config> and <delete-config> base protocol operations are supported for specific combinations of source and target datastores.

The <copy-config> operation is supported for the following combinations of sources and targets:

- <source>=<url> and <target>=<startup> (as long as both are not remote urls)
- <source>=<startup> and <target>=<url> (as long as both are not remote urls)
- <source>=<running> and <target>=<url>
 - Equivalent of “admin save <file-url>”
 - An index file is also saved if “persist on” is configured in the bof
- <source>=<running> and <target>=<startup>
 - Equivalent of “admin save”

- An index file is also saved if “persist on” is configured in the bof

The <running> datastore cannot be a <target> for a <copy-config>.

The <candidate> datastore cannot be a <target> or a <source> for a <copy-config>.

Remote URL to remote URL copies are not supported. For example, if primary-image is a remote URL, then a <startup> to copy will fail with an error.

The <copy-config> operation uses the CLI Content Layer format. The format of the source and target is block CLI.

The <delete-config> operation is supported for the following targets:

- <url>
- <startup>

The <delete-config> operation is not allowed on the <running> or <candidate> datastores.

4.3.4.5 <lock>

Taking the <candidate> datastore’s lock is equivalent to doing a CLI exclusive transaction.

Although the NETCONF protocol allows specifying a target datastore for a lock operation, the SR OS only implements a single lock:

- taking the running datastore’s lock locks both the running and candidate datastores (creating a single lock)
- taking the candidate datastore’s lock locks both the running and candidate datastores (creating a single lock)

When either the running datastore’s lock or the candidate datastore’s lock is taken by a NETCONF session:

- no NETCONF session can take the <running> datastore lock
- no NETCONF session can take the <candidate> datastore lock
- no other NETCONF session can do an <edit-config> on the running datastore
- no other NETCONF session can do an <edit-config> on the candidate datastore
- no other NETCONF session can do a <commit> on the candidate datastore
- no other NETCONF session can do a <discard-changes> on the candidate datastore

- the CLI becomes read-only
- **rollback revert** is blocked
- SNMP set requests fail on objects that are part of the urn:nokia.com:sros:ns:yang:sr:conf-* namespace

A datastore's lock is unlocked when disconnecting a NETCONF session (either from the CLI using Ctrl-c, or by performing a <kill-session> or <close-session> operation). Upon disconnecting a NETCONF session that had acquired a datastore's lock, SR OS:

- releases the lock
- discards the "uncommitted" changes (if any)



Note: The behavior is different if the disconnected NETCONF session had the "implicit" lock (see the [<edit-config> with XML Content Layer](#) section). In that case, the SR OS keeps the "uncommitted" changes in the <candidate> datastore.

Timeouts of locks are not supported. No specific admin/tools commands are provided to release the lock, but the session that holds the lock can be administratively disconnected using the CLI to release the lock.

From the CLI, the operator can configure whether users that belong to a specific profile have permission to lock NETCONF sessions; see the [NETCONF Configuration Command Reference](#).

Using CLI **show** commands, the operator can determine if either the <running> datastore's lock or the <candidate> datastore's lock is currently taken and which session has the lock; see the [NETCONF Show Command Reference](#).

4.3.4.6 <unlock>

Because there is a single lock per datastore regardless of what the scope of that lock is, the following applies.

- The <running> datastore's lock is unlocked by using the <unlock> command only on the <running> datastore. An error results and the lock stays if a different datastore is used with the <unlock> operation.
- The <candidate> datastore's lock is unlocked by using the <unlock> command only on the <candidate> datastore. An error results and the lock stays if a different datastore is used with the <unlock> operation.

Performing an <unlock> operation on the candidate datastore discards all pending (not committed) candidate datastore changes.

4.3.4.7 <commit>

The <commit> command has the following characteristics:

- It represents the equivalent of the CLI command **candidate commit**.
- When a <commit> operation fails, only the first error is returned.
- When the SR OS cannot commit all the changes in the candidate datastore, the SR OS keeps the <running> datastore unchanged; that is, no partial commit takes place.
- When a NETCONF session is disconnected (using Ctrl-c or <kill-session>) in the middle of a <commit> operation, SR OS keeps the running datastore unchanged.
- The persistency of changes made via a <commit> operation is operator-controlled. A copy of the running datastore to the startup datastore is not automatically performed after each <commit> operation.
- When some changes exist in the candidate datastore (prior to being committed to the running datastore), there are some impacts to:
 - a CLI user trying to make some immediate changes, as the SR OS blocks all CLI immediate configurations
 - an SNMP set request, as SR OS blocks it and returns an error
 - an <edit-config> to the running datastore, as SR OS blocks all <edit-config> requests to the running datastore and returns an error

4.3.4.8 <discard-changes>

The <discard-changes> operation causes the <candidate> datastore to revert back to match the <running> datastore and releases the “implicit” lock. From the CLI, the operator can do the equivalent of a <discard-changes> operation which releases the implicit lock as well (see [4.11](#)).

4.3.4.9 <validate>

The validate capability is supported in the following ways:

- The validate:1.1 and 1.0 capabilities are advertised in the NETCONF server's <hello>:
 - <capability>urn:ietf:params:netconf:capability:validate:1.0</capability>
 - <capability>urn:ietf:params:netconf:capability:validate:1.1</capability>
- The <validate> operation is supported for an XML content layer request but not for a CLI content layer request. Detection of a <config-format-cli-block> or <oper-data-format-cli-block> tag in a <validate> request will result in an “operation not supported” error response.
- A <validate> operation is supported for a selection of config (<source><config>) for both the <candidate> datastore and the <running> datastore, which only returns 'OK'. The <validate> request is not supported for URL sources or the <startup> datastore.
- A <validate> operation checks mainly the syntax. Only the first error is returned.

4.3.5 Data Model, Datastore and Operation Combinations

Table 54 shows the which operations are supported by data model and datastore combination.

Table 54 Data Model, Datastore and Operation Combinations

Operation	R13 Modules		Nokia Modules	
	<running>	<candidate>	<running>	<candidate>
<edit-config>	supported	not supported	not supported	supported
<get-config>	supported	not supported	supported	supported
<get>*	retrieves CLI content layer state data (no XML content layer)		retrieves configuration and state data (XML format only)	

* - Note that datastore is not applicable for a <get> operation

4.3.6 General NETCONF Behavior

Pressing Ctrl-c in a NETCONF session will immediately terminate the session.

The SR OS NETCONF implementation does support XML namespaces (xmlns).

If an invalid namespace is specified within the client's hello message, no error will be returned as the NETCONF server is still waiting for the client to send a valid <hello/>. Further NETCONF requests (without sending a proper hello message) even though correct, SR OS returns an error in that case mentioning "Common base capability not found."

In the <rpc> element, the allowed XML namespaces are:

- the standard NETCONF "urn:ietf:params:xml:ns:netconf:base:1.0" namespace
- the SR OS "urn:alcatel-lucent.com:sros:ns:yang:conf-r13" namespace
- the SR OS "urn:nokia.com:sros:ns:yang:sr:conf" namespace

In the <rpc> element, prefixes are accepted and have to be specified with a valid URI. If an incorrect URI is declared with a prefix, then SR OS detects the invalid URI and sends an <rpc-error> response.

If any other XML namespace is declared (or assigned to a prefix) in the RPC tag, then the SR OS returns an error.

Any prefix declarations in the rest of the request are ignored and unused. The SR OS NETCONF server puts the correct NETCONF namespace declaration ("urn:ietf:params:xml:ns:netconf:base:1.0") in all replies.

An <edit-config> request must specify which data model (Alcatel-Lucent Base-r13 or Nokia SR OS) is being used in the top level <configure> element.

- The SR OS accepts a request with only a single namespace at the top <configure> element. For example:

```
<configure xmlns="urn:nokia.com:sros:ns:yang:sr:conf">
  <system>
    ....
```

Or:

```
<configure xmlns="urn:alcatel-lucent.com:sros:ns:yang:conf-r13">
  <system>
    ....
```

- The NETCONF client can declare those two namespaces with prefixes at the <rpc> tag itself and use the corresponding prefixes later in the request's <configure/> block.
- The SR OS returns an error if the request contains one or more incorrect namespaces.

Example 1 — the standard NETCONF namespace

"urn:ietf:params:xml:ns:netconf:base:1.0" is used more than once in the <rpc> element:

```
<?xml version="1.0" encoding="UTF-8"?>
<rpc message-id="101"
xmlns="urn:ietf:params:xml:ns:netconf:base:1.0"
xmlns:alu="urn:ietf:params:xml:ns:netconf:base:1.0">
<get-config>
<source> <running/> </source>
<filter>
  <configure xmlns="urn:alcatel-lucent.com:sros:ns:yang:conf-r13">
    <router>
      <interface>
        <interface-name>"system"</interface-name>
      </interface>
    </router>
  </configure>
</filter>
</get-config>
</rpc>
]]>]]>
```

Reply (no error message):

```
<?xml version="1.0" encoding="UTF-8"?>
<rpc-reply message-id="101" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0"
xmlns:alu="urn:ietf:params:xml:ns:netconf:base:1.0">
  <data>
    <configure xmlns="urn:alcatel-lucent.com:sros:ns:yang:conf-r13">
      <router>
        <router-instance>Base</router-instance>
        <interface>
          <interface-name>system</interface-name>
          <shutdown>>false</shutdown>
        </interface>
      </router>
    </configure>
  </data>
</rpc-reply>
]]>]]>
```

Example 2 — an allowed non-default NETCONF base namespace is used in the <rpc> element:

```
<?xml version="1.0" encoding="UTF-8"?>
<rpc message-id="101"
xmlns="urn:ietf:params:xml:ns:netconf:base:1.0"
xmlns:alu="urn:alcatel-lucent.com:sros:ns:yang:conf-r13">
<get-config>
<source> <running/> </source>
<filter>
  <configure xmlns="urn:alcatel-lucent.com:sros:ns:yang:conf-r13">
    <router>
      <interface>
        <interface-name>"system"</interface-name>
      </interface>
    </router>
  </configure>
</filter>
</get-config>
```

```
</rpc>
]]>]]>
```

Reply (non-NETCONF base namespace is allowed and no error is returned):

```
<?xml version="1.0" encoding="UTF-8"?>
<rpc-reply message-id="101" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0"
  xmlns:alu="urn:alcatel-lucent.com:sros:ns:yang:conf-r13">
  <data>
    <configure xmlns="urn:alcatel-lucent.com:sros:ns:yang:conf-r13">
      <router>
        <router-instance>Base</router-instance>
        <interface>
          <interface-name>system</interface-name>
          <shutdown>>false</shutdown>
        </interface>
      </router>
    </configure>
  </data>
</rpc-reply>
]]>]]>
```

Example 3 — an invalid NETCONF namespace is declared in the <rpc> element:

```
<?xml version="1.0" encoding="UTF-8"?>
<rpc message-id="101"
  xmlns="urn:ietf:params:xml:ns:netconf:base:1.0"
  xmlns:alu="urn:alcatel-lucent.com:sros:ns:yang:sr:conf">
  <get-config>
    <source><running/></source>
    <filter>
      <configure xmlns="urn:alcatel-lucent.com:sros:ns:yang:conf-r13">
        <router>
          <interface>
            <interface-name>"system"</interface-name>
          </interface>
        </router>
      </configure>
    </filter>
  </get-config>
</rpc>
]]>]]>
```

Reply (the SR OS returns an error):

```
<?xml version="1.0" encoding="UTF-8"?>
<rpc-reply message-id="101" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0"
  xmlns:alu="urn:alcatel-lucent.com:sros:ns:yang:sr:conf">
  <rpc-error>
    <error-type>protocol</error-type>
    <error-tag>unknown-element</error-tag>
    <error-severity>error</error-severity>
    <error-info>
      <bad-element>rpc</bad-element>
      <bad-namespace>urn:alcatel-lucent.com:sros:ns:yang:sr:conf</bad-namespace>
    </error-info>
  </rpc-error>
</rpc-reply>
]]>]]>
```

```
        <error-message>
            An unexpected namespace is present.
        </error-message>
    </rpc-error>
</rpc-reply>
]]>]]>
```

Example 4 — a non-default NETCONF namespace/prefix declared in any child tag overrides the one declared under rpc tag:

```
<?xml version="1.0" encoding="UTF-8"?>
<rpc message-id="101"
xmlns="urn:ietf:params:xml:ns:netconf:base:1.0"
xmlns:alu="urn:ietf:params:xml:ns:netconf:base:1.0">
<get-config>
<source> <running/> </source>
<filter>
    <configure xmlns="urn:alcatel-lucent.com:sros:ns:yang:conf-r13">
        <router>
            <interface xmlns:alu="urn:alcatel-lucent.com:sros:ns:yang:conf-r13">
                <alu:interface-name>"system"</alu:interface-name>
            </interface>
        </router>
    </configure>
</filter>
</get-config>
</rpc>
]]>]]>
```

Reply (non-standard namespace/prefix used in tag is ignored):

```
<?xml version="1.0" encoding="UTF-8"?>
<rpc-reply message-id="101" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0"
xmlns:alu="urn:ietf:params:xml:ns:netconf:base:1.0">
    <data>
        <configure xmlns="urn:alcatel-lucent.com:sros:ns:yang:conf-r13">
            <router>
                <router-instance>Base</router-instance>
                <interface>
                    <interface-name>system</interface-name>
                    <shutdown>>false</shutdown>
                </interface>
            </router>
        </configure>
    </data>
</rpc-reply>
]]>]]>
```

The chunked framing mechanism is supported (in addition to the EOM mechanism). As per RFC 6242, Section 4.1 - Framing Protocol, “[...] If the :base:1.1 capability is advertised by both peers, the chunked framing mechanism (see Section 4.2) is used for the remainder of the NETCONF session. Otherwise, the end-of-message-based mechanism (see Section 4.3) is used.”

Example 5 — Chunked message:

```
#340
<?xml version="1.0" encoding="UTF-8"?><rpc message-id="101"
xmlns="urn:ietf:params:xml:ns:netconf:base:1.0"><get-config><source><running/>
</source>
<filter><configure xmlns="urn:alcatel-lucent.com:sros:ns:yang:conf-r13">
<router><interface>
<interface-name>system</interface-name></interface></router></configure></filter>
</get-config></rpc>
##
```

Example 6 — Chunked message:

```
#38
<?xml version="1.0" encoding="UTF-8"?>
#83
<rpc message-id="101" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
<get-config>
#101
<source><running/></source>
<filter>
<configure xmlns="urn:alcatel-lucent.com:sros:ns:yang:conf-r13">
##39
<system>
<netconf>
</netconf>
</system>
##43
</configure>
</filter>
</get-config>
</rpc>
##
```

Handling of default data (for example, 'info' vs 'info detail') uses the mechanisms detailed in RFC 6243. The SR OS NETCONF server supports the 'trim' method as the default for the Alcatel-Lucent Base-R13 SR OS YANG modules. It supports the 'explicit' method as the default for the NOKIA SROS Yang modules and also supports the 'report-all' method and advertises that in the <hello>:

```
<capability>urn:ietf:params:netconf:capability:with-defaults:1.0?basic-
mode=trim&also-supported=explicit,report-all</capability>
```

A user can save a rollback checkpoint (for example, prior to doing an <edit-config> or a series of <edit-config>) and perform a rollback revert if needed later using the <cli-action> RPC.

The set of supported actions are as follows:

- admin>rollback compare [to checkpoint2]
- admin>rollback compare checkpoint1 to checkpoint2
- admin>rollback delete checkpoint | rescue
- admin>rollback save [comment comment] [rescue]

- admin>rollback revert checkpoint | rescue [now]
- admin>rollback view [checkpoint | rescue]

Example 7 — Two rollback items with responses:

```
<?xml version="1.0" encoding="UTF-8"?>
<rpc message-id="102" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <cli-action>
    <admin>rollback compare active-cfg to 1</admin>
    <admin>rollback compare</admin>
  </cli-action>
</rpc>
]]>]]>
```

Reply:

```
<?xml version="1.0" encoding="UTF-8"?>
<rpc-reply message-id="102" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <data xmlns="urn:alcatel-lucent.com:sros:ns:yang:cli-content-layer-r13">
    <cli-action>
      <item>
        <admin>rollback compare active-cfg to 1</admin>
        <response>
          0.150 s
          0.450 s
        </response>
      </item>
    </cli-action>
  </data>
</rpc-reply>
```

```
-----
configure
router
-   mpls
-       shutdown
-       interface "system"
-           no shutdown
-       exit
-       lsp "test"
-           shutdown
-       exit
-   exit
-   rsvp
-       shutdown
-       interface "system"
-           no shutdown
-       exit
-   exit
exit
exit
```

```
-----
Finished in 0.720 s
      </response>
    </item>
    <item>
      <admin>rollback compare</admin>
      <response>
        0.160 s
        0.070 s
      </response>
    </item>
  </cli-action>
</data>
</rpc-reply>
```

```
-----
configure
router
```

```

-      mpls
-        shutdown
-        interface "system"
-          no shutdown
-        exit
-        lsp "test"
-          shutdown
-        exit
-      exit
-    rsvp
-      shutdown
-      interface "system"
-        no shutdown
-      exit
-    exit
  exit
service
-  vpls "99" customer 1 create
-    shutdown
-    stp
-      shutdown
-    exit
-  exit
  exit
exit
-----
Finished in 0.350 s
      </response>
    </item>
  </cli-action>
</data>
</rpc-reply>
]]>]]>

```

Example 8 — Syntax error in the request resulting in global rpc-error reply:

```

<?xml version="1.0" encoding="UTF-8"?>
<rpc message-id="103"
  xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <cli-action>
    <admin>rollback compare active-cfg to 1</admin>
    <admin>rollback compare flee-fly</admin>
  </cli-action>
</rpc>
]]>]]>

```

Reply:

```

<?xml version="1.0" encoding="UTF-8"?>
<rpc-reply message-id="103" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <rpc-error>
    <error-type>application</error-type>
    <error-tag>operation-failed</error-tag>
    <error-severity>error</error-severity>
    <error-info>
      <err-element>admin</err-element>
    </error-info>
  </rpc-error>
</rpc-reply>

```



```

        </error-info>
        <error-message>
            command failed - '/admin rollback compare flee-fly'
        </error-message>
    </rpc-error>
</rpc-reply>
]]>]]>

```

Example 9 — Error processing the request:

```

<?xml version="1.0" encoding="UTF-8"?>
<rpc message-id="103"
  xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <cli-action>
    <admin>rollback compare active-cfg to 1</admin>
    <admin>rollback compare 1 to flee-fly</admin>
  </cli-action>
</rpc>
]]>]]>

```

Reply:

```

<?xml version="1.0" encoding="UTF-8"?>
<rpc-reply message-id="103" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <data xmlns="urn:alcatel-lucent.com:sros:ns:yang:cli-content-layer-r13">
    <cli-action>
      <item>
        <admin>rollback compare active-cfg to 1</admin>
        <response>
          0.160 s
          0.180 s

```

```

-----
configure
router
-   mpls
-       shutdown
-       interface "system"
-           no shutdown
-       exit
-   exit
-   rsvp
-       shutdown
-       interface "system"
-           no shutdown
-       exit
-   exit
exit
exit
-----

```

```

Finished in 0.460 s
    </response>
  </item>
  <item>
    <admin>rollback compare 1 to flee-fly</admin>
    <response>
    </response>
  </rpc-error>

```

```

        <error-type>application</error-type>
        <error-tag>operation-failed</error-tag>
        <error-severity>error</error-severity>
        <error-info>
          <err-element>admin</err-element>
        </error-info>
        <error-message>
          command failed - '/admin rollback compare 1 to flee-fly'
          MINOR: CLI No such file ('flee-fly').
        </error-message>
      </rpc-error>
    </item>
  </cli-action>
</data>
</rpc-reply>
]]>]]>

```

Example 10 — Error in the 2nd item of the request, resulting in no 3rd item in the reply:

```

<?xml version="1.0" encoding="UTF-8"?>
<rpc message-id="104" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <cli-action>
    <admin>rollback compare active-cfg to 1</admin>
    <admin>rollback compare 1 to xyz</admin>
    <admin>rollback compare active-cfg to 1</admin>
  </cli-action>
</rpc>
]]>]]>

```

Reply:

```

<?xml version="1.0" encoding="UTF-8"?>
<rpc-reply message-id="104" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <data xmlns="urn:alcatel-lucent.com:sros:ns:yang:cli-content-layer-r13">
    <cli-action>
      <item>
        <admin>rollback compare active-cfg to 1</admin>
        <response>
          0.170 s
          1.350 s

```

```

-----
configure
router
-   mpls
-     shutdown
-     interface "system"
-       no shutdown
-     exit
-   exit
-   rsvp
-     shutdown
-     interface "system"
-       no shutdown
-     exit
-   exit

```

```

    exit
  exit
-----
Finished in 1.640 s
    </response>
  </item>
  <item>
    <admin>rollback compare 1 to xyz</admin>
    <response>
    </response>
    <rpc-error>
      <error-type>application</error-type>
      <error-tag>operation-failed</error-tag>
      <error-severity>error</error-severity>
      <error-info>
        <err-element>admin</err-element>
      </error-info>
      <error-message>
        command failed - '/admin rollback compare 1 to xyz'
        MINOR: CLI No such file ('xyz').
      </error-message>
    </rpc-error>
  </item>
</cli-action>
</data>
</rpc-reply>
]]>]]>

```

4.3.6.1 System-Provisioned Configuration (SPC) Objects

There is a set of configuration objects that are provisioned (added to the <running> datastore) automatically by SR OS; for example, log-id 99.

Some of these objects can be deleted/removed by a user (Deletable SPC Objects).

- In the CLI these are removed by specifying the keyword **no**, which is then visible in an **info** command or in a saved config (**admin save**); for example, **no log-id 99**.
- The Deletable SPC Objects can be removed or recreated via NETCONF <edit-config> requests, but they are not visible in a <get-config> response in the “urn:alcatel-lucent.com:sros:ns:yang:conf-*-r13” namespace (the Alcatel-Lucent Base-R13 SR OS YANG modules) when they are:
 - set to their default values (including all child leaves and objects)
 - removed or deleted
- The Deletable SPC Objects are visible in a <get-config> response in the “urn:alcatel-lucent.com:sros:ns:yang:conf-*-r13” namespace (the Alcatel-Lucent Base-R13 SR OS YANG modules) if a child leaf or object is changed away from the default value; for example, changing log-99 to time-format local.

- The Deletable SPC objects are not visible in a <get-config> response in the “urn:nokia.com:sros:ns:yang:sr:conf” namespace (the Nokia SR OS YANG modules) if the child leaves are all at default values.
- The list of Deletable SPC Objects is as follows:

```
Config system security profile default
Config system security profile default entry 10-100
Config system security profile administrative
Config system security profile administrative entry 10-112
Config system security user "admin"
Config system security user console member "default"
Config system security snmp access group xyz (a set of access groups)
Config system security ssh client-cipher-list protocol-version 1 cipher 200-210
Config system security ssh client-cipher-list protocol-version 2 cipher 190-235
Config system security ssh server-cipher-list protocol-version 1 cipher 200-205
Config system security ssh server-cipher-list protocol-version 2 cipher 190-235
Config log filter 1001
Config log filter 1001 entry 10
Config log log-id 99 & 100
```

Some SPC objects cannot be deleted (Non-Deletable SPC Objects).

- Although these objects cannot be deleted, some of them contain leaves that can be modified.
- The Non-Deletable SPC Objects are not visible in a <get-config> response in the “urn:alcatel-lucent.com:sros:ns:yang:conf-*-r13” namespace (the Alcatel-Lucent Base-R13 SR OS YANG modules) when they are set to their default values (including all child leaves and objects).
- The Non-Deletable SPC Objects are visible in a <get-config> response in the “urn:alcatel-lucent.com:sros:ns:yang:conf-*-r13” namespace (the Alcatel-Lucent Base-R13 SR OS YANG modules) if a child leaf or object is changed away from the default value; for example, setting the card-type.
- The Non-Deletable SPC objects are not visible in a <get-config> response in the “urn:nokia.com:sros:ns:yang:sr:conf” namespace (the Nokia SR OS YANG modules) if the child leaves are all at default values.
- The list of Non-Deletable SPC Objects is as follows:

```
Config system security user-template {tacplus_default|radius_default}
Config system security snmp view iso ...
Config system security snmp view li-view ...
Config system security snmp view mgmt-view ...
Config system security snmp view vprn-view ...
Config system security snmp view no-security-view ...
Config log event-control ...
Config filter log 101
Config qos ... various default policies can't be deleted
Config qos queue-group-templates ... these can't be deleted
Config card <x>
Config router network-domains network-domain "default"
Config oam-pm bin-group 1
Config call-trace trace-profile "default"
```

Some Non-Deletable SPC Objects are visible in a <get-config> request in the “urn:alcatel-lucent.com:sros:ns:yang:conf-*-r13” namespace (the Alcatel-Lucent Base-R13 SR OS YANG modules), even if they are set to default values:

```
Config system security cpu-protection policy 254 and 255
Config router interface "system"
Config service customer 1
```

4.4 Establishing a NETCONF Session

The following example shows a client on a Linux PC initiating a connection to an SR OS NETCONF server. The SSH session must be invoked using an SSH subsystem (as recommended in RFC 6242):

```
ssh -s my_username@a.b.c.d -p 830 netconf
```

The following example shows an exchange of hello messages which include advertisement of capabilities.

From the SR OS server:

```
<?xml version="1.0" encoding="UTF-8"?>
<hello xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <capabilities>
    <capability>urn:ietf:params:netconf:base:1.0</capability>
    <capability>urn:ietf:params:netconf:base:1.1</capability>
    <capability>urn:ietf:params:netconf:capability:writable-running:1.0
    </capability>
    <capability>urn:ietf:params:netconf:capability:candidate:1.0</capability>
    <capability>urn:ietf:params:netconf:capability:validate:1.0</capability>
    <capability>urn:ietf:params:netconf:capability:validate:1.1</capability>
    <capability>urn:ietf:params:netconf:capability:startup:1.0</capability>
    <capability>urn:ietf:params:netconf:capability:url:1.0?scheme=ftp,tftp,file
    </capability>
    <capability>urn:ietf:params:netconf:capability:with-defaults:1.0?basic-
    mode=trim&also-supported=explicit,report-all</capability>
    <capability>urn:ietf:params:xml:ns:netconf:base:1.0?module=ietf-
    netconf&revision=2011-06-01&features=writable-
    running,validate,startup,url&deviations=alu-netconf-deviations-r13</
    capability>
    <capability>urn:alcatel-lucent.com:sros:ns:yang:netconf-deviations-r13?
    module=alu-netconf-deviations-r13&revision=2015-01-23</capability>
    <capability>urn:alcatel-lucent.com:sros:ns:yang:cli-content-layer-r13?
    module=alu-cli-content-layer-r13&revision=2015-01-23</capability>
    <capability>urn:alcatel-lucent.com:sros:ns:yang:conf-r13?module=
    alu-conf-r13&revision=2017-01-04</capability>
    <capability>urn:alcatel-lucent.com:sros:ns:yang:conf-aaa-r13?module=
    alu-conf-aaa-r13&revision=2016-12-01</capability>
    <capability>urn:alcatel-lucent.com:sros:ns:yang:conf-aa-r13?module=
    alu-conf-aa-r13&revision=2016-12-21</capability>
    <capability>urn:nokia.com:sros:ns:yang:sr:conf?module=
    nokia-conf&revision=2016-07-06</capability>
    ...
    <capability>urn:nokia.com:sros:ns:yang:sr:sros-yang-extensions?module=
    nokia-sros-yang-extensions&revision=2016-01-01</capability>
    <capability>urn:nokia.com:sros:ns:yang:sr:state?module=
    nokia-state&revision=2016-07-06</capability>
    ...
    <capability>urn:nokia.com:sros:ns:yang:sr:types-services?module=nokia-
    types-services&revision=2016-12-27</capability>
    <capability>urn:nokia.com:sros:ns:yang:sr:types-sros?module=nokia-types-
    sros&revision=2017-01-10</capability>
    <capability>urn:nokia.com:sros:ns:yang:sr:types-system?module=nokia-types-
```

```
        system&revision=2017-01-05</capability>
      <capability>urn:nokia.com:sros:ns:yang:sr:major-release-0</capability>
    </capabilities>
    <session-id>69</session-id>
  </hello>
]]>]]>
```

A NETCONF client can reply with a hello message like the following:

```
<?xml version="1.0" encoding="UTF-8"?>
<hello xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <capabilities>
    <capability>urn:ietf:params:netconf:base:1.0</capability>
  </capabilities>
</hello>
]]>]]>
```

4.5 XML Content Layer

XML is the default content layer format for the SR OS NETCONF server. When using the XML format at the NETCONF content layer, configuration changes and configuration information retrieved are expressed as XML tags.

4.5.1 <get> with XML Content Layer

A <get> operation with an XML content layer is supported with the <candidate> datastore only. A <get> request retrieves both the configuration and state data from the “urn:nokia.com:sros:ns:yang:sr:conf” namespace (the Nokia SR OS YANG modules) only. If any nodes from the configure tree are included in a <get> request filter, then at minimum the <configure> tag must contain a namespace. If the namespace is not specified, the SR OS returns an error.

Example 1: The <configure> tag contains a namespace

```
<?xml version="1.0" encoding="UTF-8"?>
<rpc message-id="101" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <get>
    <filter>
      <configure xmlns="urn:nokia.com:sros:ns:yang:sr:conf">
        <python/>
      </configure>
    </filter>
  </get>
</rpc>
]]>]]>
```

Reply: no errors

```
<?xml version="1.0" encoding="UTF-8"?>
<rpc-reply message-id="101" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <data>
    <configure xmlns="urn:nokia.com:sros:ns:yang:sr:conf">
      <python xmlns="urn:nokia.com:sros:ns:yang:sr:conf-python">
        <python-script>
          <script-name>testing</script-name>
          <shutdown>>false</shutdown>
          <protection>
            </protection>
        </python-script>
        <python-script>
          <script-name>tested</script-name>
          <protection>
            </protection>
        </python-script>
      </python>
    </configure>
```



```

    </data>
  </rpc-reply>
]]>]]>

```

Example 2: The <configure> tag does not contain a namespace

```

<?xml version="1.0" encoding="UTF-8"?>
<rpc message-id="101" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <get>
    <filter>
      <configure>
        <python xmlns="urn:nokia.com:sros:ns:yang:sr:conf-python">
          </python>
        </configure>
      </filter>
    </get>
  </rpc>
]]>]]>

```

Reply: SR OS errors

```

<?xml version="1.0" encoding="UTF-8"?>
<rpc-reply message-id="101" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <rpc-error>
    <error-type>protocol</error-type>
    <error-tag>unknown-element</error-tag>
    <error-severity>error</error-severity>
    <error-info>
      <bad-element>configure</bad-element>
    </error-info>
    <error-message>
      Element is not valid in the specified context.
    </error-message>
  </rpc-error>
</rpc-reply>
]]>]]>

```

Example 3: The <state> tag contains a namespace

```

<?xml version="1.0" encoding="UTF-8"?>
<rpc message-id="101" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <get>
    <filter>
      <state xmlns="urn:nokia.com:sros:ns:yang:sr:state">
        </state>
      </filter>
    </get>
  </rpc>
]]>]]>

```

Reply: No errors

```

<?xml version="1.0" encoding="UTF-8"?>
<rpc-reply message-id="101" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <data>
    <state xmlns="urn:nokia.com:sros:ns:yang:sr:state">

```

```

...
...
    </state>
  </data>
</rpc-reply>
]]>]]>

```

Example 4: The <state> tag does not contain a namespace

```

<?xml version="1.0" encoding="UTF-8"?>
<rpc message-id="101" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
<get>
  <filter>
    <state>
    </state>
  </filter>
</get>
</rpc>
]]>]]>

```

Reply: SR OS errors

```

<?xml version="1.0" encoding="UTF-8"?>
<rpc-reply message-id="101" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <rpc-error>
    <error-type>protocol</error-type>
    <error-tag>bad-element</error-tag>
    <error-severity>error</error-severity>
    <error-info>
      <bad-element>state</bad-element>
    </error-info>
    <error-message>
      Element is not valid in the specified context.
    </error-message>
  </rpc-error>
</rpc-reply>
]]>]]>

```

4.5.2 <edit-config> with XML Content Layer

An <edit-config> operation is supported with the <running> datastore and the <candidate> datastore.

The <edit-config> requests to the <candidate> datastore can only write XML-formatted content.

The <edit-config> requests that specify the running datastore as a target while using the “urn:nokia.com:sros:ns:yang:sr:conf” namespace (the Nokia SR OS YANG modules) result in an error response.

Example 1: using the <running> datastore with the urn:nokia.com:sros:ns:yang:sr:conf” namespace

```
<?xml version="1.0" encoding="UTF-8"?>
<rpc message-id="101" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <edit-config>
    <target><running/></target>
    <config>
      <configure xmlns="urn:nokia.com:sros:ns:yang:sr:conf">
        <python>
          <python-script>
            <script-name>testing</script-name>
          </python-script>
        </python>
      </configure>
    </config>
  </edit-config>
</rpc>
]]>]]>
```

Reply: with SR OS errors

```
<?xml version="1.0" encoding="UTF-8"?>
<rpc-reply message-id="101" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <rpc-error>
    <error-type>protocol</error-type>
    <error-tag>operation-not-supported</error-tag>
    <error-severity>error</error-severity>
    <error-info>
      <bad-element>running</bad-element>
    </error-info>
    <error-message>
      Writing to running datastore not supported in the specified namespace
    </error-message>
  </rpc-error>
</rpc-reply>
]]>]]>
```

There is an internal “implicit” lock that has a scope of all configuration commands in the SR OS (not just the “urn:nokia.com:sros:ns:yang:sr:conf” namespace). The following actions take/release the “implicit” lock:

- The first NETCONF <edit-config> on a <candidate> datastore triggers the “implicit” lock
- The completion of a NETCONF <commit> releases the “implicit” lock
- A CLI **admin** command can release the “implicit” lock. For more information, see [4.11](#)
- The NETCONF <discard-changes> command is supported in the SR OS which releases the “implicit” lock as well

The following scenarios are impacted when an “implicit” lock is taking place:

- A NETCONF session attempting an <edit-config> (on either the Alcatel-Lucent Base-R13 SR OS data model or the Nokia SR OS data model) is blocked and the SR OS replies with an error (the <error-info> element includes the <session-id> of the lock owner).
- A CLI command (on either the Alcatel-Lucent Base-R13 configuration set or the Nokia SR OS data model) is blocked and the SR OS returns an error.
- A SNMP set request (on objects that are part of the “urn:nokia.com:sros:ns:yang:sr:conf” namespace only) is blocked and the SR OS returns an error.

One or more <edit-config> requests can be performed on the candidate datastore before the changes are committed or discarded.

NETCONF <edit-config> and <commit> operations impact the configuration of the router and, as with some CLI or SNMP configuration changes, additional actions or steps may need to occur before certain configuration changes take operational effect. Some examples include:

- Configuration changes that require a **shutdown** and then **no shutdown** to be performed by an operator in order to take operational effect also need this explicit **shutdown** and then **no shutdown** to be performed via NETCONF (in separate edit-configs/commits) in order to take operational effect after those configuration items are changed. Some examples include:
 - changes to Autonomous System or Confederation value require a BGP **shutdown** and then **no shutdown**
 - changes to VPRN Max-routes requires a **shutdown** and then **no shutdown** on the VPRN service
 - changes to OSPF/ISIS export-limit require a **shutdown** and then **no shutdown** on OSPF/ISIS
- Configuration changes to an msap-policy that normally require a **tools perform subscriber-mgmt eval-msap** command to take operational effect on subscribers that are already active. NETCONF can be used to change the msap-policy configuration, but if it must have the configuration changes applied to the active subscribers then the operator must run the **eval-msap tools** command.

The supported <edit-config> operation attribute values are listed in [Table 55](#).

Table 55 <edit-config> Operation Attribute Values

Command	Notes
urn:alcatel-lucent.com:sros:ns:yang:conf-*-r13 namespace	Alcatel-Lucent Base-R13 SR OS YANG modules

Table 55 <edit-config> Operation Attribute Values (Continued)

Command	Notes
merge (Base-R13 SR OS modules)	<ul style="list-style-type: none"> For a merge operation, the operations and tags specified in an <edit-config> request are order-aware and order-dependent, and the sequence of merge operations must follow the required sequence of the equivalent CLI commands. The <edit-config> request is processed and executed in a top-down order. The same leaf can be enabled and disabled multiple times in the request and the final result is whatever was last specified for that leaf in the <edit-config> request.
remove (Base-R13 SR OS modules)	<ul style="list-style-type: none"> A <remove> operation is not supported for boolean leaves. For example, any of the following example commands will return an error: <ul style="list-style-type: none"> <shutdown operation="remove"/> <shutdown operation="remove">>false</shutdown> <interface operation="remove"> <interface-name>abc</interface-name> <shutdown>>true</shutdown> </interface> (For this last case <shutdown operation="merge">>true</shutdown> could be used instead to make the request valid.) A <remove> operation is the equivalent of no command in the CLI. This no command is applied whether the default for <i>command</i> is enabled (<i>command</i>), disabled (no command), or a specific value. The <remove> operation is not aware of the default value of the object or leaf being removed. A <remove> operation for a leaf where the request also specifies a value for the leaf, will result in an error.

Table 55 <edit-config> Operation Attribute Values (Continued)

Command	Notes
delete (Base-R13 SR OS modules)	<ul style="list-style-type: none"> • A <delete> operation for a leaf or a presence container will not return an error if the item is already deleted. • An error is returned if attempting to delete a list node that does not exist. • A <delete> operation for a container without presence will return an error. • A <delete> operation is not supported for boolean leaves. For example, any of the following example commands will return an error: <ul style="list-style-type: none"> – <shutdown operation="delete"/> – <shutdown operation="delete">>false</shutdown> – <interface operation="delete"> <ul style="list-style-type: none"> <interface-name>abc</interface-name> <shutdown>>true</shutdown> </interface> <p>(For this last case <shutdown operation="merge">>true</shutdown> could be used instead to make the request valid.)</p> • A <delete> operation is the equivalent of no command in the CLI. This no command is applied whether the default for <i>command</i> is enabled (<i>command</i>), disabled (no command), or a specific value. The <delete> operation is not aware of the default value of the object/leaf being deleted. • A <delete> operation on a node will ignore any values provided for that node (it will not check if that value is configured or valid), and it will ignore any data below that node (it will not check if that data exists or is valid).
create (Base-R13 SR OS modules)	<ul style="list-style-type: none"> • A <create> operation for a leaf or a presence container will not return an error if the item is being set to the same value. • An error is returned if attempting to create a list node that already exists. • A <create> operation for a container without presence will result in an “OK” response (no error) but will be silently ignored. • For a <create> operation, the operations and tags specified in an <edit-config> request are order-aware and order-dependent, and the sequence of create operations must follow the required sequence of the equivalent CLI commands. The <edit-config> request is processed and executed in a top-down order. The same leaf can be enabled and disabled multiple times in the request and the final result is whatever was last specified for that leaf in the <edit-config> request.
replace (Base-R13 SR OS modules)	<ul style="list-style-type: none"> • not supported
urn:nokia.com:sros:ns:yang:sr:conf namespace Nokia SR OS YANG modules	

Table 55 <edit-config> Operation Attribute Values (Continued)

Command	Notes
merge (Nokia SR OS modules)	<ul style="list-style-type: none"> supported
remove (Nokia SR OS modules)	<ul style="list-style-type: none"> A <remove> operation removes the deleted configuration and returns it to the default value. A <remove> operation automatically removes all child objects of a deleted object (leaves, lists, containers, and so on). Explicit shutdown of the object being removed (or any child) is not required and results in an error if a merge operation is specified on a tag that inherits a <remove> operation. A <remove> operation is allowed on non-presence containers. The non-presence container and all of its children are removed (for example, a non-presence container with no child nodes, is not displayed in a <get> or <get-config> reply). A <remove> operation is allowed on an object where all child branches and dependencies are automatically removed (but the <remove> operation fails if any outside objects refer to the object being removed). A <remove> operation is allowed on a <shutdown/> leaf (which returns it to its default value). A <remove> operation is allowed on a non-boolean leaf. Upon specifying a <remove> operation on a node where none of its children belong to the urn:nokia.com:sros:ns:yang:sr:conf namespace (the Nokia SR OS YANG modules), the SR OS does not return an error and completes the node removal. A <remove> operation for a leaf where the request also specifies a value for the leaf, results in an error.

Table 55 <edit-config> Operation Attribute Values (Continued)

Command	Notes
delete (Nokia SR OS modules)	<ul style="list-style-type: none"> • The SR OS returns an error if a <delete> operation is performed on a list that does not specify a key (that is, an attempt to delete all members of a list). • The SR OS returns an error if a <delete> operation is performed on a leaf or presence container that is already deleted (or has the default value and the default-handling is trim). • The SR OS may return an error and may not complete the deletion operation when a <delete> operation is performed on a node where any of its children do not belong to the urn:nokia.com:sros:ns:yang:sr:conf namespace (the Nokia SR OS YANG modules). • A <delete> operation removes the deleted configuration and returns it to the default value. • A <delete> operation automatically deletes all child objects of a deleted object (leaves, lists, containers, and so on). • Explicit shutdown of the object being deleted (or any of its children) is not required and results in an error if a merge operation is specified on a tag that inherits a <delete> operation. • A <delete> operation is allowed on non-presence containers. The non-presence container and all of its children are deleted (for example, a non-presence container with no child nodes is not displayed in a <get> or <get-config> reply). • A <delete> operation is allowed on an object where all child branches and dependencies are automatically deleted (but the <delete> operation fails if any outside objects refer to the object being deleted). • A <delete> operation is allowed on a <shutdown/> leaf (which returns it to its default value). • A <delete> operation is allowed on a non-boolean leaf. • Upon specifying a <delete> operation on a node where none of its children belong to the urn:nokia.com:sros:ns:yang:sr:conf namespace (the Nokia SR OS YANG modules), the SR OS does not return an error and completes the node deletion. • A <delete> operation for a leaf where the request also specifies a value for the leaf, will result in an error.
create (Nokia SR OS modules)	<ul style="list-style-type: none"> • When a <create> operation for a leaf or presence container is performed, the SR OS returns an error if the leaf or presence container is being set to the same value (unless the default-handling is trim and the value being set is the default value).
replace (Nokia SR OS modules)	<ul style="list-style-type: none"> • Not supported

The <edit-config> operation's <default-operation> parameter is supported with the following values:

- merge
- none
 - In the urn:alcatel-lucent.com:sros:ns:yang:conf-*-r13 namespace (the Alcatel-Lucent Base-R13 SR OS YANG modules), an operation of "none" on a leaf node (inherited or direct) causes that leaf statement to be ignored. No error will be returned if the leaf does not exist in the data model.
 - In the urn:nokia.com:sros:ns:yang:sr:conf namespace (the Nokia SR OS YANG modules), an operation of "none" (inherited or direct) on a leaf node that does not exist in the data model causes the SR OS to return an error with an <error-tag> value of data-missing.

For <delete> and <remove> operations in the Nokia SR OS namespace, the SR OS NETCONF server will recursively "unwind" any children of the node being deleted or removed first before removing the node itself. The 'deepest' child branch of the request is examined first and any leaves are processed, after which the server works backwards out of the deepest branches back up to the object where the delete operation was specified.

For urn:alcatel-lucent.com:sros:ns:yang:conf-*-r13 namespace (the Alcatel-Lucent Base-R13 SR OS YANG modules), if child branches of an object are required to be removed before deleting the object in the CLI, then the equivalent delete request in a NETCONF <edit-config> request must contain all those children if they exist). For example:

```
<?xml version="1.0" encoding="UTF-8"?>
<rpc message-id="101" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <edit-config>
    <target><running/></target>
    <config>
      <configure xmlns="urn:alcatel-lucent.com:sros:ns:yang:conf-r13">
        <service>
          <vpls operation="delete">
            <service-id>11</service-id>
            <interface>
              <ip-int-name>test</ip-int-name>
              <shutdown operation="merge">true</shutdown>
            </interface>
            <shutdown operation="merge">true</shutdown>
          </vpls>
        </service>
      </configure>
    </config>
  </edit-config>
</rpc>
]]>]]>
```

In the example above, the SR OS will first shut down the test interface, then delete the interface, then shut down the VPLS, and then finally remove it.



Note: In the urn:alcatel-lucent.com:sros:ns:yang:conf-*r13 namespace (the Alcatel-Lucent Base-R13 SR OS YANG modules), the 'operation="merge"' is required in the shutdown nodes; otherwise the inherited operation is delete, which is not supported on boolean leaves.

In the example above, if other children of vpls 11 exist in the config besides the interface test specified in the delete request above, and those children are required in the CLI to be deleted before removing vpls 11, then the deletion request above will fail. All configured children must be specified in the delete request.

The following applies to the urn:nokia.com:sros:ns:yang:sr:conf namespace (the Nokia SR OS YANG modules).

- The SR OS returns an error if an explicitly defined <edit-config> operation (such as “delete”) is specified on a “key” leaf.
- The “operation” attribute is inherited from the parent node if not explicitly specified (same as namespaces). If no parent node is available, then the “default-operation” value is used. In other words, the “operation” attribute has a “scope” that it applies to the nested nodes until it is redefined. The following scenarios simplify the “operation” inheritance, where the first line in each scenario represents the operation value of the parent node and the following lines represent the possible operation values for the child nodes and the SR OS behavior in each case:
 1. Create
 - Create/Merge: The SR OS processes request (request succeeds/fails based on operation’s behavior)
 - Delete/Remove: The SR OS returns an error
 2. Merge
 - Create/Merge/Delete/Remove: The SR OS processes request (request succeeds/fails based on operation’s behavior)
 3. Delete/Remove
 - Create/Merge: The SR OS returns an error
 - Delete/Remove: The SR OS processes request (request succeeds/fails based on operation’s behavior)

4.5.3 <get-config> with XML Content Layer

A <get-config> operation is supported with the <running> datastore and the <candidate> datastore.

The <get-config> requests on the <candidate> datastore return only XML-formatted content.

On a <candidate> datastore, if no filter is specified, SR OS returns the Nokia SR OS configurations only.

On the <running> datastore, if no filter is specified, SR OS returns both the Alcatel-Lucent Base-R13 configurations and the Nokia SR OS configurations.

On the <running> datastore, to return configurations from the Alcatel-Lucent Base-R13 configurations only (or the Nokia SR OS configurations only), the user must specify at least a top-level tag and a namespace in the filter. If the namespace is not specified, SR OS returns an error.

The following applies to the urn:alcatel-lucent.com:sros:ns:yang:conf-*-r13 namespace (the Alcatel-Lucent Base-R13 SR OS YANG modules):

- <get-config> requests that specify a non-existing list node or presence container will result in a reply that contains no data for those list nodes or containers. An <rpc-error> is not sent in this case.

The following applies to the urn:nokia.com:sros:ns:yang:sr:conf namespace (the Nokia SR OS YANG modules):

- <get-config> requests that specify a non-existing list node or presence container result in an <rpc-error> response
- <get-config> requests that specify a list without specifying a key result in an <rpc-error> response

Using the 'report-all' value with the <with-defaults> tag (RFC 6243) in an XML-content layer <get-config>, returns the equivalent of the CLI command **info detail** (the returned data includes attributes that are set to their default values).

Example 1: use of <with-defaults> with a value of "report-all"

```
<?xml version="1.0" encoding="UTF-8"?>
<rpc message-id="101" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
<get-config>
  <source>
    <candidate/>
  </source>
  <filter>
    <configure xmlns="urn:nokia.com:sros:ns:yang:sr:conf">
```

```

        <system>
            <security>
                <cpm-filter>
                    <ipv6-filter>
                    </ipv6-filter>
                </cpm-filter>
            </security>
        </system>
    </configure>
</filter>
<with-defaults xmlns="urn:ietf:params:xml:ns:yang:ietf-netconf-with-defaults">
    report-all
</with-defaults>
</get-config>
</rpc>
]]>]]>

```

Reply: returns even attributes with default values

```

<?xml version="1.0" encoding="UTF-8"?>
<rpc-reply message-id="101" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
    <data>
        <configure xmlns="urn:nokia.com:sros:ns:yang:sr:conf">
            <system xmlns="urn:nokia.com:sros:ns:yang:sr:conf-system">
                <security>
                    <cpm-filter>
                        <ipv6-filter>
                            <shutdown>true</shutdown>
                        </ipv6-filter>
                    </cpm-filter>
                </security>
            </system>
        </configure>
    </data>
</rpc-reply>
]]>]]>

```

Example 2: without using <with-defaults>

```

<?xml version="1.0" encoding="UTF-8"?>
<rpc message-id="101" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
<get-config>
    <source>
        <candidate/>
    </source>
    <filter>
        <configure xmlns="urn:nokia.com:sros:ns:yang:sr:conf">
            <system>
                <security>
                    <cpm-filter>
                        <ipv6-filter>
                        </ipv6-filter>
                    </cpm-filter>
                </security>
            </system>
        </configure>
    </filter>

```

```
</get-config>
</rpc>
]]>]]>
```

Reply: Attributes with default values are not returned

```
<?xml version="1.0" encoding="UTF-8"?>
<rpc-reply message-id="101" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <data>
    <configure xmlns="urn:nokia.com:sros:ns:yang:sr:conf">
      <system>
        <security>
          <cpm-filter>
            <ipv6-filter>
              </ipv6-filter>
            </cpm-filter>
          </security>
        </system>
      </configure>
    </data>
  </rpc-reply>
]]>]]>
```

Subtree filtering for basic subtree selection is supported for XML content layer <get-config> requests. Post-filtering of the selected subtrees is not supported.

In the urn:alcatel-lucent.com:sros:ns:yang:conf-*-r13 namespace (the Alcatel-Lucent Base-R13 SR OS YANG modules), the subtree filtering behaves as follows.

- Attribute match expressions (section 6.2.2 of RFC 6241) are not supported. See details below about content match nodes.
- Only containers are supported as selection nodes (section 6.2.4 of RFC 6241). Empty leaf nodes or list name nodes are not supported as selection nodes.
 - Nodes that represent lists must also include content match nodes for all keys of the list; for example, <configure><router><interface><interface-name>abc</interface-name>.
 - A selection node that is a list but does not have a key specified is not supported; for example, <configure><router><interface/> is not supported. An alternative is to request the parent containment node that contains the desired list node; for example, <configure><router> instead of <configure><router><interface/>.
- Content match nodes (section 6.2.5 of RFC 6241) are only supported for key leaves; for example, <configure><router><interface><interface-name>abc</interface-name>.
 - Content match nodes that are leaves but are not also keys will result in an error (not silently ignored).

Example 3 — A non key leaf is specified (for example, shutdown)

```

<?xml version="1.0" encoding="UTF-8"?>
<rpc message-id="101" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <get-config>
    <source><running/></source>
    <filter>
      <configure xmlns="urn:alcatel-lucent.com:sros:ns:yang:conf-r13">
        <router>
          <interface>
            <interface-name>abc</interface-name>
            <shutdown>false</shutdown>
          </interface>
        </router>
      </configure>
    </filter>
  </get-config>
</rpc>
]]>]]>

```

Reply: SR OS errors

```

<?xml version="1.0" encoding="UTF-8"?>
<rpc-reply message-id="101" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <rpc-error>
    <error-type>protocol</error-type>
    <error-tag>operation-not-supported</error-tag>
    <error-severity>error</error-severity>
    <error-info>
      <bad-element>shutdown</bad-element>
    </error-info>
    <error-message>
      Leaf element specified which is not a key.
    </error-message>
  </rpc-error>
</rpc-reply>
]]>]]>

```

Multiple key leaves for the same key cannot be requested inside the same instance of the list name node; for example, `<interface-name>abc</interface-name> <interface-name>def</interface-name>`. Each key value must be inside its own instance of the list name node; for example, `<interface> <interface-name>abc</interface-name> </interface> <interface> <interface-name>def</interface-name> </interface>`.

Example 4 — A valid `<get-config>` request (content match on a list key):

```

<?xml version="1.0" encoding="UTF-8"?>
<rpc message-id="101"
  xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <get-config>
    <source>
      <running/>
    </source>
    <filter>
      <configure xmlns="urn:alcatel-lucent.com:sros:ns:yang:conf-r13">
        <router>
          <interface>
            <interface-name>abc</interface-name>

```

```

        </interface>
    </router>
</configure>
</filter>
</get-config>
</rpc>
]]>]]>

```

Reply:

```

<?xml version="1.0" encoding="UTF-8"?>
<rpc-reply message-id="101" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <data>
    <configure xmlns="urn:alcatel-lucent.com:sros:ns:yang:conf-r13">
      <router>
        <router-instance>Base</router-instance>
        <interface>
          <interface-name>abc</interface-name>
        </interface>
      </router>
    </configure>
  </data>
</rpc-reply>
]]>]]>

```

Example 5 — A valid <get-config> request (selection node that is a container):

```

<?xml version="1.0" encoding="UTF-8"?>
<rpc message-id="101"
  xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <get-config>
    <source>
      <running/>
    </source>
    <filter>
      <configure xmlns="urn:alcatel-lucent.com:sros:ns:yang:conf-r13">
        <router/>
      </configure>
    </filter>
  </get-config>
</rpc>
]]>]]>

```

The reply will contain all the configuration for all child nodes of config>router

Example 6 — An invalid <get-config> request (list name node - invalid selection node):

```

<?xml version="1.0" encoding="UTF-8"?>
<rpc message-id="101" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <get-config>
    <source><running/></source>
    <filter>
      <configure xmlns="urn:alcatel-lucent.com:sros:ns:yang:conf-r13">
        <router>
          <interface>

```

```

        </interface>
      </router>
    </configure>
  </filter>
</get-config>
</rpc>
]]>]]>

```

Reply: SR OS errors

```

<?xml version="1.0" encoding="UTF-8"?>
<rpc-reply message-id="101" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <rpc-error>
    <error-type>application</error-type>
    <error-tag>operation-failed</error-tag>
    <error-severity>error</error-severity>
    <error-info>
      <err-element>get-config</err-element>
    </error-info>
    <error-message>
      command failed - 'configure router interface'
    </error-message>
  </rpc-error>
</rpc-reply>
]]>]]>

```

Example 7 — An invalid <get-config> request (empty leaf node - invalid selection node):

```

<?xml version="1.0" encoding="UTF-8"?>
<rpc message-id="101"
  xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <get-config>
    <source>
      <running/>
    </source>
    <filter>
      <configure xmlns="urn:alcatel-lucent.com:sros:ns:yang:conf-r13">
        <system>
          <security>
            <ftp-server>
            </ftp-server>
          </security>
        </system>
      </configure>
    </filter>
  </get-config>
</rpc>
]]>]]>

```

Reply:

```

<?xml version="1.0" encoding="UTF-8"?>
<rpc-reply message-id="101" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <rpc-error>
    <error-type>protocol</error-type>

```



```

    <error-tag>operation-not-supported</error-tag>
    <error-severity>error</error-severity>
    <error-info>
      <bad-element>ftp-server</bad-element>
    </error-info>
    <error-message>
      Leaf element specified which is not a key.
    </error-message>
  </rpc-error>
</rpc-reply>
]]]]>

```

Example 8 — An invalid <get-config> request (key repeated in the same instance of the list node):

```

<?xml version="1.0" encoding="UTF-8"?>
<rpc message-id="101"
xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
<get-config>
  <source><running/></source>
  <filter>
    <configure xmlns="urn:alcatel-lucent.com:sros:ns:yang:conf-r13">
      <router>
        <interface>
          <interface-name>abc</interface-name>
          <interface-name>def</interface-name>
        </interface>
      </router>
    </configure>
  </filter>
</get-config>
</rpc>
]]]]>

```

Reply: SR OS errors

```

<?xml version="1.0" encoding="UTF-8"?>
<rpc-reply message-id="101" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <rpc-error>
    <error-type>application</error-type>
    <error-tag>operation-failed</error-tag>
    <error-severity>error</error-severity>
    <error-info>
      <err-element>get-config</err-element>
    </error-info>
    <error-message>
      command failed -
      &apos;configure router interface &quot;abc&quot; &quot;def&quot;&apos;;
    </error-message>
  </rpc-error>
</rpc-reply>
]]]]>

```

The full configuration (equivalent to the CLI command 'admin display-config') can be obtained via a <get-config> request:

- A — when the <filter> tag is not present

For example:

```
<?xml version="1.0" encoding="UTF-8"?>
<rpc message-id="101" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <get-config>
    <source><running/></source>
  </get-config>
</rpc>
]]>]]>
```

```
<?xml version="1.0" encoding="UTF-8"?>
<rpc message-id="101" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <get-config>
    <source><candidate/></source>
  </get-config>
</rpc>
]]>]]>
```

- B — when only the <configure> tag is present inside a <filter> tag

For example:

```
<?xml version="1.0" encoding="UTF-8"?>
<rpc message-id="101" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <get-config>
    <source><running/></source>
    <filter>
      <configure xmlns="urn:alcatel-lucent.com:sros:ns:yang:conf-r13"/>
    </filter>
  </get-config>
</rpc>
]]>]]>
```

```
<?xml version="1.0" encoding="UTF-8"?>
<rpc message-id="101" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <get-config>
    <source><candidate/></source>
    <filter>
      <configure xmlns="urn:nokia.com:sros:ns:yang:sr:conf"/>
    </filter>
  </get-config>
</rpc>
]]>]]>
```

4.6 XML Content Layer Examples

The following examples can be used after a NETCONF session has been established including the exchange of the <hello> messages.

The following is an example of a <get-config> request on the <running> datastore to check on whether netconf is shut down or not on the router:

```
<?xml version="1.0" encoding="UTF-8"?>
<rpc message-id="101" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <get-config>
    <source><running/></source>
    <filter>
      <configure xmlns="urn:alcatel-lucent.com:sros:ns:yang:conf-r13">
        <system>
          <netconf>
            </netconf>
          </system>
        </configure>
      </filter>
    </get-config>
  </rpc>
</></>
```

Reply:

```
<?xml version="1.0" encoding="UTF-8"?>
<rpc-reply message-id="101" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <data>
    <configure xmlns="urn:alcatel-lucent.com:sros:ns:yang:conf-r13">
      <system>
        <netconf>
          <shutdown>false</shutdown>
        </netconf>
      </system>
    </configure>
  </data>
</rpc-reply>
</></>
```

The following is an example for a <get-config> request on the <candidate> datastore to get the full configurations of the system, qos and log branches:

```
<?xml version="1.0" encoding="UTF-8"?>
<rpc message-id="101" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
<get-config>
  <source><candidate/></source>
  <filter>
    <configure xmlns="urn:nokia.com:sros:ns:yang:sr:conf">
      <system>
        </system>
      </configure>
    <configure xmlns="urn:nokia.com:sros:ns:yang:sr:conf">
      <qos>
        </qos>
      </configure>
    <configure xmlns="urn:nokia.com:sros:ns:yang:sr:conf">
      <log/>
    </configure>
  </filter>
</get-config>
</rpc>
</></>
```

Reply:

```

<?xml version="1.0" encoding="UTF-8"?>
<rpc-reply message-id="101" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <data>
    <configure xmlns="urn:nokia.com:sros:ns:yang:sr:conf">
      <system>
        <contact>tester</contact>
        <name>r2-node</name>
        <location>over-here</location>
        <lldp>
          <shutdown>>false</shutdown>
        </lldp>
        ...
        ...
      </system>
      <qos>
        <sap-ingress>
          <policy-id>1</policy-id>
          <policy-name>default</policy-name>
          <description>Default SAP ingress QoS policy.</description>
          <sub-insert-shared-pccrule>
          </sub-insert-shared-pccrule>
          <dynamic-policer>
            <range>
            </range>
            <parent>
            </parent>
          </dynamic-policer>
          <mac-criteria>
          </mac-criteria>
          <ip-criteria>
          </ip-criteria>
          <ipv6-criteria>
          </ipv6-criteria>
          ...
          ...
        </sap-ingress>
        </qos>
        <log>
          <route-preference>
          </route-preference>
          <app-route-notifications>
          </app-route-notifications>
          <event-control>
            <application-id>1</application-id>
            <event-number>4401</event-number>
            <severity-level>major</severity-level>
            <throttle>>true</throttle>
          </event-control>
          ...
          ...
        </log>
      </configure>
    </data>
  </rpc-reply>
</></></>

```

The following is an example of an <edit-config> request on the <running> datastore to create a basic VPRN service:

```
<?xml version="1.0" encoding="UTF-8"?>
<rpc message-id="101"
  xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <edit-config>
    <target>
      <running/>
    </target>
    <config>
      <configure xmlns="urn:alcatel-lucent.com:sros:ns:yang:conf-r13">
        <service>
          <vprn operation="create">
            <service-id>200</service-id>
            <customer>1</customer>
          </vprn>
        </service>
      </configure>
    </config>
  </edit-config>
</rpc>
]]>]]>
```

Reply:

```
<?xml version="1.0" encoding="UTF-8"?>
<rpc-reply message-id="101" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <ok/>
</rpc-reply>
]]>]]>
```

The following is an example of an <edit-config> request on the <candidate> datastore to create a basic epipe service:

```
<?xml version="1.0" encoding="UTF-8"?>
<rpc message-id="101" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <edit-config>
    <target><candidate/></target>
    <config>
      <configure xmlns="urn:nokia.com:sros:ns:yang:sr:conf">
        <service>
          <epipe>
            <service-id>444</service-id>
            <customer>1</customer>
            <service-mtu>1514</service-mtu>
          </epipe>
        </service>
      </configure>
    </config>
  </edit-config>
</rpc>
]]>]]>
```

Reply:

```
<?xml version="1.0" encoding="UTF-8"?>  
<rpc-reply message-id="101" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">  
  <ok/>  
</rpc-reply>  
]]>]]>
```

4.7 CLI Content Layer

When using the CLI format at the NETCONF content layer, configuration changes and configuration information retrieved are expressed as untagged (non-XML) CLI commands; for example, CLI script.

The script must be correctly ordered and has the same dependencies and behavior as CLI. The location of CR/LF (ENTER) within the CLI for an <edit-config> is significant and affects the processing of the CLI commands, such as what CLI branch is considered the “working context”. In the following two examples, the “working context” after the commands are issued are different.

Example 1:

```
exit all [-ENTER]
configure system time zone EST [-ENTER]
```

Example 2:

```
exit all [-ENTER]
configure [-ENTER]
  system [-ENTER]
    time [-ENTER]
      zone EST [-ENTER]
```

After example 1, the CLI working context is the root and immediately sending 'dst-zone CEST' would return an error. After example 2, the CLI working context is config>system>time and sending 'dst-zone CEST' would work as expected.

Configuration changes done via NETCONF trigger the same “change” log events (for example, tmnxConfigCreate) as a normal CLI user doing the same changes.

The <with-defaults> tag (RFC 6243) is not supported in a CLI content layer request.

The operator can get a full configuration including defaults for a CLI Content Layer using an empty <cli-info-detail>. The full configuration (equivalent to the CLI command 'admin display-config [detail]') can be obtained via a <get-config> request in a CLI Content Layer format with an empty <cli-info> or <cli-info-detail> tag inside a <config-format-cli-block>. <report-all> is not supported.

Post-processing commands are ignored: "| match" (pipe match), "| count" (pipe count) and ">" (redirect to file) and CLI ranges are not supported for any command; for example, show card [1..5].

4.8 CLI Content Layer Examples

The following examples can be used after a NETCONF session has been established including the exchange of the <hello> messages.

The following shows an example of a configuration change request and response.



Note: The **exit all** command is not required at the beginning of the CLI block; it is automatically assumed by the SR OS NETCONF server.

```
<?xml version="1.0" encoding="UTF-8"?>
<rpc message-id="104" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <edit-config>
    <target><running/></target>
    <config>
      <config-format-cli-block>
        configure system
          time zone EST
          location over-here
        exit all
      </config-format-cli-block>
    </config>
  </edit-config>
</rpc>
]]>]]>
```

Reply:

```
<?xml version="1.0" encoding="UTF-8"?>
<rpc-reply message-id="104"
xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <ok/>
</rpc-reply>
]]>]]>
```

The following is an example of a <get-config> request and response to retrieve configuration information:

```
<?xml version="1.0" encoding="UTF-8"?>
<rpc message-id="101" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <get-config>
    <source>
      <running/>
    </source>
    <filter>
      <config-format-cli-block>
        <cli-info>router</cli-info>
        <cli-info-detail>system login-control</cli-info-detail>
      </config-format-cli-block>
    </filter>
  </get-config>
```



```
</rpc>
]]>]]>
```

Reply:

```
<?xml version="1.0" encoding="UTF-8"?>
<rpc-reply message-id="101" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <data xmlns="urn:alcatel-lucent.com:sros:ns:yang:cli-content-layer-r13">
    <config-format-cli-block>
      <item>
        <cli-info>router</cli-info>
        <response>
          -----
          #-----
          echo "IP Configuration"
          #-----
            interface "system"
              no shutdown
            exit
          -----
          </response>
        </item>
        <item>
          <cli-info-detail>system login-control</cli-info-detail>
          <response>
            -----
            ftp
              inbound-max-sessions 3
            exit
            ssh
              no disable-graceful-shutdown
              inbound-max-sessions 5
              outbound-max-sessions 5
              no ttl-security
            exit
            telnet
              no enable-graceful-shutdown
              inbound-max-sessions 5
              outbound-max-sessions 5
              no ttl-security
            exit
            idle-timeout 30
            no pre-login-message
            no motd
            login-banner
            no exponential-backoff
          -----
          </response>
        </item>
      </config-format-cli-block>
    </data>
  </rpc-reply>
]]>]]>
```

The following example shows a <get-config> request and response to retrieve full configuration information.



Note: The `<cli-info-detail/>` request can be used to get the full configuration, including default settings.

```
<?xml version="1.0" encoding="UTF-8"?>
<rpc message-id="101" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <get-config>
    <source>
      <running/>
    </source>
    <filter>
      <config-format-cli-block>
        <cli-info/>
      </config-format-cli-block>
    </filter>
  </get-config>
</rpc>
]]>]]>
```

Reply:

```
<?xml version="1.0" encoding="UTF-8"?>
<rpc-reply message-id="101" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <data xmlns="urn:alcatel-lucent.com:sros:ns:yang:cli-content-layer-r13">
    <config-format-cli-block>
      <item>
        <cli-info></cli-info>
        <response>
# TiMOS-C-0.0.I4301 cpm/x86_64 ALCATEL SR 7750 Copyright (c) 2000-2015 Alcatel-
Lucent.
# All rights reserved. All use subject to applicable license agreements.
# Built on Sun Jan 4 19:11:11 PST 2015 by builder in /rel0.0/I4301/panos/main

# Generated WED JAN 07 01:07:43 2015 UTC

exit all
configure
#-----
echo "System Configuration"
#-----
  system
    dns
    exit
    load-balancing
      lsr-load-balancing lbl-ip
      system-ip-load-balancing
    exit
    netconf
      no shutdown
    exit
    snmp
      shutdown
      engineID "deadbeefdeadbeef"
    exit
    time
      ntp
```

```

        authentication-key 1 key "OAwgNULbZgI" hash2 type des
        no shutdown
    exit
    snmp
        shutdown
    exit
    zone EST
    exit
    thresholds
        rmon
    exit
    exit
#-----
echo "Cron Configuration"
#-----
    cron
        ...
        ...
        ...
    exit
    exit
#-----
echo "System Security Configuration"
#-----
    ...
    ...
    ...
#-----
echo "System Time NTP Configuration"
#-----
    system
        time
            ntp
        exit
    exit
    exit

exit all

# Finished WED JAN 07 01:07:43 2015 UTC
-----
-----
        </response>
    </item>
</config-format-cli-block>
</data>
</rpc-reply>
]]>]]>

```

The following is an example of a <get> request and the response to it:

```

<?xml version="1.0" encoding="UTF-8"?>
<rpc message-id="101" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <get>
    <filter>
      <oper-data-format-cli-block>
        <cli-show>system security ssh</cli-show>
      </oper-data-format-cli-block>
    </filter>
  </get>
</rpc>

```

```

        </filter>
      </get>
</rpc>
]]>]]>

```

Reply:

```

<?xml version="1.0" encoding="UTF-8"?>
<rpc-reply message-id="101" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <data xmlns="urn:alcatel-lucent.com:sros:ns:yang:cli-content-layer-r13">
    <oper-data-format-cli-block>
      <item>
        <cli-show>system security ssh</cli-show>
      <response>

```

```

=====
SSH Server
=====
Administrative State      : Enabled
Operational State       : Up
Preserve Key             : Enabled

SSH Protocol Version 1   : Disabled

SSH Protocol Version 2   : Enabled
DSA Host Key Fingerprint : ca:ce:37:90:49:7d:cc:68:22:b3:06:2c:11:cd:3c:8e
RSA Host Key Fingerprint : 49:7c:21:97:42:35:83:61:06:95:cd:a8:78:4c:1e:76

```

```

-----
Connection                               Username
  Version Cipher                          ServerName  Status
-----
135.121.143.254                          admin
  2          aes128-cbc                    netconf    connected
-----

```

```

Number of SSH sessions : 1
=====
      </response>
    </item>
  </oper-data-format-cli-block>
</data>
</rpc-reply>
]]>]]>

```

4.9 NETCONF Configuration Command Reference

This section provides the NETCONF configuration command reference. Topics in this section include:

- [Command Hierarchies](#)
- [Configuration Commands](#)

4.9.1 Command Hierarchies

4.9.1.1 NETCONF System Commands

```
config
  — system
    — netconf
      — capabilities
        — [no] candidate
        — [no] writable-running
      — [no] shutdown
      — yang-modules
        — [no] base-r13-modules
        — [no] nokia-modules
```

4.9.1.2 NETCONF Security Commands

```
config
  — system
    — security
      — profile profile-id
        — netconf
          — base-op-authorization
            — [no] kill-session
            — [no] lock
```

4.9.2 Configuration Commands

This section provides NETCONF configuration command descriptions.

4.9.2.1 NETCONF System Commands

shutdown

Syntax	[no] shutdown
Context	config>system>netconf
Description	This command disables the NETCONF server. The shutdown command is blocked if there are any active NETCONF sessions. Use the admin disconnect command to disconnect all NETCONF sessions before shutting down the NETCONF service.

candidate

Syntax	[no] candidate
Context	config>system>netconf>capabilities
Description	This command enables or disables support of the candidate datastore in the SR OS NETCONF server. If the candidate is disabled then requests that reference the candidate datastore return an error, and when a NETCONF client establishes a new session the candidate capability is not advertised in the SR OS <hello>. This command also controls support of the <commit> and <discard-changes> operations.
Default	candidate

writable-running

Syntax	[no] writable-running
Context	config>system>netconf>capabilities
Description	This command enables or disables support of the writable-running capability in the SR OS NETCONF server. If writable-running is disabled then requests that reference the running datastore as a target return an error, and when a NETCONF client establishes a new session the writable-running capability is not advertised in the SR OS <hello>.
Default	writable-running

base-r13-modules

Syntax	[no] base-r13-modules
Context	config>system>netconf>yang-modules
Description	This command enables or disables support of the Base-R13 YANG modules in the SR OS NETCONF server. If the base-r13-modules are disabled then requests that reference the Base-R13 modules return an error, and when a NETCONF client establishes a new session the Base-R13 modules are not advertised in the SR OS <hello>.
Default	base-r13-modules

nokia-modules

Syntax	[no] nokia-modules
Context	config>system>netconf>yang-modules
Description	This command enables or disables support of the Nokia YANG modules in the SR OS NETCONF server. If the nokia-modules are disabled then requests that reference the Nokia modules return an error, and when a NETCONF client establishes a new session the Nokia modules are not advertised in the SR OS <hello>.
Default	nokia-modules

4.9.2.2 NETCONF Security Commands

netconf

Syntax	netconf
Context	config>system>security>profile
Description	This command authorizes netconf capability for the user.

base-op-authorization

Syntax	base-op-authorization
Context	config>system>security>profile>netconf
Description	This command enables the context where permission to use various NETCONF operations is controlled.

kill-session

Syntax	[no] kill-session
Context	config>system>security>profile>netconf>base-op-authorization
Description	This operation authorizes a user associated with the profile to send a kill session NETCONF operation. This kill session operation allows a NETCONF client to kill another NETCONF session, but not the session in which the operation is requested.
Default	no kill-session

lock

Syntax	[no] lock
Context	config>system>security>profile>netconf>base-op-authorization
Description	This operation authorizes a user associated with the profile to send a lock NETCONF operation. This lock operation allows a NETCONF client to lock the running datastore or the candidate datastore.
Default	no lock

4.10 NETCONF Show Command Reference

4.10.1 Command Hierarchies

4.10.1.1 Show Commands

```
show
  — system
    — netconf
      — counters
```

4.10.2 Command Descriptions

4.10.2.1 Show Commands

Command outputs shown in this section are examples only; actual displays may differ depending on supported functionality and user configuration.

4.10.2.1.1 NETCONF System Commands

netconf

Syntax	netconf
Context	show>system
Description	This command displays NETCONF SSH sessions.
Output	SSH Options Output

[Table 56](#) describes the NETCONF output fields.

Table 56 Show System NETCONF Output Fields

Label	Description
Administrative State	Enabled Displays that NETCONF is enabled. Disabled Displays that NETCONF is disabled.
Operational State	Up Displays that NETCONF is operational. Down Displays that NETCONF is not operational.
Connection	The IP address of the connected router(s) (remote client).
Username	The name of the user.
Session ID	The NETCONF session ID.
Status	Connected or not connected.
Number of NETCONF sessions	Total NETCONF sessions
Running Locked?	Yes Displays that the <running> datastore is locked. No Displays that the <running> datastore is not locked.
Candidate Locked?	Yes Displays that the <candidate> datastore is locked. No Displays that the <candidate> datastore is not locked

Sample Output

```
# show system netconf
=====
NETCONF Server
=====
Administrative State      : Enabled
Operational State        : Up
-----
Connection      Username      Session Status      Running      Candidate
                Id              Id                  Locked?      Locked?
-----
135.224.26.145  admin          17      connected      no          no
135.224.26.145  admin          15      connected      no          no
-----
```

```
Number of NETCONF sessions : 2
=====
```

counters

- Syntax** **counters**
- Context** show>system>netconf
- Description** This command displays NETCONF counters.
- Output** SSH Options Output

[Table 57](#) describes the NETCONF counter output fields.

Table 57 NETCONF Counters Output Fields

Label	Description
RX Messages	Types and numbers of received messages
RX Total	Total of all received messages
TX Messages	Types and numbers of sent messages
TX Total	Total of all sent messages
failed edit-configs	Number of failed <edit-config> requests due to a lock (including implicit ones) being taken by other netconf sessions
failed locks	Number of failed <lock> requests due to a lock (including implicit ones) being taken by other netconf sessions

Sample Output

```
# show system netconf counters
=====
NETCONF counters:
=====
    Rx Messages
-----
    in gets           : 23
    in get-configs   : 19
    in edit-configs  : 35
    in copy-configs  : 0
    in delete-configs : 0
    in validates     : 0
    in close-sessions : 0
    in kill-sessions  : 0
    in locks         : 0
    in unlocks       : 0
    in commits       : 2
    in discards      : 1
```

```
-----  
Rx Total           : 80  
-----  
Tx Messages  
-----  
out rpc-errors    : 4  
-----  
Tx Total          : 9  
-----  
Failed requests due to lock being taken by other netconf sessions  
-----  
failed edit-configs: 1  
failed locks       : 0  
=====
```

4.11 NETCONF Admin Command Reference

4.11.1 Command Hierarchies

4.11.1.1 Admin Commands

```
admin
  — system
    — candidate
      — discard-changes datastore-type
```

4.11.2 Command Descriptions

4.11.2.1 Admin Commands

Command outputs shown in this section are examples only; actual displays may differ depending on supported functionality and user configuration.

discard-changes

Syntax	discard-changes <i>datastore-type</i>
Context	admin>system>candidate
Description	This operation discards uncommitted changes on the <candidate> datastore.
Parameters	<i>datastore-type</i> — The datastore type.
Values	global

5 Event and Accounting Logs

5.1 In This Chapter

This chapter provides information about configuring event and accounting logs in the system.

Topics in this chapter include:

- [Logging Overview](#)
- [Log Destinations](#)
- [Event Logs](#)
 - [Event Sources](#)
 - [Event Control](#)
 - [Log Manager and Event Logs](#)
 - [Event Filter Policies](#)
 - [Event Log Entries](#)
 - [Simple Logger Event Throttling](#)
 - [Default System Log](#)
 - [Event Handling System](#)
- [Accounting Logs](#)
 - [Accounting Records](#)
 - [Accounting Files](#)
 - [Design Considerations](#)
- [Configuration Notes](#)

5.2 Logging Overview

The two primary types of logging supported in the OS are event logging and accounting logs.

Event logging controls the generation, dissemination and recording of system events for monitoring status and troubleshooting faults within the system. The OS groups events into four major categories or event sources:

- Security events — Events that pertain to attempts to breach system security.
- Change events — Events that pertain to the configuration and operation of the node.
- Main events — Events that pertain to applications that are not assigned to other event categories/sources.
- Debug events — Events that pertain to trace or other debugging information.

The following are events within the OS and have the following characteristics:

- A time stamp in UTC or local time.
- The generating application.
- A unique event ID within the application.
- The VRF-ID.
- A subject identifying the affected object.
- A short text description.

Event control assigns the severity for each application event and whether the event should be generated or suppressed. The severity numbers and severity names supported in the OS conform to ITU standards M.3100 X.733 & X.21 and are listed in [Table 58](#).

Table 58 Event Severity Levels

Severity Number	Severity Name
1	cleared
2	indeterminate (info)
3	critical
4	major
5	minor
6	warning

Events that are suppressed by event control will not generate any event log entries. Event control maintains a count of the number of events generated (logged) and dropped (suppressed) for each application event. The severity of an application event can be configured in event control.

An event log within the OS associates the event sources with logging destinations. Examples of logging destinations include, the console session, a specific telnet or SSH session, memory logs, file destinations, SNMP trap groups and syslog destinations. A log filter policy can be associated with the event log to control which events will be logged in the event log based on combinations of application, severity, event ID range, VRF ID, and the subject of the event.

The OS accounting logs collect comprehensive accounting statistics to support a variety of billing models. The routers collect accounting data on services and network ports on a per-service class basis. In addition to gathering information critical for service billing, accounting records can be analyzed to provide insight about customer service trends for potential service revenue opportunities. Accounting statistics on network ports can be used to track link utilization and network traffic pattern trends. This information is valuable for traffic engineering and capacity planning within the network core.

Accounting statistics are collected according to the parameters defined within the context of an accounting policy. Accounting policies are applied to customer Service Access Points (SAPs) and network ports. Accounting statistics are collected by counters for individual service queues defined on the customer's SAP or by the counters within forwarding class (FC) queues defined on the network ports.

The type of record defined within the accounting policy determines where a policy is applied, what statistics are collected and time interval at which to collect statistics.

The supported destination for an accounting log is a compact flash system device. Accounting data is stored within a standard directory structure on the device in compressed XML format. It is recommended that accounting logs be configured on the cf1: or cf2: devices only. Accounting log files are not recommended on the cf3: device (cf3: is intended to be used primarily for software images and configuration related files).

5.3 Log Destinations

Both event logs and accounting logs use a common mechanism for referencing a log destination. Routers support the following log destinations:

- [Console](#)

- [Session](#)
- [Memory Logs](#)
- [Log Files](#)
- [SNMP Trap Group](#)
- [Syslog](#)

Only a single log destination can be associated with an event log or with an accounting log. An event log can be associated with multiple event sources, but it can only have a single log destination.

A file destination is the only type of log destination that can be configured for an accounting log.

5.3.1 Console

Sending events to a console destination means the message will be sent to the system console. The console device can be used as an event log destination.

5.3.2 Session

A session destination is a temporary log destination which directs entries to the active telnet or SSH session for the duration of the session. When the session is terminated, for example, when the user logs out, the “to session” configuration is removed. Event logs configured with a session destination are stored in the configuration file but the “to session” part is not stored. Event logs can direct log entries to the session destination.

5.3.3 Memory Logs

A memory log is a circular buffer. When the log is full, the oldest entry in the log is replaced with the new entry. When a memory log is created, the specific number of entries it can hold can be specified, otherwise it will assume a default size. An event log can send entries to a memory log destination.

5.3.4 Log Files

Log files can be used by both event logs and accounting logs and are stored on the compact flash devices in the file system. It is recommended that event and accounting logs be configured on the cf1: or cf2: devices only. Log files are not recommended on the cf3: device (cf3: is intended to be used primarily for software images and configuration related files).

A log file is identified with a single log file ID, but a log file will generally be composed of a number individual files in the file system. A log file is configured with a rollover parameter, expressed in minutes, which represents the length of time an individual log file should be written to before a new file is created for the relevant log file ID. The rollover time is checked only when an update to the log is performed. Thus, complying to this rule is subject to the incoming rate of the data being logged. For example, if the rate is very low, the actual rollover time may be longer than the configured value.

The retention time for a log file specifies the amount of time the file should be retained on the system based on the creation date and time of the file.

When a log file is created, only the compact flash device for the log file is specified. Log files are created in specific subdirectories with standardized names depending on the type of information stored in the log file.

Event log files are always created in the **\log** directory on the specified compact flash device. The naming convention for event log files is:

log eeff-timestamp

where:

ee is the event log ID

ff is the log file destination ID

timestamp is the timestamp when the file is created in the form of *yyyymmdd-hhmmss* where:

yyyy is the four-digit year (for example, 2007)

mm is the two digit number representing the month (for example, 12 for December)

dd is the two digit number representing the day of the month (for example, 03 for the 3rd of the month)

hh is the two digit hour in a 24-hour clock (for example, 04 for 4 a.m.)

mm is the two digit minute (for example, 30 for 30 minutes past the hour)

ss is the two digit second (for example, 14 for 14)

Accounting log files are created in the **\act-collect** directory on a compact flash device (specifically *cf1* or *cf2*). The naming convention for accounting log files is nearly the same as for log files except the prefix **act** is used instead of the prefix **log**. The naming convention for accounting logs is:

act aa`ff`-timestamp.xml.gz

where:

aa is the accounting policy ID

ff is the log file destination ID

timestamp is the timestamp when the file is created in the form of *yyyymmdd-hhmmss* where:

yyyy is the four-digit year (for example, 2007)

mm is the two digit number representing the month (for example, 12 for December)

dd is the two digit number representing the day of the month (for example, 03 for the 3rd of the month)

hh is the two digit hour in a 24-hour clock (for example, 04 for 4 a.m.)

mm is the two digit minute (for example, 30 for 30 minutes past the hour)

ss is the two digit second (for example, 14 for 14 seconds)

Accounting logs are .xml files created in a compressed format and have a .gz extension.

The **\act-collect** directory is where active accounting logs are written. When an accounting log is rolled over, the active file is closed and archived in the **\act** directory before a new active accounting log file created in **\act-collect**.

When creating a new log file on a Compact Flash disk card, the system will check the amount of free disk space and that amount must be greater than or equal to the lesser of 5.2 MB or 10% of the Compact Flash disk capacity.

5.3.5 SNMP Trap Group

An event log can be configured to send events to SNMP trap receivers by specifying an SNMP trap group destination.

An SNMP trap group can have multiple trap targets. Each trap target can have different operational parameters.

A trap destination has the following properties:

- The IP address of the trap receiver.
- The UDP port used to send the SNMP trap.
- SNMP version (v1, v2c, or v3) used to format the SNMP notification.
- SNMP community name for SNMPv1 and SNMPv2c receivers.
- Security name and level for SNMPv3 trap receivers.

For SNMP traps that will be sent out-of-band through the Management Ethernet port on the SF/CPM, the source IP address of the trap is the IP interface address defined on the Management Ethernet port. For SNMP traps that will be sent in-band, the source IP address of the trap is the system IP address of the router.

Each trap target destination of a trap group receives the identical sequence of events as defined by the log ID and the associated sources and log filter applied.

5.3.6 Syslog

An event log can be configured to send events to one syslog destination. Syslog destinations have the following properties:

- Syslog server IP address.
- The UDP port used to send the syslog message.
- The Syslog Facility Code (0 - 23) (default 23 - local 7).
- The Syslog Severity Threshold (0 - 7) - events exceeding the configured level will be sent.

Because syslog uses eight severity levels whereas the router uses six internal severity levels, the severity levels are mapped to syslog severities. [Table 59](#) displays the severity level mappings to syslog severities.

Table 59 Router to Syslog Severity Level Mappings

SR OS Event Severity	Syslog Severity Numerical Code	Syslog Severity name	Syslog Severity Definition
--	0	emergency	System is unusable
critical (3)	1	alert	Action must be taken immediately
major (4)	2	critical	Critical conditions
minor (5)	3	error	Error conditions
warning (6)	4	warning	Warning conditions
--	5	notice	Normal but significant condition
cleared (1) indeterminate (2)	6	info	Informational messages
--	7	debug	Debug-level messages

The general format of an SR OS syslog message is as follows (see RFC3164). The '<' and '>' are informational delimiters to make reading and understanding the format easier and they do not appear in the actual syslog message except as part of the 'PRI':

<PRI> <HEADER><MSG>

where:

- <PRI> (the "<" and ">" are included in the syslog message) is the configured facility*8+severity (as described in the 7450 ESS, 7750 SR, and 7950 XRS System Management Guide and RFC3164).
- <HEADER> is "MMM DD HH:MM:SS <source IP addr>" (without the quotes). There are always 2 characters for the day (DD). Single digit days are preceded with a space character.
- <MSG> is <log-prefix>: <seq> <router-name> <application>-<severity>-<Event Name>-<Event ID> [<subject>]: <message>\n

where:

- <log-prefix> is an optional 32 characters of text (default = 'TMNX') as configured in the log-prefix command.
- <seq> is the log event sequence number (always preceded by a colon and a space char)
- <router-name> is vprn1, vprn2, ... | Base | management | vpls-management
- <subject> may be empty resulting in []:

- \n is the standard ASCII new line character (hex 0A)

Examples (from different nodes):

default log-prefix (TMNX):

```
<188>Jan 2 18:43:23 135.221.38.108 TMNX: 17 Base SYSTEM-WARNING-tmnxStateChange-2009 [CHASSIS]: Status of Card 1 changed administrative state: inService, operational state: outOfService\n<186>Jan 2 18:43:23 135.221.38.108 TMNX: 18 Base CHASSIS-MAJOR-tmnxEqCardRemoved-2003 [Card 1]: Class IO Module : removed\n
```

no log-prefix:

```
<188>Jan 11 18:48:12 135.221.38.108 : 32 Base SYSTEM-WARNING-tmnxStateChange-2009 [CHASSIS]: Status of Card 1 changed administrative state: inService, operational state: outOfService\n<186>Jan 11 18:48:12 135.221.38.108 : 33 Base CHASSIS-MAJOR-tmnxEqCardRemoved-2003 [Card 1]: Class IO Module : removed\n
```

log-prefix "test":

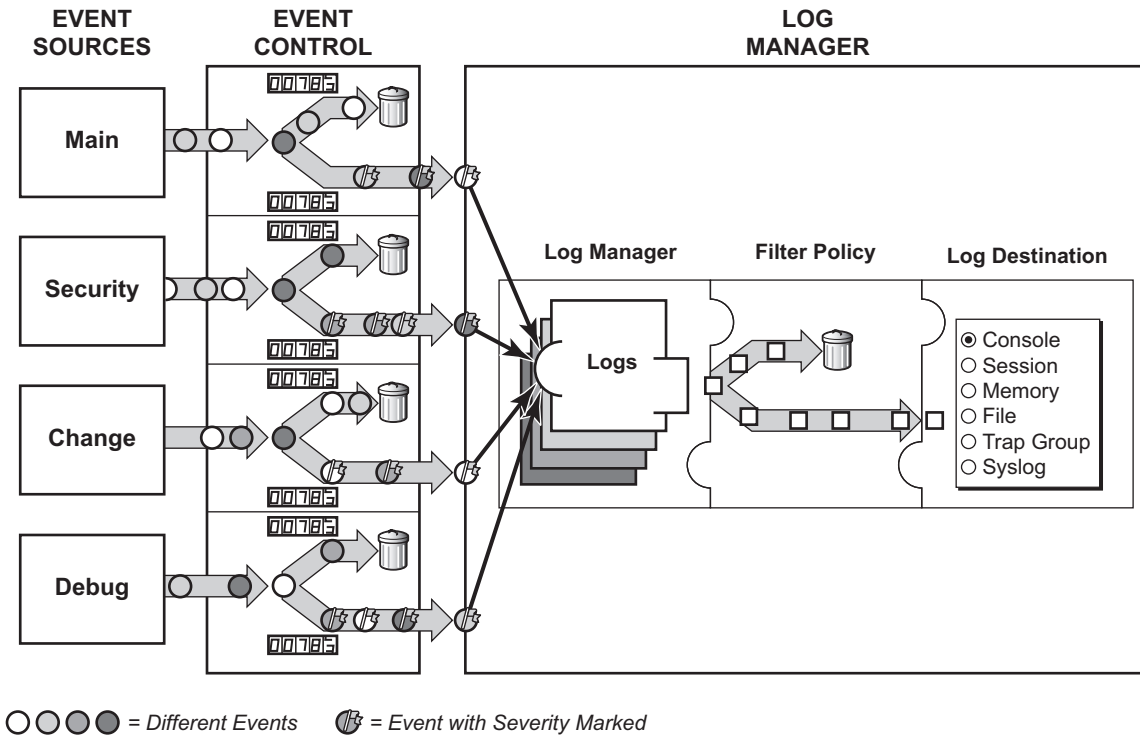
```
<186>Jan 11 18:51:22 135.221.38.108 test: 47 Base CHASSIS-MAJOR-tmnxEqCardRemoved-2003 [Card 1]: Class IO Module : removed\n<188>Jan 11 18:51:22 135.221.38.108 test: 48 Base SYSTEM-WARNING-tmnxStateChange-2009 [CHASSIS]: Status of Card 1 changed administrative state: inService, operational state: outOfService\n
```

5.4 Event Logs

Event logs are the means of recording system generated events for later analysis. Events are messages generated by the system by applications or processes within the router.

[Figure 17](#) depicts a function block diagram of event logging.

Figure 17 Event Logging Block Diagram



CL10001B

5.4.1 Event Sources

In [Figure 17](#), the event sources are the main categories of events that feed the log manager.

- **Security** — The security event source is all events that affect attempts to breach system security such as failed login attempts, attempts to access MIB tables to which the user is not granted access or attempts to enter a branch of the CLI to which access has not been granted. Security events are generated by the SECURITY application and the authenticationFailure event in the SNMP application.
- **Change** — The change activity event source is all events that directly affect the configuration or operation of the node. Change events are generated by the USER application. The Change event stream also includes the tmnxConfigModify (#2006), tmnxConfigCreate (#2007), tmnxConfigDelete (#2008) and tmnxStateChange (#2009) change events from the SYSTEM application.

- Debug — The debug event source is the debugging configuration that has been enabled on the system. Debug events are generated by the DEBUG application.
- Main — The main event source receives events from all other applications within the router.

Examples of applications within the system include IP, MPLS, OSPF, CLI, services, etc. The following example displays a partial sample of the **show log applications** command output which displays all applications.

```
*A:ALA-48# show log applications
=====
Log Event Application Names
=====
Application Name
-----
...
BGP
CCAG
CFLOWD
CHASSIS
...
MPLS
MSDP
NTP
...
USER
VRRP
VRTR
=====
*A:ALA-48#
```

5.4.2 Event Control

Event control pre-processes the events generated by applications before the event is passed into the main event stream. Event control assigns a severity to application events and can either forward the event to the main event source or suppress the event. Suppressed events are counted in event control, but these events will not generate log entries as it never reaches the log manager.

Simple event throttling is another method of event control and is configured similarly to the generation and suppression options. See [Simple Logger Event Throttling](#).

Events are assigned a default severity level in the system, but the application event severities can be changed by the user.

Application events contain an event number and description that explains why the event is generated. The event number is unique within an application, but the number can be duplicated in other applications.

The following example, generated by querying event control for application generated events, displays a partial list of event numbers and names.

```
router# show log event-control
=====
Log Events
=====
Application
ID#      Event Name                               P  g/s   Logged   Dropped
-----
show
BGP:
  2001  bgpEstablished                          MI  gen    1         0
  2002  bgpBackwardTransition                    WA  gen    7         0
  2003  tBgpMaxPrefix90                          WA  gen    0         0
...
CCAG:
CFLOWD:
  2001  cflowdCreated                            MI  gen    1         0
  2002  cflowdCreateFailure                      MA  gen    0         0
  2003  cflowdDeleted                            MI  gen    0         0
...
CHASSIS:
  2001  cardFailure                              MA  gen    0         0
  2002  cardInserted                             MI  gen    4         0
  2003  cardRemoved                              MI  gen    0         0
...
'''
DEBUG:
L 2001  traceEvent                              MI  gen    0         0
DOT1X:
FILTER:
  2001  filterPBRPacketsDropped                 MI  gen    0         0
IGMP:
  2001  vRtrIcmpIfRxQueryVerMismatch            WA  gen    0         0
  2002  vRtrIcmpIfCModeRxQueryMismatch          WA  gen    0         0
IGMP_SNOOPING:
IP:
L 2001  clearRTMError                           MI  gen    0         0
L 2002  ipEtherBroadcast                         MI  gen    0         0
L 2003  ipDuplicateAddress                       MI  gen    0         0
...
ISIS:
  2001  vRtrIcmpDatabaseOverload                WA  gen    0         0
```

5.4.3 Log Manager and Event Logs

Events that are forwarded by event control are sent to the log manager. The log manager manages the event logs in the system and the relationships between the log sources, event logs and log destinations, and log filter policies.

An event log has the following properties:

- A unique log ID
The log ID is a short, numeric identifier for the event log. A maximum of 15 logs can be configured at a time.
- One or more log sources
The source stream or streams to be sent to log destinations can be specified. The source must be identified before the destination can be specified. The events can be from the main event stream, events in the security event stream, or events in the user activity stream.
- One event log destination
A log can only have a single destination. The destination for the log ID destination can be one of console, session, syslog, snmp-trap-group, memory, or a file on the local file system.
- An optional event filter policy
An event filter policy defines whether to forward or drop an event or trap-based on match criteria.

5.4.4 Event Filter Policies

The log manager uses event filter policies to allow fine control over which events are forwarded or dropped based on various criteria. Like other filter policies in the SR OS, filter policies have a default action. The default actions are either:

- Forward
- Drop

Filter policies also include a number of filter policy entries that are identified with an entry ID and define specific match criteria and a forward or drop action for the match criteria.

Each entry contains a combination of matching criteria that define the application, event number, router, severity, and subject conditions. The entry's action determines how the packets should be treated if they have met the match criteria.

Entries are evaluated in order from the lowest to the highest entry ID. The first matching event is subject to the forward or drop action for that entry.

Valid operators are displayed in [Table 60](#):

Table 60 Valid Filter Policy Operators

Operator	Description
eq	equal to
neq	not equal to
lt	less than
lte	less than or equal to
gt	greater than
gte	greater than or equal to

A match criteria entry can include combinations of:

- Equal to or not equal to a given system application.
- Equal to, not equal to, less than, less than or equal to, greater than or greater than or equal to an event number within the application.
- Equal to, not equal to, less than, less than or equal to, greater than or greater than or equal to a severity level.
- Equal to or not equal to a router name string or regular expression match.
- Equal to or not equal to an event subject string or regular expression match.

5.4.5 Event Log Entries

Log entries that are forwarded to a destination are formatted in a way appropriate for the specific destination whether it be recorded to a file or sent as an SNMP trap, but log event entries have common elements or properties. All application generated events have the following properties:

- A time stamp in UTC or local time.
- The generating application.
- A unique event ID within the application.
- A router name identifying the VRF-ID that generated the event.
- A subject identifying the affected object.
- A short text description.

The general format for an event in an event log with either a memory, console or file destination is as follows.

```
nnnn YYYY/MM/DD HH:MM:SS.SS TZONE <severity>: <application> #<event_id> <router-
name>
<subject>
<message>
```

The following is an event log example:

```
252 2013/05/07 16:21:00.76 UTC WARNING: SNMP #2005 Base my-interface-abc
"Interface my-interface-abc is operational"
```

The specific elements that compose the general format are described in [Table 61](#).

Table 61 Log Entry Field Descriptions

Label	Description
nnnn	The log entry sequence number.
YYYY/MM/DD	The UTC date stamp for the log entry. YYYY — Year MM — Month DD — Date
HH:MM:SS.SS	The UTC time stamp for the event. HH — Hours (24 hour format) MM — Minutes SS.SS — Seconds
TZONE	The timezone (for example, UTC, EDT) as configured by configure log log-id x time-format .
<severity>	The severity level name of the event. CLEARED — A cleared event (severity number 1). INFO — An indeterminate/informational severity event (severity level 2). CRITICAL — A critical severity event (severity level 3). MAJOR — A major severity event (severity level 4). MINOR — A minor severity event (severity level 5). WARNING — A warning severity event (severity 6).
<application>	The application generating the log message.
<event_id>	The application's event ID number for the event.
<router>	The router name representing the VRF-ID that generated the event.
<subject>	The subject/affected object for the event.
<message>	A text description of the event.

5.4.6 Simple Logger Event Throttling

Simple event throttling provides a mechanism to protect event receivers from being overloaded when a scenario causes many events to be generated in a very short period of time. A throttling rate, # events/# seconds, can be configured. Specific event types can be configured to be throttled. Once the throttling event limit is exceeded in a throttling interval, any further events of that type cause the dropped events counter to be incremented. Dropped events counts are displayed by the **show>log>event-control** context. Events are dropped before being sent to one of the logger event collector tasks. There is no record of the details of the dropped events and therefore no way to retrieve event history data lost by this throttling method.

A particular event type can be generated by multiple managed objects within the system. At the point this throttling method is applied the logger application has no information about the managed object that generated the event and cannot distinguish between events generated by object "A" from events generated by object "B". If the events have the same event-id, they are throttled regardless of the managed object that generated them. It also does not know which events may eventually be logged to destination log-id <n> from events that will be logged to destination log-id <m>.

Throttle rate applies commonly to all event types. It is not configurable for a specific event-type.

A timer task checks for events dropped by throttling when the throttle interval expires. If any events have been dropped, a TIMETRA-SYSTEM-MIB::tmnxTrapDropped notification is sent.

5.4.7 Default System Log

Log 99 is a pre-configured memory-based log which logs events from the main event source (not security, debug, etc.). Log 99 exists by default.

The following example displays the log 99 configuration.

```
ALA-1>config>log# info detail
#-----
echo "Log Configuration "
#-----
...
    snmp-trap-group 7
    exit
...
    log-id 99
```



```
description "Default system log"
no filter
from main
to memory 500
no shutdown
exit
-----
ALA-1>config>log#
```

5.4.8 Event Handling System

The Event Handling System (EHS) is a tool that allows operator-defined behavior to be configured on the router. EHS adds user-controlled programmatic exception handling by allowing a CLI script to be executed upon the detection of a log event (the 'trigger'). Regexp style expression matching is available on various fields in the log event to give flexibility in the trigger definition.

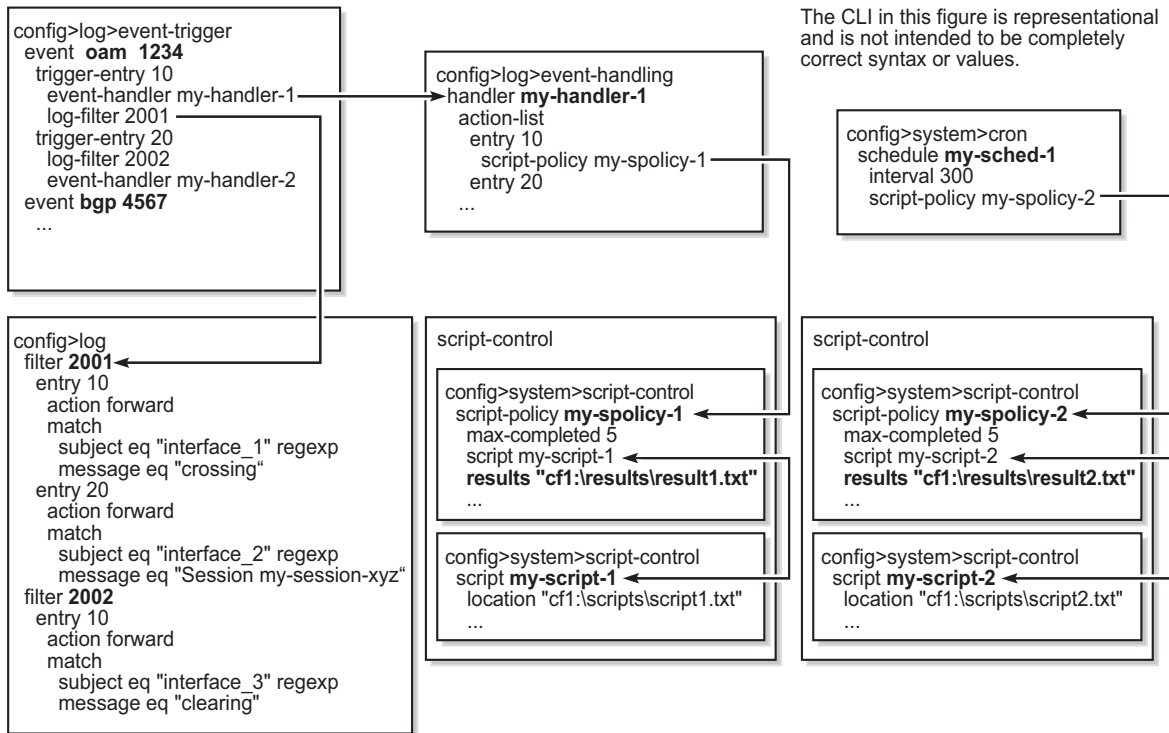
EHS handler objects are used to tie together:

- trigger events (typically log events that match some configurable criteria)
- a set of actions to perform (typically one or more CLI scripts)

EHS, along with CRON, makes use of the generic SR OS CLI script-control functions for scripts. Any command available in CLI (with some limited exceptions such as 'candidate' commands) can be executed in a script as the result of an EHS handler being triggered.

The following figure illustrates the relationships between the different configurable objects used by EHS (and CRON).

Figure 18 EHS Object Relationships



The CLI in this figure is representational and is not intended to be completely correct syntax or values.

24884

Complex rules can be configured to match on log events as a trigger for an EHS handler.

When a log event is generated in SR OS it will be subject to discard via suppression and throttling (**config>log>event-control**) before it is evaluated as a trigger for EHS:

- EHS will not trigger on log events that are suppressed through **config>log>event-control**
- EHS will not trigger on log events that are throttled by the logger

EHS will trigger on log events that are dropped by user configured log filters that are assigned to individual logs (**config>log>filter**). The EHS event trigger logic occurs before the distribution of log event streams into individual logs.

A triggering log event's common parameters and varbinds are passed in to the triggered EHS script and can be used within the EHS script as passed in (dynamic) variables. Passed in (dynamic) variables are:

- the common event parameters, such as, severity, subject, appid, eventid, gentime, and so on.
- the predefined varbinds in a log event's message.

For example, the following are the passed in (dynamic) variables for an event:

- appid
- eventid
- severity
- subject
- gentime
- event_varbind_1
- event_varbind_2
- ...
- ...
- event_varbind_N



Note:

- For more information about showing event parameters, see the show commands in the "Log Command Reference" section.
- See the *7750 SR Log Events Guide* for any event's predefined varbinds
- The passed in event's **gentime** is always UTC
- The event's sequence number is not passed in to the script

An EHS script has the ability to define local (static) variables and use some basic .if and .set commands inside the script. The use of variables with .if and .set commands within an EHS script adds more logic to EHS scripting and allows the reuse of a single EHS script for more than one trigger or action.

Both imported and local variables can be used within the EHS script either as part of the CLI commands or as part of the .if or .set commands.

The following applies to both CLI commands and .if or .set commands.

- Using \$X, without using single or double quotes, replaces the variable X with its string or integer value.
- Using "X", with double quotes, means the literal string X.
- Using "\$X", with double quotes, replaces the variable X with its string or integer value.
- Using 'X', with single quotes, means the literal string X.
- Using '\$X', with single quotes, does not replace the variable X with its value but means the literal string \$X.

In summary:

- All characters within single quotes are interpreted as a string character.
- All characters within double quotes are interpreted as regular characters except for \$, which replaces the variable with its value (for example, shell expansion inside a string).

Some supported shell command scenarios are (the following are pseudo commands):

- `.if $string_variable==string_value_or_string_variable {`
 `CLI_commands_set1`
 `.} else {`
 `CLI_commands_set2`
 `.} endif`
- `.if ($string_variable==string_value_or_string_variable) {`
 `CLI_commands_set1`
 `.} else {`
 `CLI_commands_set2`
 `.} endif`
- `.if $integer_variable==integer_value_or_integer_variable {`
 `CLI_commands_set1`
 `.} else {`
 `CLI_commands_set2`
 `.} endif`
- `.if ($integer_variable==integer_value_or_integer_variable) {`
 `CLI_commands_set1`
 `.} else {`
 `CLI_commands_set2`
 `.} endif`
- `.if $string_variable!=string_value_or_string_variable {`
 `CLI_commands_set1`
 `.} else {`
 `CLI_commands_set2`
 `.} endif`
- `.if ($string_variable!=string_value_or_string_variable) {`
 `CLI_commands_set1`
 `.} else {`
 `CLI_commands_set2`

```
.} endif
• .if $integer_variable!=integer_value_or_integer_variable {
    CLI_commands_set1
.} else {
    CLI_commands_set2
.} endif
• .if ($integer_variable!=integer_value_or_integer_variable) {
    CLI_commands_set1
.} else {
    CLI_commands_set2
.} endif
• .set $string_variable = string_value_or_string_variable
• .set ($string_variable = string_value_or_string_variable)
• .set $integer_variable = integer_value_or_integer_variable
• .set ($integer_variable = integer_value_or_integer_variable)
```

where:

- *CLI_commands_set1* is a set of one or more CLI commands
- *CLI_commands_set2* is a set of one or more CLI commands
- *string_variable* is a local (static) string variable
- *string_value_or_string_variable* is a string value/variable
- *integer_variable* is a local (static) integer variable
- *integer_value_or_integer_variable* is an integer value/variable

**Note:**

- A limit of 100 local (static) variables per EHS script is imposed. Exceeding this limit may result in an error and partial execution of the script.
- When a set statement is used to set a `string_variable` to a `string_value`, the `string_value` can be any non-integer value not surrounded by single/double quotes or it can be surrounded by single/double quotes
- A "." preceding a directive (e.g. `if`, `set`...etc) is always expected to start a new line
- An end of line is always expected after `{`
- A CLI command is always expected to start a new line
- Passed in (dynamic) variables are always read only inside an EHS script and cannot be overwritten using a set statement
- `.if` commands support `==` and `!=` operators only
- `.if` and `.set` commands support addition, subtraction, multiplication, and division of integers
- `.if` and `.set` commands support addition of strings which means "concatenation" of strings

Valid Examples:

- `configure service epipe $serviceID`
where `$serviceID` is either a local (static) integer variable or passed in (dynamic) integer variable
- `echo srcAddr is $srcAddr`
where `$srcAddr` is a passed in (dynamic) string variable
- `.set $ipAddr = "10.0.0.1"`
where `$ipAddr` is a local (static) string variable
- `.set $ipAddr = $srcAddr`
where `$srcAddr` is a passed in (dynamic) string variable
`$ipAddr` is a local (static) string variable.
- `.set ($customerID = 50)`
where `$customerID` is a local (static) integer variable
- `.set ($totalPackets = $numIngrPackets + $numEgrPackets)`
where `$totalPackets`, `$numIngrPackets`, `$numEgrPackets` are local (static) integer variables
- `.set ($portDescription = $portName + $portLocation)`
where `$portDescription`, `$portName`, `$portLocation` are local (static) string variables

- `if ($srcAddr == "CONSOLE") {`
 `CLI_commands_set1`
• `.else {`
 `CLI_commands_set2`
• `.} endif`
 where `$srcAddr` is a passed in (dynamic) string variable
 `CLI_commands_set1` is a set of one or more CLI commands
 `CLI_commands_set2` is a set of one or more CLI commands
- `.if ($customerID == 10) {`
 `CLI_commands_set1`
• `.else {`
 `CLI_commands_set2`
• `.} endif`
 where `$customerID` is a passed in (dynamic) integer variable
 `CLI_commands_set1` is a set of one or more CLI commands
 `CLI_commands_set2` is a set of one or more CLI commands
- `.if ($numIngrPackets == $numEgrPackets) {`
 `CLI_commands_set1`
• `.else {`
 `CLI_commands_set2`
• `.} endif`
 where `$numIngrPackets` and `$numEgrPackets` are local (static) integer variables
 `CLI_commands_set1` is a set of one or more CLI commands
 `CLI_commands_set2` is a set of one or more CLI commands

Invalid Examples:

- `.set $srcAddr = "10.0.0.1"`
 where `$srcAddr` is a passed in (dynamic) string variable
 Reason: passed in variables are read only inside an EHS script.
- `.set ($ipAddr = '$numIngrPackets' + $numEgrPackets)`
 where `$ipAddr` is a local (static) string variable
 `$numIngrPackets` and `$numEgrPackets` are local (static) integer variables
 Reason: variable types do not match, cannot assign a string to an integer.
- `.set ($numIngrPackets = $ipAddr + $numEgrPackets)`

where *\$ipAddr* is a local (static) string variable

\$numIngrPackets and *\$numEgrPackets* are local (static) integer variables

Reason: variable types do not match, cannot concatenate a string to an integer.

- `.set $ipAddr = "10.0.0.1"100`

where *\$ipAddr* is a local (static) string variable

Reason: when double quotes are used, they have to surround the entire string.

- `.if ($totalPackets == "10.1.1.1") {
 .} endif`

where *\$totalPackets* is a local (static) integer variables

Reason: cannot compare an integer variable to a string value.

- `.if ($ipAddr == 10) {
 .} endif`

where *\$ipAddr* is a local (static) string variable

Reason: cannot compare a string variable to an integer value.

- `.if ($totalPackets == $ipAddr) {`

where *\$totalPackets* is a local (static) integer variables

\$ipAddr is a local (static) string variable

Reason: cannot compare an integer variable to a string variable.

EHS debounce

EHS bounce is the ability to trigger an action (for example an EHS script), if an event happens (N) times within a specific time window (S).

N = [2..15]

S = [1..604800]



Note:

- Triggering happens with the Nth event not at the end of S
- There is no sliding window (for example a trigger at Nth event, N+1 event, and so on), as N is reset after a trigger and count is restarted
- When EHS debouncing/dampening is used, the varbinds passed in to an EHS script at script triggering time are from the Nth event occurrence (the Nth triggering event)
- If S is not specified then the SR OS will continue to trigger every Nth event

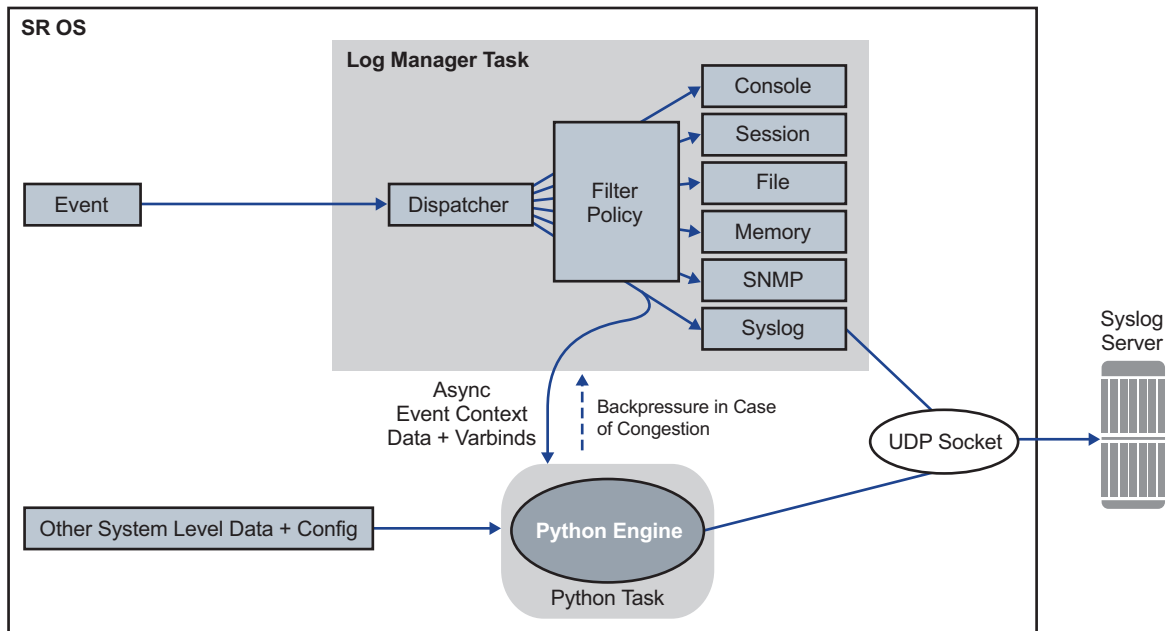
Example:

When linkDown occurs N times in S sec, an EHS script is triggered to shut down the port.

5.5 Customizing Syslog Messages Using Python

Log events in SR OS can be customized by a Python script before they are sent to a syslog server. The log events that are subject to Python processing are selected via log filters. This allows only a preferred subset of log messages to be customized (Figure 19).

Figure 19 Interaction between the Logger and the Python Engine



sw0029

5.5.1 Python Engine for Syslog

This section discusses syslog-specific aspects of Python processing. Refer to the “Python Script Support for ESM” section of the *7450 ESS, 7750 SR, and 7950 XRS Triple Play Guide* for an introduction to Python.

When an event is dispatched to the log manager in SR OS, the log manager asynchronously passes the event context data and varbinds to the Python engine, that is, the logger task is not waiting for feedback from Python. Varbinds are variable bindings that represent the variable number of values that are included in the event. Each varbind consists of a triplet (OID, type, value). Along with other system-level variables, the Python engine constructs a syslog message and sends it to the syslog destination. During this process, the operator can modify the format of the syslog message or leave it intact, as if it was generated by the syslog process within the log manager.

The tasks of the Python engine in a syslog context are as follows:

- assembles custom syslog messages (including PRI, HEADER and MSG fields) based on the received event context data, varbinds specific to the event, system-level data, and the configuration parameters (syslog server IP address, syslog facility, log-prefix and the destination UDP port)
- reformats timestamps in a syslog message
- sends the original or modified message to the syslog server
- drops the message

5.5.1.1 Python Syslog APIs

Python APIs are used to assemble a syslog message which, in SR OS, has the following generic format:

```
PRI> <HEADER><MSG>
```

where:

- **<PRI>** (the “<” and “>” are included in the syslog message) is the configured facility x 8+severity (as described in the *7450 ESS*, *7750 SR*, and *7950 XRS System Management Guide* and RFC 3164, *The BSD syslog Protocol*)
- **<HEADER>** is MMM DD HH:MM:SS <hostname>. There are always two characters for the day (DD). Single digit days are preceded with a space character.
- **<MSG>** is <log-prefix>: <seq> <router-name> <application>-<severity>-<Event Name>-<Event ID> [<subject>]: <message>\n

where:

- *<log-prefix>* is an optional set of 32 characters (default = 'TMNX') as configured in the **log-prefix** command

- `<seq>` is the log event sequence number. It always preceded by a colon and a space character.
- `<router-name>` is the name of the router, for example, vprn1, vprn2, Base, management, vpls-management
- `<subject>` is the topic and can be empty, resulting in []:
- `\n` is the standard ASCII new line character (hex 0A)

Table 62 describes Python information that can be used to manipulate syslog messages.

Table 62 Manipulating Python Syslog Messages

Imported Nokia (ALC) Modules	Access Rights	Comments
event (from alc import event)	—	The method used to retrieve generic event information.
syslog (from alc import syslog)	—	The method used to retrieve syslog-specific parameters.
system (from alc import system)	—	The method used to retrieve system-specific information. Currently, the only parameter retrieved is the system name.
<p>Events use the following format as they are written into memory, file, console, and system: nnnn YYYY/MM/DD HH:MM:SS.SS <severity>:<application> # <event_id> <router-name> <subject> <message></p> <p>The event-related information received in the context data from the log manager is retrieved via the following Python methods:</p>		
event.sequence	RO	The sequence number of the event (nnnn).
event.timestamp	RO	The timestamp of the event. (YYYY/MM/DD HH:MM:SS.SS).
event.routerName	RO	The router name, for example, BASE, VPRN1, and so on.
event.application	RO	The application generating the event, for example, NA.
event.severity	RO	The severity of the event. This is configurable in SR OS (CLEARED [1], INFO [2], CRITICAL [3], MAJOR [4], MINOR [5], WARNING [6]).

Table 62 Manipulating Python Syslog Messages (Continued)

Imported Nokia (ALC) Modules	Access Rights	Comments
event.eventId	RO	The event ID, for example, 2012.
event.eventName	RO	The event Name, for example, tmnxNatPIBlockAllocationLsn.
event.subject	RO	An optional field, for example, [NAT].
event.message	RO	The event-specific message, for example, "{2} Map 192.168.20.29 [2001-2005] MDA 1/2 -- 276824064 classic-lsn-sub %3 vprn1 10.10.10.101 at 2015/08/31 09:20:15".
Syslog Methods		
syslog.hostName	RO	The IP address of the SR OS node sending the syslog message. This is used in the Syslog HEADER.
syslog.logPrefix	RO	The log prefix which is configurable and optional, for example, TMNX:
syslog.severityToPRI(event.severity)	—	The Python method used to derive the PRI field in syslog header based on event severity and a configurable syslog facility.
syslog.severityToName(event.severity)	—	An SR OS event severity to syslog severity name. For more information, see the 5.3.6 section.
syslog.timestampToUnix(timestamp)		The Python method that takes a timestamp in the format if YYYY/MM/DD HH:MM:SS and converts it into a UNIX-based format (seconds since Jan 01 1970 – UTC).

Table 62 Manipulating Python Syslog Messages (Continued)

Imported Nokia (ALC) Modules	Access Rights	Comments
syslog.set(newSyslogPdu)	—	The Python method used to send the syslog message in the newSyslogPdu. This variable must be constructed manually via string manipulation. In the absence of the command, the SR OS assembles the default syslog message (as if Python was not configured) and sends it to the syslog server, assuming that the message is not explicitly dropped.
syslog.drop()	—	The Python method used to drop a syslog message. This method must be called before the syslog.set<newSyslogPdu method.
System Methods		
system.name	RO	The Python method used to retrieve the system name

For example, assume that the syslog format is:

```
<PRI><timestamp> <hostname> <log-prefix>: <sequence> <router-name> <appid>-
<severity>-<name>-<eventId> [<subject>]: <text>
```

Then the following is an example of the syslogPdu constructed via Python:

```
syslogPdu = "<" + syslog.severityToPRI(event.severity) + ">" \
+ event.timestamp + " " \
+ syslog.hostname + " " + syslog.logPrefix + ": " + \
event.sequence + " " + event.routerName + " " + \
event.application + "-" + \
syslog.severityToName(event.severity) + "-" + \
event.eventName + "-" + event.eventId + " [" + \
event.subject + "]: " + event.message
```

5.5.1.2 Timestamp Format Manipulation

Certain logging environments require customized formatting of the timestamp. Nokia provides a timestamp conversion method in the alu.syslog Python module to convert a timestamp from the format YYYY/MM/DD hh:mm:ss into a UNIX-based timestamp format (seconds since Jan 01 1970 – UTC).

For example, an operator can use the following Python method to convert a timestamp from the YYYY/MM/DD hh:mm:ss.ss or YYYY/MM/DD hh:mm:ss (no centiseconds) format into either the UNIX timestamp format or the MMM DD hh:mm:ss format.

```

from alc import event
from alc import syslog
from alc import system
#input format: YYYY/MM/DD hh:mm:ss.ss or YYYY/MM/DD hh:mm:ss
#output format 1: MMM DD hh:mm:ss
#output format 2: unixTimestamp (TBD)
def timeFormatConversion(timestamp,format):
    if format not in range(1,2):
        raise NameError('Unexpected format, expected:' \
            '0<format<3 got: '+str(format))
    try:
        dat,tim=timestamp.split(' ')
    except:
        raise NameError('Unexpected timestamp format, expected:' \
            'YYYY/MM/DD hh:mm:ss got: '+timestamp)
    try:
        YYYY,MM,DD=dat.split('/')
    except:
        raise NameError('Unexpected timestamp format, expected:' \
            'YYYY/MM/DD hh:mm:ss got: '+timestamp)
    try:
        hh,mm,ss=tim.split(':')
        ss=ss.split('.')[0] #just in case that the time format is hh:mm:ss.ss
    except:
        raise NameError('Unexpected timestamp format, expected:' \
            'YYYY/MM/DD hh:mm:ss got: '+timestamp)
    if not (1970<=int(YYYY)<2100 and
        1<=int(MM)<=12 and
        1<=int(DD)<=31 and
        0<=int(hh)<=24 and
        0<=int(mm)<=60 and
        0<=int(ss)<=60):
        raise NameError('Unexpected timestamp format, or values out of the range' \
            'Expected: YYYY/MM/DD hh:mm:ss got: '+timestamp)
    if format == 1:
        MMM={1:'Jan',
            2:'Feb',
            3:'Mar',
            4:'Apr',
            5:'May',
            6:'Jun',
            7:'Jul',
            8:'Aug',
            9:'Sep',
            10:'Oct',
            11:'Nov',
            12:'Dec'}[int(MM)]
        timestamp=MMM+' '+DD+' '+hh+':'+mm+':'+ss
    if format == 2:
        timestamp=syslog.timestampToUnix(timestamp)
    return timestamp

```

The timeFormatConversion method can accept the event.timestamp value in the format:

```
YYYY/MM/DD HH:MM:SS.SS
```

and return a new timestamp in the format determined by the format parameter:

```
1 ? MMM DD HH:MM:SS  
2 ? Unix based time format
```

This method accepts the input format in either of the two forms YYYY/MM/DD HH:MM:SS.SS or YYYY/MM/DD HH:MM:SS and simply ignores the centisecond part in the former form.

5.5.2 Python Processing Efficiency

Python retrieves event-related variables from the log manager, as opposed to retrieving pre-assembled syslog messages. This eliminates the need for string parsing of the syslog message to manipulate its constituent parts, increasing the speed of Python processing.

To further improve processing performance, Nokia recommends performing string manipulation via the Python native string method, when possible.

5.5.3 Python Backpressure

A Python task assembles syslog messages based on the context information received from the logger and sends them to the syslog server independent of the logger. If the Python task is congested due to a high volume of received data, the backpressure should be sent to the ISA so that the ISA stops allocating NAT resources. This behavior matches the current behavior in which NAT resource allocation is blocked if that logger is congested.

5.5.4 Event Selection for Python Processing

Events destined for Python processing are configured through a log ID that references a Python policy. The selection of the events are performed via a filter associated with this log ID. The remainder of the events destined to the same syslog server can bypass Python processing by redirecting them to a different log ID. The following example clarifies this point:

1. Creating the Python policy

```
A:dut-a# configure python python-policy PyForLogEvents create
*A:dut-a>config>python>py-policy$
[no] description      - Configure the description of this policy
[no] dhcp             - Configure scripts to handle dhcp messages jritter
[no] dhcp6           - Configure scripts to handle dhcp6 messages
[no] diameter        - Configure scripts to handle diameter messages
[no] gtpv1-c         - Configure scripts to handle GTPv1-C messages
[no] gtpv2-c         - Configure scripts to handle GTPv2-C messages
[no] pppoe           - Configure scripts to handle PPPoE messages
[no] radius          - Configure scripts to handle RADIUS messages
[no] vsd             - Configure scripts to handle VSD messages
[no] syslog          - Configure a script to handle outgoing syslog messages
*A:dut-a>config>python>py-policy$ syslog
- syslog script <name>
- no syslog
<name> :[32 chars max]
```

The detailed Python policy description is explained in the “Python Script Support for ESM” section in the *7450 ESS, 7750 SR, and 7950 XRS Triple Play Guide*.

2. Log filters identify the events that are subject to Python processing

```
A:dut-a>config>log# info
-----
      filter 6
        default-action drop
        entry 1
          action forward
          match
            application eq "nat"
            number eq 2012
          exit
        exit
      exit
      filter 7
        default-action forward
        entry 1
          action drop
          match
            application eq "nat"
            number eq 2012
          exit
        exit
      exit
```

3. Syslog destination

```
syslog 1
    address 192.168.1.1
exit
```

4. Applying Python syslog policy to selected events via filter 6:

```
log-id 33 Note: Process log events with id of 2012 with Python before
sending them to syslog server.
    filter 6
    from main
    to syslog 1
    python-policy "PyForLogEvents"
    no shutdown
exit
log-id 34 Note: Log events that are not processed by Python.
    filter 7
    from main
    to syslog 1
    no shutdown
exit
```

In the example above, the configuration-only event 2012 from application "nat" will be sent to log-id 33. All other events are forwarded to the same syslog destination via log-id 34, without any modification. As a result, all events (modified via log-id 33 and unmodified via log-id 34) are sent to the syslog 1 destination.

This configuration may cause reordering of syslog messages at the syslog 1 destination due to slight delay of messages processed by Python.

5.5.5 Modifying a Log File

The following displays the current log configuration:

```
ALA-12>config>log>log-id# info
-----
...
log-id 2
      description "This is a test log file."
      filter 1
      from main security
      to file 1

exit
...
-----
ALA-12>config>log>log-id#
```

The following displays an example to modify log file parameters:

```
Example:config# log
config>log# log-id 2
config>log>log-id# description "Chassis log file."
config>log>log-id# filter 2
config>log>log-id# from security
config>log>log-id# exit
```

The following displays the modified log file configuration:

```
A:ALA-12>config>log# info
-----
...
log-id 2
      description "Chassis log file."
      filter 2
      from security
      to file 1

exit
...
-----
A:ALA-12>config>log#
```

5.5.6 Deleting a Log File

The log ID must be shutdown first before it can be deleted. In a previous example, **file 1** is associated with **log-id 2**.

```
A:ALA-12>config>log# info
-----
file-id 1
      description "LocationTest."
```

```

        location cfl:
        rollover 600 retention 24
    exit
...
log-id 2
        description "Chassis log file."
        filter 2
        from security
        to file 1
exit
...
-----
A:ALA-12>config>log#

```

The following displays an example to delete a log file:

```

Example:config# log
config>log# log-id 2
config>log>log-id# shutdown
config>log>log-id# exit
config>log# no log-id 2

```

5.5.7 Modifying a File ID

The following displays the current log configuration:

```

A:ALA-12>config>log# info
-----
        file-id 1
        description "This is a log file."
        location cfl:
        rollover 600 retention 24
    exit
-----
A:ALA-12>config>log#

```

The following displays an example to modify log file parameters:

```

Example:config# log
config>log# file-id 1
config>log>file-id# description "LocationTest."
config>log>file-id# rollover 2880 retention 500
config>log>file-id# exit

```

The following displays the file modifications:

```

A:ALA-12>config>log# info
-----
...

```

```
file-id 1
  description "LocationTest."
  rollover 2880 retention 500
exit
...
-----
A:ALA-12>config>log#
```

The following displays an example to modify log file parameters:

```
Example:config# log
config>log# file-id 1
config>log>file-id# description "LocationTest."
config>log>file-id# location cf2:
config>log>file-id# rollover 2880 retention 500
config>log>file-id# exit
```

The following displays the file modifications:

```
A:ALA-12>config>log# info
-----
...
file-id 1
  description "LocationTest."
  location cf2:
  rollover 2880 retention 500
  exit
...
-----
A:ALA-12>config>log#
```

5.5.8 Modifying a Syslog ID

The following displays an example of the syslog ID modifications:

```
Example:config# log
config>log# syslog 1
config>log>syslog$ description "Test syslog."
config>log>syslog# address 10.10.0.91
config>log>syslog# facility mail
config>log>syslog# level info
```

The following displays the syslog configuration:

```
A:ALA-12>config>log# info
-----
...
  syslog 1
    description "Test syslog."
```

```

        address 10.10.10.91
        facility mail
        level info
    exit
    ...
-----
A:ALA-12>config>log#

```

5.5.9 Modifying an SNMP Trap Group

The following displays the current SNMP trap group configuration:

```

A:ALA-12>config>log# info
-----
...
snmp-trap-group 10
    trap-target 10.10.10.104:5 "snmpv3" notify-community "ccommunitystring"
exit
...
-----
A:ALA-12>config>log#

```

The following displays an example of the command usage to modify an SNMP trap group:

```

Example:config# log
config>log# snmp-trap-group 10
config>log>snmp-trap-group# no trap-target
10.10.10.104:5
config>log>snmp-trap-group# snmp-trap-group# trap-
target 10.10.0.91:1 snmpv2c notify-community "com1"

```

The following displays the SNMP trap group configuration:

```

A:ALA-12>config>log# info
-----
...
    snmp-trap-group 10
        trap-target 10.10.0.91:1 "snmpv2c" notify-community "com1"
    exit
...
-----
A:ALA-12>config>log#

```

5.5.10 Deleting an SNMP Trap Group

The following displays the SNMP trap group configuration:

```
A:ALA-12>config>log# info
-----
...
    snmp-trap-group 10
        trap-target 10.10.0.91:1 "snmpv2c" notify-community "com1"
    exit
...
-----
A:ALA-12>config>log#
```

The following displays an example to delete a trap target and an SNMP trap group.

```
Example:config>log# snmp-trap-group 10
config>log>snmp-trap-group# no trap-target 10.10.0.91:1
config>log>snmp-trap-group# exit
config>log# no snmp-trap-group 10
```

5.5.11 Modifying a Log Filter

The following output displays the current log filter configuration:

```
ALA-12>config>log# info
#-----
echo "Log Configuration "
#-----
...
    filter 1
        default-action drop
        description "This is a sample filter."
        entry 1
            action forward
            match
                application eq "mirror"
                severity eq critical
            exit
        exit
    exit
...
-----
ALA-12>config>log#
```

The following displays an example of the log filter modifications:

```
Example:config# log
config>log# filter 1
config>log>filter# description "This allows <n>."
config>log>filter# default-action forward
config>log>filter# entry 1
config>log>filter>entry$ action drop
```

```

config>log>filter>entry# match
config>log>filter>entry>match# application eq user
config>log>filter>entry>match# number eq 2001
config>log>filter>entry>match# no severity
config>log>filter>entry>match# exit

```

The following displays the log filter configuration:

```

A:ALA-12>config>log>filter# info
-----
...
    filter 1
      description "This allows <n>."
      entry 1
        action drop
        match
          application eq "user"
          number eq 2001
        exit
      exit
    exit
  ...
-----
A:ALA-12>config>log>filter#

```

5.5.12 Modifying Event Control Parameters

The following displays the current event control configuration:

```

A:ALA-12>config>log# info
-----
...
event-control "bgp" 2014 generate critical
...
-----
A:ALA-12>config>log#

```

The following displays an example of an event control modification:

```

Example:config# log
config>log# event-control bgp 2014 suppress

```

The following displays the log filter configuration:

```

A:ALA-12>config>log# info
-----
...
    event-control "bgp" 2014 suppress
  ...
-----

```



```
A:ALA-12>config>log#
```

The following displays the current event control configuration:

```
A:ALA-12>config>log# info
-----
...
event-control "ospf" 2014 generate critical
...
-----
A:ALA-12>config>log#
```

The following displays an example of an event control modification:

```
Example:config# log
config>log# event-control ospf 2014 suppress
```

The following displays the log filter configuration:

```
A:ALA-12>config>log# info
-----
...
event-control "ospf" 2014 suppress
...
-----
A:ALA-12>config>log#
```

5.5.13 Returning to the Default Event Control Configuration

The **no** form of the **event-control** command returns modified values back to the default values.

Use the following CLI syntax to modify event control parameters:

```
config>log
no event-control application [event-name
|event-number]
```

The following displays an example of the command usage to return to the default values:

```
Example:config# log
config>log# no event-control "bgp" 2001
config>log# no event-control "bgp" 2002
config>log# no event-control "bgp" 2014
```

```

A:ALA-12>config>log# info detail
-----
#-----
echo "Log Configuration"
#-----
    event-control "bgp" 2001 generate minor
    event-control "bgp" 2002 generate warning
    event-control "bgp" 2003 generate warning
    event-control "bgp" 2004 generate critical
    event-control "bgp" 2005 generate warning
    event-control "bgp" 2006 generate warning
    event-control "bgp" 2007 generate warning
    event-control "bgp" 2008 generate warning
    event-control "bgp" 2009 generate warning
    event-control "bgp" 2010 generate warning
    event-control "bgp" 2011 generate warning
    event-control "bgp" 2012 generate warning
    event-control "bgp" 2013 generate warning
    event-control "bgp" 2014 generate warning
    event-control "bgp" 2015 generate critical
    event-control "bgp" 2016 generate warning
...
-----
A:ALA-12>config>log#

```

5.6 Accounting Logs

Before an accounting policy can be created a target log file must be created to collect the accounting records. The files are stored in system memory on compact flash (*cf1:* or *cf2:*) in a compressed (tar) XML format and can be retrieved using FTP or SCP.

A file ID can only be assigned to either one event log ID or one accounting log.

5.6.1 Accounting Records

An accounting policy must define a record name and collection interval. Only one record name can be configured per accounting policy. Also, a record name can only be used in one accounting policy.

The record name, sub-record types, and default collection period for service and network accounting policies are shown in [Table 63](#). [Table 65](#) (fields per policer stat-mode are given in the **stat-mode** command descriptions in the *Quality of Service Guide*), [Table 66](#), and [Table 67](#) provide field descriptions.

Table 63 Accounting Record Name and Collection Periods

Record Name	Sub-Record Types	Accounting Object	Platform	Default Collection Period (minutes)
service-ingress-octets	sio	SAP	All	5
service-egress-octets	seo	SAP	All	5
service-ingress-packets	sip	SAP	All	5
service-egress-packets	sep	SAP	All	5
network-ingress-octets	nio	Network port	All	15
network-egress-octets	neo	Network port	All	15
network-egress-packets	nep	Network port	All	15
network-ingress-packets	nio	Network port	All	15
compact-service-ingress-octets	ctSio	SAP	All	5
combined-service-ingress	cmSipo	SAP	All	5
combined-network-ing-egr-octets	cmNio & cmNeo	Network port	All	15
combined-service-ing-egr-octets	cmSio & cmSeo	SAP	All	5
complete-network-ingr-egr	cpNipo & cpNepo	Network port	All	15
complete-service-ingress-egress	cpSipo & cpSepo	SAP	All	5
combined-sdp-ingress-egress	cmSdpipo and cmSdpepo	SDP and SDP binding	All	5
complete-sdp-ingress-egress	cmSdpipo, cmSdpepo, cpSdpipo and cpSdpepo	SDP and SDP binding	All	5
complete-subscriber-ingress-egress	cpSBipo & cpSBepo	Subscriber profile	7750 SR	5
aa-protocol	aaProt	AA ISA Group	7750 SR	15
aa-application	aaApp	AA ISA Group	7750 SR	15
aa-app-group	aaAppGrp	AA ISA Group	7750 SR	15

Table 63 Accounting Record Name and Collection Periods (Continued)

Record Name	Sub-Record Types	Accounting Object	Platform	Default Collection Period (minutes)
aa-subscriber-protocol	aaSubProt	Special study AA subscriber	7750 SR	15
aa-subscriber-application	aaSubApp	Special study AA subscriber	7750 SR	15
custom-record-aa-sub	aaSubCustom	AA subscriber	All	15
combined-mpls-lsp-egress	mplsLspEgr	LSP	All	5
combined-mpls-lsp-ingress	mplsLspIn	LSP	All	5
saa	saa png trc hop	SAA or SAA test	All	5
complete-ethernet-port	enet	Ethernet port	All	15

When creating accounting policies, one service accounting policy and one network accounting policy can be defined as default. If statistics collection is enabled on a SAP or network port and no accounting policy is applied, then the respective default policy is used. If no default policy is defined, then no statistics are collected unless a specifically defined accounting policy is applied.

Each accounting record name is composed of one or more sub-records which is in turn composed of multiple fields.

Refer to the Application Assurance Statistics Fields Generated per Record table Integrated Services Adapter Guide for fields names for Application Assurance records.

The availability of the records listed in [Table 64](#) depends on your specific platform functionality and user configuration.

Table 64 Accounting Record Name Details

Record Name	Sub-Record	Field	Field Description
Service-ingress-octets (sio) ¹	sio	svc	Svcld
		sap	Sapld
		qid	QueueId
		hoo	OfferedHiPrioOctets
		hod	DroppedHiPrioOctets
		loo	LowOctetsOffered
		lod	LowOctetsDropped
		uco	UncoloredOctetsOffered
		iof	InProfileOctetsForwarded
		oof	OutOfProfileOctetsForwarded
Service-egress-octets (seo) ¹	seo	svc	Svcld
		sap	Sapld
		qid	QueueId
		iof	InProfileOctetsForwarded
		iod	InProfileOctetsDropped
		oof	OutOfProfileOctetsForwarded
		ood	OutOfProfileOctetsDropped
Service-ingress-packets (sip) ^{1 2}	sip	svc	Svcld
		sap	Sapld
		qid	QueueId
		hpo	HighPktsOffered
		hpd	HighPktsDropped
		lpo	LowPktsOffered
		lpd	LowPktsDropped
		ucp	UncoloredPacketsOffered
		ipf	InProfilePktsForwarded
		opf	OutOfProfilePktsForwarded

Table 64 Accounting Record Name Details (Continued)

Record Name	Sub-Record	Field	Field Description
Service-egress-packets (sep) ^{1 2}	sep	svc	SvcId
		sap	SapId
		qid	QueueId
		ipf	InProfilePktsForwarded
		ipd	InProfilePktsDropped
		opf	OutOfProfilePktsForwarded
		opd	OutOfProfilePktsDropped
Network-ingress-octets (nio)	nio	port	PortId
		qid	QueueId
		iof	InProfileOctetsForwarded
		iod	InProfileOctetsDropped
		oof	OutOfProfileOctetsForwarded
		ood	OutOfProfileOctetsDropped
Network-egress-octets (neo)	neo	port	PortId
		qid	QueueId
		iof	InProfileOctetsForwarded
		iod	InProfileOctetsDropped
		oof	OutOfProfileOctetsForwarded
		ood	OutOfProfileOctetsDropped
Network-ingress-packets (nip)	nip	port	PortId
		qid	QueueId
		ipf	InProfilePktsForwarded
		ipd	InProfilePktsDropped
		opf	OutOfProfilePktsForwarded
		opd	OutOfProfilePktsDropped

Table 64 Accounting Record Name Details (Continued)

Record Name	Sub-Record	Field	Field Description
Network Egress Packets (nep)	nep	port	PortId
		qid	QueueId
		ipf	InProfilePktsForwarded
		ipd	InProfilePktsDropped
		opf	OutOfProfilePktsForwarded
		opd	OutOfProfilePktsDropped
Compact-service-ingress-octets (ctSio)	ctSio	svc	SvcId
		sap	SapId
		qid	QueueId
		hoo	OfferedHiPrioOctets
		hod	DroppedHiPrioOctets
		loo	LowOctetsOffered
		lod	LowOctetsDropped
		uco	UncoloredOctetsOffered

Table 64 Accounting Record Name Details (Continued)

Record Name	Sub-Record	Field	Field Description
Combined-service-ingress (cmSipo)	cmSipo	svc	SvcId
		sap	SapId
		qid	QueueId
		hpo	HighPktsOffered
		hpd	HighPktsDropped
		lpo	LowPktsOffered
		lpd	LowPktsDropped
		ucp	UncoloredPacketsOffered
		hoo	OfferedHiPrioOctets
		hod	DroppedHiPrioOctets
		loo	LowOctetsOffered
		lod	LowOctetsDropped
		uco	UncoloredOctetsOffered
		ipf	InProfilePktsForwarded
		opf	OutOfProfilePktsForwarded
		iof	InProfileOctetsForwarded
oof	OutOfProfileOctetsForwarded		

Table 64 Accounting Record Name Details (Continued)

Record Name	Sub-Record	Field	Field Description
Combined-network-ing-egr-octets (cmNio & cmNeo)	cmNio	port	PortId
		qid	QueueId
		iof	InProfileOctetsForwarded
		iod	InProfileOctetsDropped
		oof	OutOfProfileOctetsForwarded
		ood	OutOfProfileOctetsDropped
	cmNeo	port	PortId
		qid	QueueId
		iof	InProfileOctetsForwarded
		iod	InProfileOctetsDropped
		oof	OutOfProfileOctetsForwarded
		ood	OutOfProfileOctetsDropped

Table 64 Accounting Record Name Details (Continued)

Record Name	Sub-Record	Field	Field Description
Combined-service-ingr-egr-octets (cmSio & CmSeo)	cmSio	svc	Svcld
		sap	Sapld
		qid	QueueId
		hoo	OfferedHiPrioOctets
		hod	DroppedHiPrioOctets
		loo	LowOctetsOffered
		lod	LowOctetsDropped
		uco	UncoloredOctetsOffered
		iof	InProfileOctetsForwarded
		oof	OutOfProfileOctetsForwarded
	cmSeo	svc	Svcld
		sap	Sapld
		qid	QueueId
		iof	InProfileOctetsForwarded
		iod	InProfileOctetsDropped
		oof	OutOfProfileOctetsForwarded
		ood	OutOfProfileOctetsDropped

Table 64 Accounting Record Name Details (Continued)

Record Name	Sub-Record	Field	Field Description
Complete-network-ingr-egr (cpNipo & cpNepo)	cpNipo	port	PortId
		qid	QueueId
		ipf	InProfilePktsForwarded
		ipd	InProfilePktsDropped
		opf	OutOfProfilePktsForwarded
		opd	OutOfProfilePktsDropped
		iof	InProfileOctetsForwarded
		iod	InProfileOctetsDropped
		oof	OutOfProfileOctetsForwarded
		ood	OutOfProfileOctetsDropped
	cpNepo	port	PortId
		qid	QueueId
		ipf	InProfilePktsForwarded
		ipd	InProfilePktsDropped
		opf	OutOfProfilePktsForwarded
		opd	OutOfProfilePktsDropped
		iof	InProfileOctetsForwarded
		iod	InProfileOctetsDropped
		oof	OutOfProfileOctetsForwarded
		ood	OutOfProfileOctetsDropped

Table 64 Accounting Record Name Details (Continued)

Record Name	Sub-Record	Field	Field Description
Complete-service-ingress-egress (cpSipo & cpSepa)	cpSipo	svc	SvcId
		sap	SapId
		pid	PolicerId
		hpo	HighPktsOffered
		hpd	HighPktsDropped
		lpo	LowPktsOffered
		lpd	LowPktsDropped
		ucp	UncoloredPacketsOffered
		hoo	OfferedHiPrioOctets
		hod	DroppedHiPrioOctets
		loo	LowOctetsOffered
		lod	LowOctetsDropped
		uco	UncoloredOctetsOffered
		apo	AllPacketsOffered
		aoo	AllOctetsOffered
		apd	AllPacketsDropped
		aod	AllOctetsDropped
		apf	AllPacketsForwarded
		aof	AllOctetsForwarded
		ipd	InProfilePktsDropped
		iod	InProfileOctetsDropped
opd	OutOfProfilePktsDropped		
ood	OutOfProfileOctetsDropped		
hpf	HighPriorityPacketsForwarded		
hof	HighPriorityOctetsForwarded		

Table 64 Accounting Record Name Details (Continued)

Record Name	Sub-Record	Field	Field Description		
Complete-service-ingress-egress (cpSipo & cpSepo) (Continued)	cpSipo (Continued)	lpf	LowPriorityPacketsForwarded		
		lof	LowPriorityOctetsForwarded		
		ipf	InProfilePktsForwarded		
		opf	OutOfProfilePktsForwarded		
		iof	InProfileOctetsForwarded		
		oof	OutOfProfileOctetsForwarded		
	cpSepo	svc	SvcId		
		sap	SapId		
		qid	QueueId		
		ipf	InProfilePktsForwarded		
		ipd	InProfilePktsDropped		
		opf	OutOfProfilePktsForwarded		
		opd	OutOfProfilePktsDropped		
		iof	InProfileOctetsForwarded		
		iod	InProfileOctetsDropped		
		oof	OutOfProfileOctetsForwarded		
		ood	OutOfProfileOctetsDropped		
		Complete-sdp-ingress-egress (cpSdpipe & cpSdpepo)	cpSdpipe	sdp	SdpID
				tpf	TotalPacketsForwarded
tpd	TotalPacketsDropped				
tof	TotalOctetsForwarded				
tod	TotalOctetsDropped				
cpSdpepo	sdp		SdpID		
	tpd		TotalPacketsDropped		
	tod		TotalOctetsDropped		
	tod		TotalOctetsDropped		

Table 64 Accounting Record Name Details (Continued)

Record Name	Sub-Record	Field	Field Description
Combined-sdp-ingress-egress (cmSdpipo & cmSdpepo)	cmSdpipo	svc	SvcID
		sdp	SdpID
		tpf	TotalPacketsForwarded
		tpd	TotalPacketsDropped
		tof	TotalOctetsForwarded
		tod	TotalOctetsDropped
	cmSdpepo	svc	SvcID
		sdp	SdpID
		tpf	TotalPacketsForwarded
		tof	TotalOctetsForwarded

Table 64 Accounting Record Name Details (Continued)

Record Name	Sub-Record	Field	Field Description
Complete-sdp-ingress-egress (cmSdpipo & cmsdpepo) (cpSdpip & cpSdpepo)	cmSdpipo	svc	SvcID
		sdp	SdpID
		tpf	TotalPacketsForwarded
		tpd	TotalPacketsDropped
		tof	TotalOctetsForwarded
		tod	TotalOctetsDropped
	cmSdpepo	svc	SvcID
		sdp	SdpID
		tpf	TotalPacketsForwarded
		tof	TotalOctetsForwarded
	cpSdpipo	sdp	SdpID
		tpf	TotalPacketsForwarded
		tpd	TotalPacketsDropped
		tof	TotalOctetsForwarded
		tod	TotalOctetsDropped
	cpSdpepo	sdp	SdpID
		tpf	TotalPacketsForwarded
		tof	TotalOctetsForwarded
Complete-subscriber-ingress-egress (cpSBipo & cpSBepo) (cpSBipoc & cpSBepoc) ³	SubscriberInformation	subId	SubscriberId
		subProfile	SubscriberProfile
	Sla-Information ⁴	svc	SvcId
		sap	SapId
		slaProfile	SlaProfile

Table 64 Accounting Record Name Details (Continued)

Record Name	Sub-Record	Field	Field Description
Complete-subscriber-ingress-egress (cpSBipo & cpSBepo) (cpSBipooc & cpSBepooc) ³ (Continued)	cpSBipo	qid	QueueId
		hpo	HighPktsOffered ⁴
		hpd	HighPktsDropped
		lpo	LowPktsOffered ⁴
		lpd	LowPktsDropped
		ucp	UncolouredPacketsOffered
		hoo	OfferedHiPrioOctets ⁴
		hod	DroppedHiPrioOctets
		loo	LowOctetsOffered ⁴
		lod	LowOctetsDropped
		apo	AllPktsOffered ⁴
		aoo	AllOctetsOffered ⁴
		uco	UncolouredOctetsOffered
		ipf	InProfilePktsForwarded
		opf	OutOfProfilePktsForwarded
		iof	InProfileOctetsForwarded
		oof	OutOfProfileOctetsForwarded
		v4pf	IPv4PktsForwarded
		v6pf	IPv6PktsForwarded
		v4pd	IPv4PktsDropped
		v6pd	IPv6PktsDropped
		v4of	IPv4OctetsForwarded
		v6of	IPv6OctetsForwarded
v4od	IPv4OctetsDropped		
v6od	IPv6OctetsDropped		

Table 64 Accounting Record Name Details (Continued)

Record Name	Sub-Record	Field	Field Description
Complete-subscriber-ingress-egress (cpSBipo & cpSBepo) (cpSBipooc & cpSBepooc) ³ (Continued)	cpSBepo	qid	QueueId
		ipf	InProfilePktsForwarded
		ipd	InProfilePktsDropped
		opf	OutOfProfilePktsForwarded
		opd	OutOfProfilePktsDropped
		iof	InProfileOctetsForwarded
		iod	InProfileOctetsDropped
		oof	OutOfProfileOctetsForwarded
		ood	OutOfProfileOctetsDropped
		v4pf	IPv4PktsForwarded
		v6pf	IPv6PktsForwarded
		v4pd	IPv4PktsDropped
		v6pd	IPv6PktsDropped
		v4of	IPv4OctetsForwarded
		v6of	IPv6OctetsForwarded
		v4od	IPv4OctetsDropped
v6od	IPv6OctetsDropped		

Table 64 Accounting Record Name Details (Continued)

Record Name	Sub-Record	Field	Field Description
Complete-subscriber-ingress-egress (cpSBipo & cpSBepo) (cpSBipoooc & cpSBepoooc) ³ (Continued)	cpSBipoooc ³	cid	OverrideCounterId
		apo	AllPktsOffered
		hpd	HighPktsDropped
		lpd	LowPktsDropped
		aoo	AllOctetsOffered
		hod	DroppedHiPrioOctets
		lod	LowOctetsDropped
		ipf	InProfilePktsForwarded
		opf	OutOfProfilePktsForwarded
		iof	InProfileOctetsForwarded
		oof	OutOfProfileOctetsForwarded
		ucp	UncolouredPacketsOffered
		uco	UncolouredOctetsOffered
	cpSBepoooc ³	cid	OverrideCounterId
		ipf	InProfilePktsForwarded
		ipd	InProfilePktsDropped
		ofp	OutOfProfilePktsForwarded
		opd	OutOfProfilePktsDropped
		ipd	InProfileOctetsForwarded
		iod	InProfileOctetsDropped
oof		OutOfProfileOctetsForwarded	
ood	OutOfProfileOctetsDropped		

Table 64 Accounting Record Name Details (Continued)

Record Name	Sub-Record	Field	Field Description
saa	saa	tmd	TestMode
		own	OwnerName
		tst	TestName
		png	PingRun subrecord
		rid	RunIndex
		trr	TestRunResult
		mnr	MinRtt
		mrx	MaxRtt
		avr	AverageRtt
		rss	RttSumOfSquares
		pbr	ProbeResponses
		spb	SentProbes
		mnt	MinOutTt
		mxt	MaxOutTt
		avt	AverageOutTt
		tss	OutTtSumOfSquares
		mni	MinInTt
		mxi	MaxInTt
		avi	AverageInTt
		iss	InTtSumOfSqrS
ojt	OutJitter		
ijt	InJitter		
rjt	RtJitter		
prt	ProbeTimeouts		
prf	ProbeFailures		

Table 64 Accounting Record Name Details (Continued)

Record Name	Sub-Record	Field	Field Description
saa (Continued)	trc	rid	RunIndex
		trr	TestRunResult
		lgp	LastGoodProbe
	hop	hop	TraceHop
		hid	HopIndex
		mnr	MinRtt
		mxx	MaxRtt
		avr	AverageRtt
		rss	RttSumOfSquares
		pbr	ProbeResponses
		spb	SentProbes
		mnt	MinOutTt
		mxt	MaxOutTt
		avt	AverageOutTt
		tss	OutTtSumOfSquares
		mni	MinInTt
		mxi	MaxInTt
		avi	AverageInTt
		iss	InTtSumOfSqrS
		ojt	OutJitter
		ijt	InJitter
		rjt	RtJitter
		prt	ProbeTimeouts
prf	ProbeFailures		
tat	TraceAddressType		
tav	TraceAddressValue		

Table 64 Accounting Record Name Details (Continued)

Record Name	Sub-Record	Field	Field Description
Complete-ethernet-port (enet)	enet	port	PortId
		to	EtherStatsOctets
		tp	EtherStatsPkts
		de	EtherStatsDropEvents
		tbcP	EtherStatsBroadcastPkts
		mcp	EtherStatsMulticastPkts
		cae	EtherStatsCRCAlignErrors
		up	EtherStatsUndersizePkts
		op	EtherStatsOversizePkts
		fgm	EtherStatsFragments
		jab	EtherStatsJabbers
		col	EtherStatsCollisions
		p64o	EtherStatsPkts64Octets
		p127o	EtherStatsPkts65to127Octets
		p255o	EtherStatsPkts128to255Octets
		p511o	EtherStatsPkts256to511Octets
		p1023o	EtherStatsPkts512to1023Octets
		p1518o	EtherStatsPkts1024to1518Octets
		po1518o	EtherStatsPktsOver1518Octets
		ae	Dot3StatsAlignmentErrors
		fe	Dot3StatsFCSErrors
		scf	Dot3StatsSingleCollisionFrames
		mcf	Dot3StatsMultipleCollisionFrames
sqe	Dot3StatsSQETestErrors		
dt	Dot3StatsDeferredTransmissions		

Table 64 Accounting Record Name Details (Continued)

Record Name	Sub-Record	Field	Field Description
Complete-ethernet-port (enet) (Continued)	enet (Continued)	lcc	Dot3StatsLateCollisions
		exc	Dot3StatsExcessiveCollisions
		imt	Dot3StatsInternalMacTransmitErrors
		cse	Dot3StatsCarrierSenseErrors
		ftl	Dot3StatsFrameTooLongs
		imre	Dot3StatsInternalMacReceiveErrors
		se	Dot3StatsSymbolErrors
		ipf	Dot3InPauseFrames
		opf	Dot3OutPauseFrames

Notes:

1. The number of octets in an ATM sap excludes the Header Error Control (HEC) byte, thus meaning each packet/cell has only 52 bytes instead of the usual 53.
2. For a SAP in AAL5 SDU mode, packet counters refer to the number of SDU. For a SAP in N-to-1 cell mode, packet counters refer to the number of cells.
3. If override counters on the HSMDA are configured (see the Quality of Service Guide).
4. Not used to identify stats from HSMDA due to MDA architecture. If the statistics are from HSMDA: apo, aoo else lpo/hpo, loo/hoo.

[Table 65](#), [Table 66](#), and [Table 67](#) provide field descriptions.

Table 65 Policer Stats Field Descriptions

Field	Field Description
pid	PolicerId
statmode	PolicerStatMode
aod	AllOctetsDropped
aof	AllOctetsForwarded
aoo	AllOctetsOffered
apd	AllPacketsDropped
apf	AllPacketsForwarded

Table 65 Policer Stats Field Descriptions (Continued)

Field	Field Description
apo	AllPacketsOffered
hod	HighPriorityOctetsDropped
hof	HighPriorityOctetsForwarded
hoo	HighPriorityOctetsOffered
hpd	HighPriorityPacketsDropped
hpf	HighPriorityPacketsForwarded
hpo	HighPriorityPacketsOffered
iod	InProfileOctetsDropped
iof	InProfileOctetsForwarded
ioo	InProfileOctetsOffered
ipd	InProfilePacketsDropped
ipf	InProfilePacketsForwarded
ipo	InProfilePacketsOffered
lod	LowPriorityOctetsDropped
lof	LowPriorityOctetsForwarded
loo	LowPriorityOctetsOffered
lpd	LowPriorityPacketsDropped
lpf	LowPriorityPacketsForwarded
lpo	LowPriorityPacketsOffered
opd	OutOfProfilePacketsDropped
opf	OutOfProfilePacketsForwarded
opo	OutOfProfilePacketsOffered
ood	OutOfProfileOctetsDropped
oof	OutOfProfileOctetsForwarded
ooo	OutOfProfileOctetsOffered
xpd	ExceedProfilePktsDropped
xpf	ExceedProfilePktsForwarded

Table 65 Policer Stats Field Descriptions (Continued)

Field	Field Description
xpo	ExceedProfilePktsOffered
xod	ExceedProfileOctetsDropped
xof	ExceedProfileOctetsForwarded
xoo	ExceedProfileOctetsOffered
ppd	InplusProfilePacketsDropped
ppf	InplusProfilePacketsForwarded
ppo	InplusProfilePacketsOffered
pod	InplusProfileOctetsDropped
pof	InplusProfileOctetsForwarded
poo	InplusProfileOctetsOffered
uco	UncoloredOctetsOffered
ucp	UncoloredPacketsOffered
v4po	IPv4PktsOffered *
v4oo	IPv4OctetsOffered *
v6po	IPv6PktsOffered *
v6oo	IPv6OctetsOffered *
v4pf	IPv4PktsForwarded *
v6pf	IPv6PktsForwarded *
v4pd	IPv4PktsDropped *
v6pd	IPv6PktsDropped *
v4of	IPv4OctetsForwarded *
v6of	IPv6OctetsForwarded *
v4od	IPv4OctetsDropped *
v6od	IPv6OctetsDropped *

* Enhanced Subscriber Management (ESM) only.

Table 66 Queue Group Record Types

Record Name	Description
qgone	PortQueueGroupOctetsNetworkEgress
qgosi	PortQueueGroupOctetsServiceIngress
qgose	PortQueueGroupOctetsServiceEgress
qgpne	PortQueueGroupPacketsNetworkEgress
qgpsi	PortQueueGroupPacketsServiceIngress
qgpse	PortQueueGroupPacketsServiceEgress
fpqgosi	ForwardingPlaneQueueGroupOctetsServiceIngress
fpqgoni	ForwardingPlaneQueueGroupOctetsNetworkIngress
fpqgpsi	ForwardingPlaneQueueGroupPacketsServiceIngress
fpqgpni	ForwardingPlaneQueueGroupPacketsNetworkIngress

Table 67 Queue Group Record Type Fields

Field	Field Description
data port	Port (used for port based Queue Groups)
member-port	LAGMemberPort (used for port based Queue Groups)
data slot	Slot (used for Forwarding Plane based Queue Groups)
forwarding-plane	ForwardingPlane (used for Forwarding Plane based Queue Groups)
queue-group	QueueGroupName
instance	QueueGroupInstance
qid	QueueId
pid	PolicerId
statmode	PolicerStatMode
aod...ucp	same as above

5.6.2 Accounting Files

When a policy has been created and applied to a service or network port, the accounting file is stored on the compact flash in a compressed XML file format. The router creates two directories on the compact flash to store the files. The following output displays a directory named **act-collect** that holds accounting files that are open and actively collecting statistics. The directory named **act** stores the files that have been closed and are awaiting retrieval.

```
ALA-1>file cf1:\# dir act*
12/19/2006 06:08a      <DIR>          act-collect
12/19/2006 06:08a      <DIR>          act

ALA-1>file cf1:\act-collect\ # dir
Directory of cf1:\act-collect#

12/23/2006 01:46a      <DIR>          .
12/23/2006 12:47a      <DIR>          ..
12/23/2006 01:46a                112 act1111-20031223-014658.xml.gz
12/23/2006 01:38a                197 act1212-20031223-013800.xml.gz
```

Accounting files always have the prefix **act** followed by the accounting policy ID, log ID and timestamp. The accounting log file naming and log file destination properties like rollover and retention are discussed in more detail in [Log Files](#).

5.6.3 Design Considerations

The router has ample resources to support large scale accounting policy deployments. When preparing for an accounting policy deployment, verify that data collection, file rollover, and file retention intervals are properly tuned for the amount of statistics to be collected.

If the accounting policy collection interval is too brief there may be insufficient time to store the data from all the services within the specified interval. If that is the case, some records may be lost or incomplete. Interval time, record types, and number of services using an accounting policy are all factors that should be considered when implementing accounting policies.

The rollover and retention intervals on the log files and the frequency of file retrieval must also be considered when designing accounting policy deployments. The amount of data stored depends on the type of record collected, the number of services that are collecting statistics, and the collection interval that is used. For example, with a 1GB CF and using the default collection interval, the system is expected to hold 48 hours worth of billing information.

5.6.4 Reporting and Time-Based Accounting

SR OS on the 7750 SR platform has support for volume accounting and time-based accounting concepts, and provides an extra level of intelligence at the network element level in order to provide service models such as “prepaid access” in a scalable manner. This means that the network element gathers and stores per-subscriber accounting information and compares it with “pre-defined” quotas. Once a quota is exceeded, the pre-defined action (such as re-direction to a web portal or disconnect) is applied.

5.6.5 Overhead Reduction in Accounting: Custom Record

Custom records can be used to decrease accounting messaging overhead as follows:

- [User Configurable Records](#)
- [Changed Statistics Only](#)
- [Configurable Accounting Records](#)
- [Significant Change Only Reporting](#)

5.6.5.1 User Configurable Records

Users can define a collection of fields that make up a record. These records can be assigned to an accounting policy. These are user-defined records rather than being limited to pre-defined record types. The operator can select what queues and the counters within these queues that need to be collected. Refer to the predefined records containing a given field for XML field name of a custom record field.

5.6.5.2 Changed Statistics Only

A record is only generated if a significant change has occurred to the fields being written in a given the record. This capability applies to both ingress and egress records regardless on the method of delivery (such as RADIUS and XML). The capability also applies to Application Assurance records; however without an ability to specify different significant change values and per-field scope (for example, all fields of a custom record are collected if any activity was reported against any of the statistics that are part of the custom record).

5.6.5.3 Configurable Accounting Records

- [XML Accounting Files for Service and ESM-Based Accounting](#)
- [RADIUS Accounting in Networks Using ESM](#)

5.6.5.3.1 XML Accounting Files for Service and ESM-Based Accounting

The **custom-record** command in the **config>log>accounting-policy** context provide the flexibility to reduce the volume of data generated, network operators can define the record that needs to be collected. This can eliminate queues or selected counters within these queues that are not relevant for billing.

ESM-based accounting applies to the 7750 SR only.

Record headers including information such as service-ID, SAP-ID, etc., will always be generated.

5.6.5.3.2 RADIUS Accounting in Networks Using ESM

The **custom-record** command in the **config>subscr-mgmt>radius-accounting-policy** context provide the flexibility to include individual counters in RADIUS accounting messages. See the CLI tree for commands and syntax. This functionality applies to the 7750 SR only.

5.6.5.4 Significant Change Only Reporting

Another way to decrease accounting messaging related to overhead is to include only “active” objects in a periodical reporting. An “active object” in this context is an object which has seen a “significant” change in corresponding counters. A significant change is defined in terms of a cumulative value (the sum of all reference counters).

This concept is applicable to all methods used for gathering accounting information, such as an XML file and RADIUS, as well as to all applications using accounting, such as service-acct, ESM-acct, and Application Assurance.

Accounting records are reported at the periodical intervals. This periodic reporting is extended with an internal filter which omits periodical updates for objects whose counter change experienced lower changes than a defined (configurable) threshold.

Specific to RADIUS accounting the **significant-change** command does not affect ACCT-STOP messages. ACCT-STOP messages will be always sent, regardless the amount of change of the corresponding host.

For Application Assurance records, a significant change of 1 in any field of a customized record (send a record if any field changed) is supported. When configured, if any statistic field records activity, an accounting record containing all fields will be collected.

5.6.6 Immediate Completion of Records

5.6.6.1 Record Completion for XML Accounting

For ESM RADIUS accounting, an accounting stop message is sent when:

- A subscriber/subscriber-host is deleted.
- An SLA profile instance (non-HSMDA) or subscriber instance (HSMDA) is changed.

A similar concept is also used for XML accounting. In case the accounted object is deleted or changed, the latest information will be written in the XML file with a “final” tag indication in the record header. This functionality applies to the 7750 SR only.

5.6.7 AA Accounting per Forwarding Class

This feature allows the operator to report on protocol/application/app-group volume usage per forwarding class by adding a bitmap information representing the observed FC in the XML accounting files. In case the accounted object is deleted or changed, the latest information will be written in the XML file with a “final” tag indication in the record header.

5.7 Configuration Notes

This section describes logging configuration caveats.

- A file or filter cannot be deleted if it has been applied to a log.
- File IDs, syslog IDs, or SNMP trap groups must be configured before they can be applied to a log ID.
- A file ID can only be assigned to *either* one log ID *or* one accounting policy.
- Accounting policies must be configured in the **config>log** context before they can be applied to a service SAP or service interface, or applied to a network port.
- The **snmp-trap-id** must be the same as the **log-id**.

5.8 Configuring Logging with CLI

5.9 Log Configuration Overview

Configure logging parameters to save information in a log file or direct the messages to other devices. Logging does the following:

- Provides you with logging information for monitoring and troubleshooting.
- Allows you to select the types of logging information to be recorded.
- Allows you to assign a severity to the log messages.
- Allows you to select the source and target of logging information.

5.9.1 Log Types

Logs can be configured in the following contexts:

- Log file — Log files can contain log event message streams or accounting/billing information. Log file IDs are used to direct events, alarms/traps and debug information to their respective targets.
- SNMP trap groups — SNMP trap groups contain an IP address and community names which identify targets to send traps following specified events.
- Syslog — Information can be sent to a syslog host that is capable of receiving selected syslog messages from a network element.
- Event control — Configures a particular event or all events associated with an application to be generated or suppressed.
- Event filters — An event filter defines whether to forward or drop an event or trap based on match criteria.
- Accounting policies — An accounting policy defines the accounting records that will be created. Accounting policies can be applied to one or more service access points (SAPs).
- Event logs — An event log defines the types of events to be delivered to its associated destination.
- Event throttling rate — Defines the rate of throttling events.

5.10 Basic Event Log Configuration

The most basic log configuration must have the following:

- Log ID or accounting policy ID
- A log source
- A log destination

The following displays a log configuration example for the 7750 SR.

```
A:ALA-12>config>log# info
#-----
echo "Log Configuration "
#-----
    event-control "bgp" 2001 generate critical
    file-id 1
        description "This is a test file-id."
        location cf1:
    exit
    file-id 2
        description "This is a test log."
        location cf1:
    exit
    snmp-trap-group 7
        trap-target 11.22.33.44 "snmpv2c" notify-community "public"
    exit
    log-id 2
        from main
        to file 2
    exit
-----
A:ALA-12>config>log#
```

5.11 Common Configuration Tasks

The following sections describe basic system tasks that must be performed.

- [Configuring a File ID](#)
- [Configuring an Event Log](#)
- [Configuring an Accounting Policy](#)
- [Configuring Event Control](#)
- [Configuring a Log Filter](#)
- [Configuring an SNMP Trap Group](#)
- [Configuring a Syslog Target](#)

5.11.1 Configuring an Event Log

A event log file contains information used to direct events, alarms, traps, and debug information to their respective destinations. One or more event sources can be specified. File IDs, SNMP trap groups, or syslog IDs must be configured before they can be applied to an event log ID.

Use the following CLI syntax to configure a log file:

```
config>log
      log-id log-id
      description description-string
      filter filter-id
      from {[main] [security] [change] [debug-trace]}
      to console
      to file file-id
      to memory [size]
      to session
      to snmp [size]
      to syslog syslog-id}
      time-format {local | utc}
      no shutdown
```

The following displays a log file configuration example:

```
ALA-12>config>log>log-id# info
-----
...
log-id 2
      description "This is a test log file."
      filter 1
      from main security
      to file 1
exit
...
-----
ALA-12>config>log>log-id#
```

5.11.2 Configuring a File ID

To create a log file a file ID is defined, specifies the target CF drive, and the rollover and retention interval period for the file. The rollover interval is defined in minutes and determines how long a file will be used before it is closed and a new log file is created. The retention interval determines how long the file will be stored on the CF before it is deleted.

When creating new log files in a compact flash disk card, the minimum amount of free space is the MINIMUM of 10% of Compact Flash disk capacity OR 5Mb (5,242,880 = 5 * 1024 * 1024).

The following displays a log file configuration example:

```
A:ALA-12>config>log# info
-----
      file-id 1
      description "This is a log file."
      location cf1:
      rollover 600 retention 24
      exit
-----
A:ALA-12>config>log#
```

5.11.3 Configuring an Accounting Policy

Before an accounting policy can be created a target log file must be created to collect the accounting records. The files are stored in system memory of compact flash (cf1: or cf2:) in a compressed (tar) XML format and can be retrieved using FTP or SCP. See [Configuring an Event Log](#) and [Configuring a File ID](#).

Accounting policies must be configured in the **config>log** context before they can be applied to a service SAP or service interface, or applied to a network port.

The default accounting policy statement cannot be applied to LDP nor RSVP statistics collection records.

An accounting policy must define a record type and collection interval. Only one record type can be configured per accounting policy.

When creating accounting policies, one service accounting policy and one network accounting policy can be defined as default. If statistics collection is enabled on a SAP or network port and no accounting policy is applied, then the respective default policy is used. If no default policy is defined, then no statistics are collected unless a specifically defined accounting policy is applied.

By default, the subscriber host volume accounting data are based on the 14-byte Ethernet DLC header, 4-byte or 8-byte VLAN Tag (optional), 20-byte IP header, IP payload, and the 4-byte CRC (everything except the preamble and inter-frame gap). See [Figure 20](#). This default can be altered by the `packet-byte-offset` configuration option.

Figure 20 Subscriber Host Volume Accounting Data

Destination MAC	Source MAC	802.1Q tag (optional)	802.1Q tag (optional)	Ethertype (Ethernet II) or length (IEEE 802.3)	Payload	CRC/FCS
6 octets	6 octets	(4 octets)	(4 octets)	2 octets	46-1500 octets	4 octets

0971

The following displays an accounting policy configuration example:

```
A:ALA-12>config>log# info
-----
accounting-policy 4
description "This is the default accounting policy."
record complete-service-ingress-egress
default
to file 1
exit
accounting-policy 5
description "This is a test accounting policy."
record service-ingress-packets
to file 3
exit
```

5.11.4 Configuring Event Control

The following displays an example of an event control configuration:

```
A:ALA-12>config>log# info
#-----
echo "Log Configuration"
#-----
      throttle-rate 500 interval 10
      event-control "oam" 2001 generate throttle
      event-control "ospf" 2001 suppress
      event-control "ospf" 2003 generate cleared
      event-control "ospf" 2014 generate critical
..
-----
A:ALA-12>config>log>filter#
```

5.11.5 Configuring a Log Filter

The following displays a log filter configuration example:

```
A:ALA-12>config>log# info
#-----
echo "Log Configuration "
#-----
      file-id 1
```

```

        description "This is our log file."
        location cf1:
        rollover 600 retention 24
    exit
    filter 1
        default-action drop
        description "This is a sample filter."
        entry 1
            action forward
            match
                application eq "mirror"
                severity eq critical
            exit
        exit
    exit
...
log-id 2
    shutdown
    description "This is a test log file."
    filter 1
    from main security
    to file 1
    exit
...
-----
A:ALA-12>config>log#

```

5.11.6 Configuring an SNMP Trap Group

The associated *log-id* does not have to be configured before a **snmp-trap-group** can be created, however, the **snmp-trap-group** must exist before the *log-id* can be configured to use it.

The following displays a basic SNMP trap group configuration example:

```

A:ALA-12>config>log# info
-----
...
snmp-trap-group 2
trap-target 10.10.10.104:5 "snmpv3" notify-community "communitystring"
    exit
...
log-id 2
    description "This is a test log file."
    filter 1
    from main security
    to file 1
    exit
...
-----
A:ALA-12>config>log#

```

The following displays a SNMP trap group, log, and interface configuration examples:

```
A:SetupCLI>config>log# snmp-trap-group 44
A:SetupCLI>config>log>snmp-trap-group# info
-----
      trap-target "xyz-test" address xx.xx.x.x snmpv2c notify-community "xyztesting"
      trap-target "test2" address xx.xx.xx.x snmpv2c notify-community "xyztesting"
-----
*A:SetupCLI>config>log>log-id# info
-----
      from main
      to snmp
-----
*A:SetupCLI>config>router# interface xyz-test
*A:SetupCLI>config>router>if# info
-----
      address xx.xx.xx.x/24
      port 1/1/1
-----
*A:SetupCLI>config>router>if#
```

5.11.6.1 Setting the Replay Parameter

For this example the replay parameter was set by a SNMP SET request for the trap-target address 10.10.10.3 which is bound to port-id 1/1/1.

```
A:SetupCLI>config>log>snmp-trap-group 44
A:SetupCLI>config>log>snmp-trap-group# info
-----
trap-target "xyz-test" address 10.10.10.3 snmpv2c notify-
community "xyztesting" replay
trap-target "test2" address 20.20.20.5 snmpv2c notify-community "xyztesting"
-----
A:SetupCLI>config>log>snmp-trap-group#
```

In the following output, the **Replay** field changed from disabled to enabled.

```
A:SetupCLI>config>log>snmp-trap-group# show log snmp-trap-group 44
=====
SNMP Trap Group 44
=====
Description : none
-----
Name       : xyz-test
Address    : 10.10.10.3
Port       : 162
Version    : v2c
Community  : xyztesting
Sec. Level : none
Replay     : enabled
Replay from : n/a
Last replay : never
```

```

-----
Name       : test2
Address    : 20.20.20.5
Port       : 162
Version    : v2c
Community  : xyztesting
Sec. Level : none
Replay     : disabled
Replay from : n/a
Last replay : never
=====

```

```
A:SetupCLI>config>log>snmp-trap-group#
```

Since no events are waiting to be replayed, the log displays as before.

```

A:SetupCLI>config>log>snmp-trap-group# show log log-id 44
=====
Event Log 44
=====
SNMP Log contents [size=100 next event=3819 (wrapped)]

3818 2008/04/22 23:35:39.89 UTC WARNING: SYSTEM #2009 Base IP
"Status of vRtrIfTable: router Base (index 1) interface xyz-test (index 35) changed
administrative state: inService, operational state: inService"

3817 2008/04/22 23:35:39.89 UTC WARNING: SNMP #2005 Base xyz-test
"Interface xyz-test is operational"

3816 2008/04/22 23:35:39.89 UTC WARNING: SNMP #2005 Base 1/1/1
"Interface 1/1/1 is operational"

3815 2008/04/22 23:35:39.71 UTC WARNING: SYSTEM #2009 Base CHASSIS
"Status of Mda 1/1 changed administrative state: inService, operational state:
inService"

3814 2008/04/22 23:35:38.88 UTC MINOR: CHASSIS #2002 Base Mda 1/2
"Class MDA Module : inserted"

3813 2008/04/22 23:35:38.88 UTC MINOR: CHASSIS #2002 Base Mda 1/1

```

5.11.6.2 Shutdown In-Band Port

A **shutdown** on the in-band port that the trap-target address is bound to causes the route to that particular trap target to be removed from the route table. When the SNMP module is notified of this event, it marks the trap-target as inaccessible and saves the sequence-id of the first SNMP notification that will be missed by the trap-target.

```
Example: config>log>snmp-trap-group# exit all
#configure port 1/1/1 shutdown
#
# tools perform log test-event
#
```

The **Replay from** field is updated with the sequence-id of the first event that will be replayed when the trap-target address is added back to the route table.

```
*A:SetupCLI# show log snmp-trap-group 44
=====
SNMP Trap Group 44
=====
Description : none
-----
Name       : xyz-test
Address    : 10.10.10.3
Port      : 162
Version   : v2c
Community  : xyztesting
Sec. Level : none
Replay    : enabled
Replay from : event #3819
Last replay : never
-----
Name       : test2
Address    : 20.20.20.5
Port      : 162
Version   : v2c
Community  : xyztesting
Sec. Level : none
Replay    : disabled
Replay from : n/a
Last replay : never
=====
*A:SetupCLI#
```

A display of the event log indicates which trap targets are not accessible and waiting for notification replay and the sequence ID of the first notification that will be replayed.



Note: If there are more missed events than the log size, the replay will actually start from the first available missed event.

```
*A:SetupCLI# show log log-id 44
=====
Event Log 44
=====
SNMP Log contents [size=100 next event=3821 (wrapped)]
Cannot send to SNMP target address 10.10.10.3.
Waiting to replay starting from event #3819
```

```

3820 2008/04/22 23:41:28.00 UTC INDETERMINATE: LOGGER #2011 Base Event Test
"Test event has been generated with system object identifier tmnxModelSR12Reg.
System description: TiMOS-B-0.0.private both/i386 Nokia 7750 SR Copyright (c)
2000-2016 Nokia. All rights reserved. All use subject to applicable license
agreements. Built on Tue Apr 22 14:41:18 PDT 2008 by test123 in /test123/ws/panos/
main"

3819 2008/04/22 23:41:20.37 UTC WARNING: MC_REDUNDANCY #2022 Base operational state
of peer chan*
"The MC-Ring operational state of peer 2.2.2.2 changed to outOfService."

3818 2008/04/22 23:35:39.89 UTC WARNING: SYSTEM #2009 Base IP
"Status of vRtrIfTable: router Base (index 1) interface xyz-test (index 35) changed
administrative state: inService, operational state: inService"

3823 2008/04/22 23:41:49.82 UTC WARNING: SNMP #2005 Base xyz-test
"Interface xyz-test is operational"

```

5.11.6.3 No Shutdown Port

A **no shutdown** command executed on the in-band port to which the trap-target address is bound will cause the route to that trap target to be re-added to the route table. When the SNMP trap module is notified of this event, it resends the notifications that were missed while there was no route to the trap-target address.

Example: configure# port 1/1/1 no shutdown

tools perform log test-event

After the notifications have been replayed the **Replay from** field indicates n/a because there are no more notifications waiting to be replayed and the **Last replay** field timestamp has been updated.

```

*A:SetupCLI# show log snmp-trap-group 44
=====
SNMP Trap Group 44
=====
Description : none
-----
Name       : xyz-test
Address    : 10.10.10.3
Port      : 162
Version   : v2c
Community  : xyztesting
Sec. Level : none
Replay    : enabled
Replay from : n/a
Last replay : 04/22/2008 18:52:36
-----
Name       : test2
Address    : 20.20.20.5
Port      : 162

```



```
Version      : v2c
Community    : xyztesting
Sec. Level   : none
Replay       : disabled
Replay from  : n/a
Last replay  : never
```

```
=====
*A:SetupCLI#
```

A display of the event log shows that it is no longer waiting to replay notifications to one or more of its trap target addresses. An event message has been written to the logger that indicates the replay to the trap-target address has happened and displays the notification sequence ID of the first and last replayed notifications.

```
*A:SetupCLI# show log log-id 44
```

```
=====
Event Log 44
```

```
=====
SNMP Log contents [size=100 next event=3827 (wrapped)]
```

```
3826 2008/04/22 23:42:02.15 UTC MAJOR: LOGGER #2015 Base Log-id 44
"Missed events 3819 to 3825 from Log-id 44 have been resent to SNMP notification
target address 10.10.10.3."
```

```
3825 2008/04/22 23:42:02.15 UTC INDETERMINATE: LOGGER #2011 Base Event Test
"Test event has been generated with system object identifier tmnxModelSR12Reg.
System description: TiMOS-B-0.0.private both/i386 Nokia 7750 SR Copyright (c)
2000-2016 Nokia.
All rights reserved. All use subject to applicable license agreements.
Built on Tue Apr 22 14:41:18 PDT 2008 by test123 in /test123/ws/panos/main"
```

```
3824 2008/04/22 23:41:49.82 UTC WARNING: SYSTEM #2009 Base IP
"Status of vRtrIfTable: router Base (index 1) interface xyz-test (index 35) changed
administrative state: inService, operational state: inService"
```

```
3823 2008/04/22 23:41:49.82 UTC WARNING: SNMP #2005 Base xyz-test
"Interface xyz-test is operational"
```

5.11.7 Configuring a Syslog Target

Log events cannot be sent to a syslog target host until a valid syslog ID exists.

The following displays a syslog configuration example:

```
A:ALA-12>config>log# info
-----
...
    syslog 1
      description "This is a syslog file."
      address 10.10.10.104
      facility user
      level warning
    exit
```

```

...
-----
A:ALA-12>config>log#

```

5.11.7.1 Configuring an Accounting Custom Record

```

A:ALA-48>config>subscr-mgmt>acct-plcy# info
-----
..
    custom-record
    queue 1
    i-counters
        high-octets-discarded-count
        low-octets-discarded-count
        in-profile-octets-forwarded-count
        out-profile-octets-forwarded-count
    exit
    e-counters
        in-profile-octets-forwarded-count
        in-profile-octets-discarded-count
        out-profile-octets-forwarded-count
        out-profile-octets-discarded-count
    exit
    significant-change 20
    ref-queue all
    i-counters
        in-profile-packets-forwarded-count
        out-profile-packets-forwarded-count
    exit
    e-counters
        in-profile-packets-forwarded-count
        out-profile-packets-forwarded-count
    exit
    exit
..
-----
A:ALA-48>config>subscr-mgmt>acct-plcy#

```

The following is an example custom record configuration.

```

Dut-C>config>log>acct-policy>cr# info
-----
    aa-specific
    aa-sub-counters
        short-duration-flow-count
        medium-duration-flow-count
        long-duration-flow-count
        total-flow-duration
        total-flows-completed-count
    exit
    from-aa-sub-counters
        flows-admitted-count
        flows-denied-count
        flows-active-count

```

```
        packets-admitted-count
        octets-admitted-count
        packets-denied-count
        octets-denied-count
        max-throughput-octet-count
        max-throughput-packet-count
        max-throughput-timestamp
        forwarding-class
    exit
to-aa-sub-counters
    flows-admitted-count
    flows-denied-count
    flows-active-count
    packets-admitted-count
    octets-admitted-count
    packets-denied-count
    octets-denied-count
    max-throughput-octet-count
    max-throughput-packet-count
    max-throughput-timestamp
    forwarding-class
exit
exit
significant-change 1
ref-aa-specific-counter any
-----
```

5.12 Log Configuration Command Reference

This section provides the log configuration command reference.

5.12.1 Command Hierarchies

- [Log Configuration Command Reference](#)
 - [Log Configuration Commands](#)
 - [Accounting Policy Commands](#)
 - [Custom Record Commands](#)
 - [File ID Commands](#)
 - [Event Filter Commands](#)
 - [Event Handling System \(EHS\) Commands](#)
 - [Event Trigger Commands](#)
 - [Log ID Commands](#)
 - [SNMP Trap Group Commands](#)
 - [Syslog Commands](#)
 - [Show Commands](#)
 - [Clear Command](#)

5.12.1.1 Log Configuration Commands

```
config
  — log
    — app-route-notifications
      — cold-start-wait seconds
      — no cold-start-wait
      — route-recovery-wait seconds
      — no route-recovery-wait
    — event-control application-id [event-name | event-number] [generate [severity-level]
      [throttle] [specific-throttle-rate events-limit interval seconds | disable-specific-throttle]
    — event-control application-id [event-name | event-number] suppress
    — no event-control application [event-name | event-number]
    — [no] event-damping
    — route-preference primary {inband | outband} secondary {inband | outband | none}
    — no route-preference
    — throttle-rate events [interval seconds]
    — no throttle-rate
```

5.12.1.2 Accounting Policy Commands

```

config
  — log
    — collection-interval minutes
    — no collection-interval
    — accounting-policy acct-policy-id
    — no accounting-policy acct-policy-id
      — [no] auto-bandwidth
      — [no] default
      — description description-string
      — no description
      — [no] include-router-info
      — [no] include-system-info
      — record record-name
      — no record
      — [no] shutdown
      — to file log-file-id

```

5.12.1.3 Custom Record Commands

```

config
  — log
    — accounting-policy acct-policy-id [interval minutes]
    — no accounting-policy acct-policy-id
      — collection-interval minutes
      — no collection-interval
      — [no] custom-record
        — [no] aa-specific
          — aa-sub-counters [all]
          — no aa-sub-counters
            — [no] long-duration-flow-count
            — [no] medium-duration-flow-count
            — [no] short-duration-flow-count
            — [no] total-flow-duration
            — [no] total-flows-completed-count
          — from-aa-sub-counters [all]
          — no from-aa-sub-counters
            — all
            — [no] flows-active-count [all]
            — [no] flows-admitted-count
            — [no] flows-denied-count
            — [no] forwarding-class
            — [no] max-throughput-octet-count
            — [no] max-throughput-packet-count
            — [no] max-throughput-packet-count
            — [no] octets-admitted-count
            — [no] octets-denied-count
            — [no] packets-admitted-count
            — [no] packets-denied-count

```

- **to-aa-sub-counters** [all]
- **to-aa-sub-counters**
 - all
 - [no] **flows-active-count** [all]
 - [no] **flows-admitted-count**
 - [no] **flows-denied-count**
 - [no] **forwarding-class**
 - [no] **max-throughput-octet-count**
 - [no] **max-throughput-packet-count**
 - [no] **max-throughput-packet-count**
 - [no] **octets-admitted-count**
 - [no] **octets-denied-count**
 - [no] **packets-admitted-count**
 - [no] **packets-denied-count**
- [no] **override-counter** *override-counter-id*
 - **e-counters** [all]
 - **no e-counters**
 - [no] **in-profile-octets-discarded-count**
 - [no] **in-profile-octets-forwarded-count**
 - [no] **in-profile-packets-discarded-count**
 - [no] **in-profile-packets-forwarded-count**
 - [no] **out-profile-octets-discarded-count**
 - [no] **out-profile-octets-forwarded-count**
 - [no] **out-profile-packets-discarded-count**
 - [no] **out-profile-packets-forwarded-count**
 - **i-counters** [all]
 - **no i-counters**
 - [no] **in-profile-octets-discarded-count**
 - [no] **in-profile-octets-forwarded-count**
 - [no] **in-profile-packets-discarded-count**
 - [no] **in-profile-packets-forwarded-count**
 - [no] **out-profile-octets-discarded-count**
 - [no] **out-profile-octets-forwarded-count**
 - [no] **out-profile-packets-discarded-count**
 - [no] **out-profile-packets-forwarded-count**
- [no] **queue** *queue-id*
 - **e-counters** [all]
 - **no e-counters**
 - [no] **in-profile-octets-discarded-count**
 - [no] **in-profile-octets-forwarded-count**
 - [no] **in-profile-packets-discarded-count**
 - [no] **in-profile-packets-forwarded-count**
 - [no] **out-profile-octets-discarded-count**
 - [no] **out-profile-octets-forwarded-count**
 - [no] **out-profile-packets-discarded-count**
 - [no] **out-profile-packets-forwarded-count**
 - **i-counters** [all]
 - **no i-counters**
 - [no] **all-octets-offered-count**
 - [no] **all-packets-offered-count**
 - [no] **high-octets-discarded-count**
 - [no] **high-octets-offered-count**
 - [no] **high-packets-discarded-count**
 - [no] **high-packets-offered-count**

- [no] **in-profile-octets-forwarded-count**
- [no] **in-profile-packets-forwarded-count**
- [no] **low-octets-discarded-count**
- [no] **low-packets-discarded-count**
- [no] **low-octets-offered-count**
- [no] **low-packets-offered-count**
- [no] **out-profile-octets-forwarded-count**
- [no] **out-profile-packets-forwarded-count**
- [no] **uncoloured-octets-offered-count**
- [no] **uncoloured-packets-offered-count**
- **ref-aa-specific-counter any**
- **no ref-aa-specific-counter**
- **ref-override-counter** *ref-override-counter-id*
- **ref-override-counter all**
- **no ref-override-counter**
 - **e-counters [all]**
 - **no e-counters**
 - [no] **in-profile-octets-discarded-count**
 - [no] **in-profile-octets-forwarded-count**
 - [no] **in-profile-packets-discarded-count**
 - [no] **in-profile-packets-forwarded-count**
 - [no] **out-profile-octets-discarded-count**
 - [no] **out-profile-octets-forwarded-count**
 - [no] **out-profile-packets-discarded-count**
 - [no] **out-profile-packets-forwarded-count**
 - **i-counters [all]**
 - **no i-counters**
 - [no] **all-octets-offered-count**
 - [no] **all-packets-offered-count**
 - [no] **high-octets-discarded-count**
 - [no] **high-octets-offered-count**
 - [no] **high-packets-discarded-count**
 - [no] **high-packets-offered-count**
 - [no] **in-profile-octets-forwarded-count**
 - [no] **in-profile-packets-forwarded-count**
 - [no] **low-octets-discarded-count**
 - [no] **low-packets-discarded-count**
 - [no] **low-octets-offered-count**
 - [no] **low-packets-offered-count**
 - [no] **out-profile-octets-forwarded-count**
 - [no] **out-profile-packets-forwarded-count**
 - [no] **uncoloured-octets-offered-count**
 - [no] **uncoloured-packets-offered-count**
- **ref-queue** *queue-id*
- **ref-queue all**
- **no ref-queue**
 - **e-counters [all]**
 - **no e-counters**
 - [no] **in-profile-octets-discarded-count**
 - [no] **in-profile-octets-forwarded-count**
 - [no] **in-profile-packets-discarded-count**
 - [no] **in-profile-packets-forwarded-count**
 - [no] **out-profile-octets-discarded-count**
 - [no] **out-profile-octets-forwarded-count**

- [no] **out-profile-packets-discarded-count**
- [no] **out-profile-packets-forwarded-count**
- **i-counters** [all]
- **no i-counters**
 - [no] **all-octets-offered-count**
 - [no] **all-packets-offered-count**
 - [no] **high-octets-discarded-count**
 - [no] **high-octets-offered-count**
 - [no] **high-packets-discarded-count**
 - [no] **high-packets-offered-count**
 - [no] **in-profile-octets-forwarded-count**
 - [no] **in-profile-packets-forwarded-count**
 - [no] **low-octets-discarded-count**
 - [no] **low-packets-discarded-count**
 - [no] **low-octets-offered-count**
 - [no] **low-packets-offered-count**
 - [no] **out-profile-octets-forwarded-count**
 - [no] **out-profile-packets-forwarded-count**
- **significant-change** *delta*
- **no significant-change**

5.12.1.4 File ID Commands

- ```

config
 — log
 — [no] file-id log-file-id
 — description description-string
 — no description
 — location cflash-id [backup-cflash-id]
 — rollover minutes [retention hours]
 — no rollover

```

### 5.12.1.5 Event Filter Commands

Refer to the SR OS Services Guide for information about configuring log filters in a VPRN service.

- ```

config
  — log
    — [no] filter filter-id
      — default-action {drop | forward}
      — no default-action
      — description description-string
      — no description
      — [no] entry entry-id
        — action {drop | forward}
        — no action
        — description description-string
  
```

- no **description**
- [no] **match**
 - **application** {eq | neq} *application-id*
 - no **application**
 - **message** {eq | neq} *pattern* [**regexp**]
 - no **message**
 - **number** {eq | neq | lt | lte | gt | gte} *event-id*
 - no **number**
 - **router** {eq | neq} *router-instance* [**regexp**]
 - no **router**
 - **severity** {eq | neq | lt | lte | gt | gte} *severity-level*
 - no **severity**
 - **subject** {eq | neq} *subject* [**regexp**]
 - no **subject**

5.12.1.6 Event Handling System (EHS) Commands

- ```

config
 — log
 — event-handling
 — [no] handler event-handler-name
 — action-list
 — [no] entry entry-id
 — description description-string
 — no description
 — min-delay [delay]
 — no min-delay
 — [no] script-policy script-policy-name [owner owner-name]
 — no script-policy
 — description description-string
 — no description
 — [no] shutdown

```

### 5.12.1.7 Event Trigger Commands

- ```

config
  — log
    — event-trigger
      — [no] event application-id event-name-id
        — description description-string
        — no description
        — [no] shutdown
      — [no] trigger-entry entry-id
        — debounce occurrences [within seconds]
        — no debounce
        — event-handler event-handler-name
  
```

- [no] **event-handler**
- **description** *description-string*
- **no description**
- **log-filter** *filter-id*
- [no] **log-filter**

5.12.1.8 Log ID Commands

Refer to the SR OS Services Guide for information about configuring logs in a VPRN service.

```

config
  — log
    — [no] log-id log-id
      — description description-string
      — no description
      — filter filter-id
      — no filter
      — from {[main] [security] [change] [debug-trace]}
      — no from
      — python-policy policy-name
      — no python-policy
      — [no] shutdown
      — time-format {local | utc}
      — to console
      — to file log-file-id
      — to memory [size]
      — to session
      — to snmp [size]
      — to syslog syslog-id

```

5.12.1.9 SNMP Trap Group Commands

Refer to the SR OS Services Guide for information about configuring SNMP trap groups in a VPRN service.

```

config
  — log
    — [no] snmp-trap-group log-id
      — description description-string
      — no description
      — trap-target name [address ip-address] [port port] [snmpv1 | snmpv2c | snmpv3] notify-community communityName | snmpv3SecurityName [security-level {no-auth-no-privacy | auth-no-privacy | privacy}] [replay]
      — no trap-target name

```

5.12.1.10 Syslog Commands

Refer to the SR OS Services Guide for information about configuring syslogs in a VPRN service.

```
config
  — log
    — [no] syslog syslog-id
      — address ip-address
      — no address
      — description description-string
      — no description
      — facility syslog-facility
      — no facility
      — level {emergency | alert | critical | error | warning | notice | info | debug}
      — no level
      — log-prefix log-prefix-string
      — no log-prefix
      — port port
      — no port
```

5.12.2 Command Descriptions

- [Generic Commands](#)
- [File ID Commands](#)
- [Log Filter Commands](#)
- [Log Filter Entry Commands](#)
- [Log Filter Entry Match Commands](#)
- [Event Handling System \(EHS\) Commands](#)
- [Event Trigger Commands](#)
- [Syslog Commands](#)
- [SNMP Trap Groups](#)
- [Accounting Policy Commands](#)

5.12.2.1 Generic Commands

description

Syntax	description <i>string</i> no description
Context	config>log>filter config>log>filter>entry config>log>log-id config>log>accounting-policy config>log>event-handling>handler config>log>event-handling>handler>action-list>entry config>log>event-trigger>event config>log>event-trigger>event>trigger-entry config>log>file-id config>log>syslog config>log>snmp-trap-group
Description	This command creates a text description stored in the configuration file for a configuration context. The description command associates a text string with a configuration context to help identify the content in the configuration file. The no form of the command removes the string from the configuration.
Default	No text description is associated with this configuration. The string must be entered.

Parameters *string* — The description can contain a string of up to 80 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

shutdown

Syntax [no] shutdown

Context config>log>log-id
config>log>accounting-policy
config>log>event-handling>handler
config>log>event-trigger>event

Description This command administratively disables an entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics. The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they may be deleted.

The **no** form of this command administratively enables an entity.

Default no shutdown

Parameters *log-id log-id* — When a *log-id* is shut down, no events are collected for the entity. This leads to the loss of event data.

accounting-policy accounting Policy — When an accounting policy is shut down, no accounting data is written to the destination log ID. Counters in the billing data reflect totals, not increments, so when the policy is re-enabled (**no shutdown**) the counters include the data collected during the period the policy was shut down.

5.12.2.2 Log Configuration Commands

app-route-notifications

Syntax	app-route-notifications
Context	config>log
Description	Specific system applications in SR OS can take action based on a route to certain IP destinations being available. This CLI branch contains configuration related to these route availability notifications. A delay can be configured between the time that a route is determined as available in the CPM, and the time that the application is notified of the available route. For example, this delay may be used to increase the chances that other system modules (such as IOMs/XCMs/MDAs/XMAs) are fully programmed with the new route before the application takes action. Currently, the only application that acts upon these <i>route available</i> or <i>route changed</i> notifications with their configurable delays is the SNMP replay feature, which receives notifications of route availability to the SNMP trap receiver destination IP address.

cold-start-wait

Syntax	cold-start-wait <i>seconds</i> no cold-start-wait
Context	config>log>app-route-notifications
Description	The time delay that must pass before notifying specific CPM applications that a route is available after a cold reboot.
Default	no cold-start-wait
Parameters	<i>seconds</i> — time delay in seconds
	Values 1 to 300
	Default 0

route-recovery-wait

Syntax	route-recovery-wait <i>seconds</i> no route-recovery-wait
Context	config>log>app-route-notifications
Description	The time delay that must pass before notifying specific CPM applications after the recovery or change of a route during normal operation.

Default	no route-recovery-wait
Parameters	<i>seconds</i> — time delay in seconds
Values	1 to 100
Default	0

event-control

Syntax	event-control <i>application-id</i> [event-name <i>event-number</i>] [generate] [<i>severity-level</i>] [throttle] [specific-throttle-rate <i>events-limit interval seconds</i> disable-specific-throttle] event-control <i>application-id</i> [event-name <i>event-number</i>] suppress no event-control <i>application</i> [event-name <i>event-number</i>]
Context	config>log
Description	<p>This command is used to specify that a particular event or all events associated with an application is either generated or suppressed.</p> <p>Events are generated by an application and contain an event number and description explaining the cause of the event. Each event has a default designation which directs it to be generated or suppressed.</p> <p>Events are generated with a default severity level that can be modified by using the <i>severity-level</i> option.</p> <p>Events that are suppressed by default are typically used for debugging purposes. Events are suppressed at the time the application requests the event's generation. No event log entry is generated regardless of the destination. While this feature can save processor resources, there may be a negative effect on the ability to troubleshoot problems if the logging entries are squelched. In reverse, indiscriminate application may cause excessive overhead.</p> <p>The rate of event generation can be throttled by using the throttle parameter.</p> <p>The no form of the command reverts the parameters to the default setting for events for the application or a specific event within the application. The severity, generate, suppress, and throttle options will also be reset to the initial values.</p>
Default	Each event has a set of default settings. To display a list of all events and the current configuration use the event-control command.
Parameters	<p><i>application-id</i> — The application whose events are affected by this event control filter.</p> <p>Values A valid application name. To display a list of valid application names, use the show log applications command.</p> <p>Some examples of valid applications are:</p> <p>Default None, this parameter must be explicitly specified.</p>

event-name | *event-number* — To generate, suppress, or revert to default for a single event, enter the specific number or event short name. If no event number or name is specified, the command applies to all events in the application. To display a list of all event short names use the [event-control](#) command.

Default none

Values A valid event name or event number.

generate — Specifies that logger event is created when this event occurs. The generate keyword can be used with two optional parameters, *severity-level* and **throttle**.

Default generate

severity-name — An ASCII string representing the severity level to associate with the specified generated events

Default The system assigned severity name

Values One of: cleared, indeterminate, critical, major, minor, warning.

throttle — Specifies whether or not events of this type will be throttled. By default, event throttling is on for most event types.

suppress — This keyword indicates that the specified events will not be logged. If the **suppress** keyword is not specified then the events are generated by default. For example on the 7750 SR, **event-control bgp suppress** will suppress all BGP events. If a log event is a raising event for a Facility Alarm, and the associated Facility Alarm is raised, then changing the log event to **suppress** clears the associated Facility Alarm.

Default generate

specific-throttle-rate *events-limit* — The log event throttling rate can be configured independently for each log event using this keyword. This specific-throttle-rate overrides the globally configured throttle rate (**configure>log>throttle-rate**) for the specific log event.

Values 1 to 20000

interval *seconds* — specifies the number of seconds that the specific throttling intervals lasts.

Values 1 to 1200

disable-specific-throttle — Specifies to disable the **specific-throttle-rate**.

event-damping

Syntax [no] **event-damping**

Context config>log

Description This command allows the user to set the event damping algorithm to suppress QoS or filter change events.



Note: While this event damping is original behavior for some modules such as service manager, QoS, and filters, it can result in the NMS system database being out of sync because of missed change events. On the other hand, if the damping is disabled (**no event-damping**), it may take much longer to **exec** a large CLI configuration file after system bootup.

route-preference

Syntax	route-preference primary {inband outband} secondary {inband outband none} no route-preference
Context	config>log
Description	This command specifies the primary and secondary routing preference for traffic generated for SNMP notifications and syslog messages. If the remote destination is not reachable through the routing context specified by primary route preference then the secondary routing preference will be attempted. The no form of the command reverts to the default values.
Default	no route-preference
Parameters	<p>primary — Specifies the primary routing preference for traffic generated for SNMP notifications and syslog messages.</p> <p>Default outband</p> <p>secondary — Specifies the secondary routing preference for traffic generated for SNMP notifications and syslog messages. The routing context specified by the secondary route preference will be attempted if the remote destination was not reachable by the primary routing preference, specified by primary route preference. The value specified for the secondary routing preference must be distinct from the value for primary route preference.</p> <p>Default inband</p> <p>inband — Specifies that the logging utility will attempt to use the base routing context to send SNMP notifications and syslog messages to remote destinations.</p> <p>outband — Specifies that the logging utility will attempt to use the management routing context to send SNMP notifications and syslog messages to remote destinations.</p> <p>none — Specifies that no attempt will be made to send SNMP notifications and syslog messages to remote destinations.</p>

throttle-rate

Syntax	throttle-rate events [interval seconds] no throttle-rate
---------------	---

Context	config>log
Description	This command configures the number of events and interval length to be applied to all event types that have throttling enabled by the event-control command and do not have a specific-throttle-rate configured.
Parameters	<p><i>events</i> — Specifies the number of log events that can be logged within the specified interval for a specific event. Once the limit has been reached, any additional events of that type will be dropped, for example, the event drop count will be incremented. At the end of the throttle interval if any events have been dropped a trap notification will be sent.</p> <p>Values 1 to 20000</p> <p>Default 2000</p> <p><i>interval seconds</i> — Specifies the number of seconds that an event throttling interval lasts.</p> <p>Values 1 to 1200</p> <p>Default 1</p>

5.12.2.3 File ID Commands

file-id

Syntax	[no] file-id <i>file-id</i>
Context	config>log
Description	<p>This command creates the context to configure a file ID template to be used as a destination for an event log or billing file.</p> <p>This command defines the file location and characteristics that are to be used as the destination for a log event message stream or accounting/billing information. The file defined in this context is subsequently specified in the to command under log-id or accounting-policy to direct specific logging or billing source streams to the file destination.</p> <p>A file ID can only be assigned to either <i>one</i> log-id or <i>one</i> accounting-policy. It cannot be reused for multiple instances. A file ID and associated file definition must exist for each log and billing file that must be stored in the file system.</p> <p>A file is created when the file ID defined in this command is selected as the destination type for a specific log or accounting record. Log files are collected in a “log” directory. Accounting files are collected in an “act” directory.</p> <p>The file names for a log are created by the system as summarized in Table 68.</p>

Table 68 Log File Names

File Type	File Name
Log File	log// <i>ff</i> - <i>timestamp</i>
Accounting File	acta// <i>ff</i> - <i>timestamp</i> .gz

Where:

- *ll* is the *log-id*
- *aa* is the accounting *policy-id*
- *ff* is the file-id
- The *timestamp* is the actual timestamp when the file is created. The format for the timestamp is *yyyymmdd-hhmmss* where:
 - *yyyy* is the year (for example, 2006)
 - *mm* is the month number (for example, 12 for December)
 - *dd* is the day of the month (for example, 03 for the 3rd of the month)
 - *hh* is the hour of the day in 24 hour format (for example, 04 for 4 a.m.)
 - *mm* is the minutes (for example, 30 for 30 minutes past the hour)
 - *ss* is the number of seconds (for example, 14 for 14 seconds)
- The accounting file is compressed and has a gz extension.

When initialized, each file will contain:

- *The log-id* description.
- *The* time the file was opened.
- The reason the file was created.
- If the event log file was closed properly, the sequence number of the last event stored on the log is recorded.

If the process of writing to a log file fails (for example, the compact flash card is full) and if a backup location is not specified or fails, the log file will not become operational even if the compact flash card is replaced. Enter either a **clear log** command or a **shutdown/no shutdown** command to reinitialize the file.

If the primary location fails (for example, the compact flash card fills up during the write process), a trap is sent and logging continues to the specified backup location. This can result in truncated files in different locations.

The **no** form of the command removes the *file-id* from the configuration. A *file-id* can only be removed from the configuration if the file is not the designated output for a log destination. The actual file remains on the file system.

Default No default file IDs are defined.

Parameters *file-id* — The file identification number for the file, expressed as a decimal integer.

Values 1 to 99

location

Syntax **location** *cflash-id* [*backup-cflash-id*]
no location

Context config>log>file *file-id*

Description This command specifies the primary and optional backup location where the log or billing file will be created.

The **location** command is optional. If the location command not explicitly configured, log files will be created on cf1: and accounting files will be created on cf2: without overflow onto other devices. Generally, cf3: is reserved for system files (configurations, images, etc.).

When multiple location commands are entered in a single file ID context, the last command overwrites the previous command.

When the location of a file ID that is associated with an active log ID is changed, the log events are not immediately written to the new location. The new location does not take affect until the log is rolled over either because the rollover period has expired or a **clear log log-id** command is entered to manually rollover the log file.

When creating files, the primary location is used as long as there is available space. If no space is available, an attempt is made to delete unnecessary files that are past their retention date.

If sufficient space is not available an attempt is made to remove the oldest to newest closed log or accounting files. After each file is deleted, the system attempts to create the new file.

A medium severity trap is issued to indicate that a compact flash is either not available or that no space is available on the specified flash and that the backup location is being used.

A high priority alarm condition is raised if none of the configured compact flash devices for this file ID are present or if there is insufficient space available. If space does becomes available, then the alarm condition will be cleared.

Use the **no** form of this command to revert to default settings.

Default Log files are created on cf1: and accounting files are created on cf2:

Parameters *cflash-id* — Specify the primary location.

Values cflash-id:cf1:, cf2:, cf3:

backup-cflash-id — Specify the secondary location.

Values cflash-id: cf1:, cf2:, cf3:

rollover

Syntax	rollover <i>minutes</i> [retention <i>hours</i>] no rollover
Context	config>log>file <i>file-id</i>
Description	<p>This command configures how often an event or accounting log is rolled over or partitioned into a new file.</p> <p>An event or accounting log is actually composed of multiple, individual files. The system creates a new file for the log based on the rollover time, expressed in minutes.</p> <p>The retention option, expressed in hours, allows you to modify the default time to keep the file in the system. The retention time is based on the rollover time of the file.</p> <p>When multiple rollover commands for a <i>file-id</i> are entered, the last command overwrites the previous command.</p>
Default	rollover 1440 retention 12
Parameters	<p><i>minutes</i> — The rollover time, in minutes.</p> <p>Values 5 to 10080</p> <p><i>retention hours</i> — The retention period in hours, expressed as a decimal integer. The retention time is based on the time creation time of the file. The file becomes a candidate for removal once the creation datestamp + rollover time + retention time is less than the current timestamp.</p> <p>Default 12</p> <p>Values 1 to 500</p>

5.12.2.4 Log Filter Commands

filter

Syntax	<code>[no] filter filter-id</code>
Context	<code>config>log</code>
Description	<p>This command creates a context for an event filter. An event filter specifies whether to forward or drop an event or trap based on the match criteria.</p> <p>Filters are configured in the filter filter-id context and then applied to a log in the log-id log-id context. Only events for the configured log source streams destined to the log ID where the filter is applied are filtered.</p> <p>Any changes made to an existing filter, using any of the sub-commands, are immediately applied to the destinations where the filter is applied.</p> <p>The no form of the command removes the filter association from log IDs which causes those logs to forward all events.</p>
Default	No event filters are defined.
Parameters	<i>filter-id</i> — The filter ID uniquely identifies the filter. Values 1 to 1000

default-action

Syntax	<code>default-action {drop forward}</code> <code>no default-action</code>
Context	<code>config>log>filter filter-id</code>
Description	<p>The default action specifies the action that is applied to events when no action is specified in the event filter entries or when an event does not match the specified criteria.</p> <p>When multiple default-action commands are entered, the last command overwrites the previous command.</p> <p>The no form of the command reverts the default action to the default value (forward).</p>
Default	<code>default-action forward</code>
Parameters	drop — The events which are not explicitly forwarded by an event filter match are dropped. forward — The events which are not explicitly dropped by an event filter match are forwarded.

5.12.2.5 Log Filter Entry Commands

action

Syntax	action {drop forward} no action
Context	config>log>filter filter-id>entry entry-id
Description	<p>This command specifies a drop or forward action associated with the filter entry. If neither drop nor forward is specified, the default-action will be used for traffic that conforms to the match criteria. This could be considered a No-Op filter entry used to explicitly exit a set of filter entries without modifying previous actions.</p> <p>Multiple action statements entered will overwrite previous actions.</p> <p>The no form of the command removes the specified action statement.</p>
Default	Action specified by the default-action command will apply.
Parameters	drop — Specifies packets matching the entry criteria will be dropped. forward — Specifies packets matching the entry criteria will be forwarded.

entry

Syntax	[no] entry entry-id
Context	config>log>filter filter-id
Description	<p>This command is used to create or edit an event filter entry. Multiple entries may be created using unique <i>entry-id</i> numbers. The TiMOS implementation exits the filter on the first match found and executes the action in accordance with the action command.</p> <p>Comparisons are performed in an ascending entry ID order. When entries are created, they should be arranged sequentially from the most explicit entry to the least explicit. Matching ceases when a packet matches an entry. The entry action is performed on the packet, either drop or forward. To be considered a match, the packet must meet all the conditions defined in the entry.</p> <p>An entry may not have any match criteria defined (in which case, everything matches) but must have at least the keyword action for it to be considered complete. Entries without the action keyword will be considered incomplete and are rendered inactive.</p> <p>The no form of the command removes the specified entry from the event filter. Entries removed from the event filter are immediately removed from all log-id's where the filter is applied.</p>

Default	No event filter entries are defined. An entry must be explicitly configured.
Parameters	<i>entry-id</i> — The entry ID uniquely identifies a set of match criteria corresponding action within a filter. Entry ID values should be configured in staggered increments so you can insert a new entry in an existing policy without renumbering the existing entries.
Values	1 to 999

5.12.2.6 Log Filter Entry Match Commands

match

Syntax	[no] match
Context	config>log>filter filter-id>entry entry-id
Description	<p>This command creates context to enter/edit match criteria for a filter entry. When the match criteria is satisfied, the action associated with the entry is executed.</p> <p>If more than one match parameter (within one match statement) is specified, then all the criteria must be satisfied (AND functional) before the action associated with the match is executed.</p> <p>Use the application command to display a list of the valid applications.</p> <p>Match context can consist of multiple match parameters (application, event-number, severity, subject), but multiple match statements cannot be entered per entry.</p> <p>The no form of the command removes the match criteria for the <i>entry-id</i>.</p>
Default	No match context is defined.

application

Syntax	application {eq neq} application-id no application
Context	config>log>filter filter-id>entry entry-id>match
Description	<p>This command adds an OS application as an event filter match criterion.</p> <p>An OS application is the software entity that reports the event. Applications include IP, MPLS, OSPF, CLI, SERVICES etc. Only one application can be specified. The latest application command overwrites the previous command.</p> <p>The no form of the command removes the application as a match criterion.</p>

Default	no application
Parameters	eq neq — The operator specifying the type of match. Valid operators are listed in Table 69 .

Table 69 Valid Operators

Operator	Notes
eq	equal to
neq	not equal to

application-id — The application name string.

Values application_assurance, aps, atm, bgp, cflowd, chassis, debug, dhcp, dhcps, diameter, dynsvc, efm_oam, elmi, ering, eth_cfm, etun, fiter, gsmp, igh, igmp, igmp_snooping, ip, ipsec, isis, l2tp, lag, ldp, li, lldp, logger, mcpath, mc_redundancy, mirror, mld, mld_snooping, mpls, mpls_tp, msdp, nat, ntp, oam, open_flow, ospf, pim, pim_snooping, port, ppp, pppoe, ptp, radius, rip, rip_ng, route_policy, rsvp, security, snmp, stp, svcmgr, system, user, video, vrrp, vrtr, wlan_gw, wpp

message

Syntax	message { eq neq } pattern <i>pattern</i> [regex] no message
Context	config>log>filter>entry>match
Description	This command adds system messages as a match criterion. The no form of the command removes messages as a match criterion.
Parameters	eq — Determines if the matching criteria should be equal to the specified value. neq — Determines if the matching criteria should not be equal to the specified value. pattern <i>pattern</i> — Specifies a message up to 400 characters to be used in the match criteria. regex — Specifies the type of string comparison to use to determine if the log event matches the value of message command parameters. When the regex keyword is not specified, the default matching algorithm used is a basic substring match.

number

Syntax **number** {**eq** | **neq** | **lt** | **lte** | **gt** | **gte**} *event-id*

- no number**
- Context** config>log>filter filter-id>entry entry-id>match
- Description** This command adds an SR OS application event number as a match criterion.
- SR OS event numbers uniquely identify a specific logging event within an application.
- Only one **number** command can be entered per event filter entry. The latest **number** command overwrites the previous command.
- The **no** form of the command removes the event number as a match criterion.
- Default** no event-number
- Parameters** **eq | neq | lt | lte | gt | gte** — This operator specifies the type of match. Valid operators are listed in [Table 70](#).

Table 70 Valid Operators

Operator	Notes
eq	equal to
neq	not equal to
lt	less than
lte	less than or equal to
gt	greater than
gte	greater than or equal to

event-id — The event ID, expressed as a decimal integer.

Values 1 to 4294967295

router

- Syntax** **router {eq | neq} router-instance [regex]**
no router
- Context** config>log>filter>entry>match
- Description** This command specifies the log event matches for the router.
- Parameters** **eq** — Determines if the matching criteria should be equal to the specified value.
neq — Determines if the matching criteria should not be equal to the specified value.

router-instance — Specifies a router name up to 32 characters to be used in the match criteria.

regexp — Specifies the type of string comparison to use to determine if the log event matches the value of **router** command parameters. When the **regexp** keyword is specified, the string in the **router** command is a regular expression string that will be matched against the subject string in the log event being filtered.

severity

Syntax	severity { eq neq lt lte gt gte } <i>severity-level</i> no severity
Context	config>log>filter>entry>match
Description	This command adds an event severity level as a match criterion. Only one severity command can be entered per event filter entry. The latest severity command overwrites the previous command. The no form of the command removes the severity match criterion.
Default	no severity
Parameters	eq neq lt lte gt gte — This operator specifies the type of match. Valid operators are listed in Table 71 .

Table 71 Valid Operators

Operator	Notes
eq	equal to
neq	not equal to
lt	less than
lte	less than or equal to
gt	greater than
gte	greater than or equal to

severity-name — The ITU severity level name. [Table 72](#) lists severity names and corresponding numbers per ITU standards M.3100 X.733 & X.21 severity levels.

Table 72 ITU Severity Information

Severity Number	Severity Name
1	cleared
2	indeterminate (info)
3	critical
4	major
5	minor
6	warning

Values cleared, intermediate, critical, major, minor, warning

subject

Syntax **subject** {**eq** | **neq**} *subject* [*regex*]
no subject

Context config>log>filter filter-id>entry entry-id>match

Description This command adds an event subject as a match criterion.

The subject is the entity for which the event is reported, such as a port. In this case the port-id string would be the subject. Only one **subject** command can be entered per event filter entry. The latest **subject** command overwrites the previous command.

The **no** form of the command removes the subject match criterion.

Default no subject

Parameters **eq** | **neq** — This operator specifies the type of match. Valid operators are listed in [Table 73](#).

Table 73 Valid Operators

Operator	Notes
eq	equal to
neg	not equal to

subject — A string used as the subject match criterion.

regex — Specifies the type of string comparison to use to determine if the log event matches the value of **subject** command parameters. When the **regex** keyword is specified, the string in the **subject** command is a regular expression string that will be matched against the subject string in the log event being filtered. When the **regex** keyword is not specified, the **subject** command string is matched exactly by the event filter.

5.12.2.7 Event Handling System (EHS) Commands

event-handling

Syntax	event-handling
Context	config>log
Description	This command enables the context to configure event handling within the Event Handler System (EHS).

handler

Syntax	[no] handler <i>event-handler-name</i>
Context	config>log>event-handling
Description	This command configures an EHS handler. The no form of the command removes the specified EHS handler.
Parameters	<i>event-handler-name</i> — Specifies the name of the EHS handler. Can be up to 32 characters maximum.

action-list

Syntax	action-list
Context	config>log>event-handling>handler
Description	This command enables the context to configure the EHS handler action list.

entry

Syntax	[no] entry <i>entry-id</i>
---------------	-----------------------------------

- Context** config>log>event-handling>handler>action-list
- Description** This command configures an EHS handler action-list entry. A handler can have multiple actions where each action, for example, could request the execution of a different script. When the handler is triggered it will walk through the list of configured actions.
- The **no** form of the command removes the specified EHS handler action-list entry.
- Parameters** *entry-id* — Specifies the identifier of the EHS handler entry.
- Values** 1 to 1500

min-delay

- Syntax** **min-delay** [*delay*]
no min-delay
- Context** config>log>event-handling>handler>action-list>entry
- Description** This command specifies the minimum delay in seconds between subsequent executions of the action specified in this entry. This is useful, for example, to ensure that a script doesn't get triggered to execute too often.
- Parameters** *delay* — Specifies the unit in seconds.
- Default** no min-delay

script-policy

- Syntax** **script-policy** *policy-name* [*owner policy-owner*]
no script-policy
- Context** config>log>event-handling>handler>action-list>entry
- Description** This command configures the script policy parameters to use for this EHS handler action-list entry. The associated script is launched when the handler is triggered.
- Parameters** *policy-name* — Specifies the script policy name. Can be up to 32 characters maximum.
- owner policy-owner* — Specifies the script policy owner. Can be up to 32 characters maximum.
- Default** "TiMOS CLI"

5.12.2.8 Event Trigger Commands

event-trigger

Syntax	event-trigger
Context	config>log
Description	This command enables the context to configure log events as triggers for Event Handling System (EHS) handlers.

event

Syntax	[no] event <i>application-id event-name-id</i>
Context	config>log>event-trigger
Description	This command configures a specific log event as a trigger for one or more EHS handlers. Further matching criteria can be applied to only trigger certain handlers with certain instances of the log event. The no form of the command removes the specified trigger event.
Parameters	<i>application-id</i> — Specifies the type of application that triggers the event. Values application_assurance, aps, atm, bgp, calltrace, cflowd, chassis, debug, dhcp, dhcps, diameter, dynsvc, efm_oam, elmi, ering, eth_cfm, etun, filter, gsmp, gmpls, igh, igmp, igmp_snooping, ip, ipsec, isis, l2tp, lag, ldp, li, lldp, lmp, logger, mcpath, mc_redundancy, mirror, mld, mld_snooping, mpls, mpls_tp, msdp, nat, ntp, oam, open_flow, ospf, pim, pim_snooping, port, ppp, pppoe, radius, rip, rip_ng, route_policy, rsvp, security, snmp, stp, svcmgr, system, user, video, vrrp, vrtr, wlan_gw, wpp <i>event-name-id</i> — Specifies the name or numerical identifier of the event. Values 0 to 4294967295 <i>event-name</i> : 32 characters max

trigger-entry

Syntax	[no] trigger-entry <i>entry-id</i>
Context	config>log>event-trigger>event
Description	This command configures an instance of a trigger for an EHS handler. A trigger entry binds a set of matching criteria for a log event to a particular handler. If the log event occurs in the system and matches the criteria configured in the associated log filter then the handler will be executed.

The **no** form of the command removes the specified trigger entry.

Parameters *entry-id* — Specifies the identifier of the EHS event trigger entry.
Values 1 to 1500

debounce

Syntax **debounce** *occurrences* [**within** *seconds*]
no debounce

Context config>log>event-trigger>event>trigger-entry

Description This command configures when to trigger, for example after one or more event occurrences. The number of occurrences of an event can be bounded by a time window or left open.

Default no debounce

Parameters *occurrences* — specifies the number of times an event must occur for EHS to trigger a response
Values 2 to 15

within seconds — specifies the time window within which a specific event must occur a number of times equivalent to the specified *occurrences* for EHS to trigger a response
Values 1 to 604800

event-handler

Syntax **event-handler** *event-handler*
no event-handler

Context config>log>event-trigger>event>trigger-entry

Description This command configures the event handler to be used for this trigger entry.

Parameters *event-handler* — Specifies the name of the event handler. Can be up to 32 characters maximum.

log-filter

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

Context config>log>event-trigger>event>trigger-entry

Description This command configures the log filter to be used for this trigger entry. The log filter defines the matching criteria that must be met in order for the log event to trigger the handler execution. The log filter is applied to the log event and, if the filtering decision results in a 'forward' action, then the handler is triggered.

It is typically unnecessary to configure match criteria for 'application' or 'number' in the log filter used for EHS since the particular filter is only applied for a specific log event application and number, as configured under **config>log>event-trigger**

Parameters *filter-id* — Specifies the identifier of the filter.

Values 1 to 1500

5.12.2.9 Syslog Commands

syslog

Syntax [**no**] **syslog** *syslog-id*

Context config>log

Description This command creates the context to configure a syslog target host that is capable of receiving selected syslog messages from this network element.

A valid *syslog-id* must have the target syslog host address configured.

A maximum of 10 syslog-id's can be configured.

No log events are sent to a syslog target address until the syslog-id has been configured as the log destination (**to**) in the log-id node.

The syslog ID configured in the **configure/service/vprn** context has a local VPRN scope and only needs to be unique within the specific VPRN instance. The same ID can be reused under a different VPRN service or in the global log context under **config>log**.

Default No syslog IDs are defined.

Parameters *syslog-id* — The syslog ID number for the syslog destination, expressed as a decimal integer.

Values 1 to 10

address

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

Context	config>log>syslog				
Description	<p>This command adds the syslog target host IP address to/from a syslog ID.</p> <p>This parameter is mandatory. If no address is configured, syslog data cannot be forwarded to the syslog target host.</p> <p>Only one address can be associated with a <i>syslog-id</i>. If multiple addresses are entered, the last address entered overwrites the previous address.</p> <p>The same syslog target host can be used by multiple log IDs.</p> <p>The no form of the command removes the syslog target host IP address.</p>				
Default	no address				
Parameters	<p><i>ip-address</i> — The IP address of the syslog target host in dotted decimal notation. ipv6-address apply to the 7750 SR only.</p> <p>Values</p> <table border="0"> <tr> <td style="padding-left: 20px;">ipv4-address</td> <td>a.b.c.d</td> </tr> <tr> <td style="padding-left: 20px;">ipv6-address</td> <td> x:x:x:x:x:x[-interface] x:x:x:x:x:d.d.d.d[-interface] x: [0..FFFF]H d: [0..255]D interface: 32 characters maximum, mandatory for link local addresses ipv6-address x:x:x:x:x:x[-interface] x:x:x:x:x:d.d.d.d[-interface] x: [0..FFFF]H d: [0..255]D interface: 32 characters maximum, mandatory for link local addresses </td> </tr> </table>	ipv4-address	a.b.c.d	ipv6-address	x:x:x:x:x:x[-interface] x:x:x:x:x:d.d.d.d[-interface] x: [0..FFFF]H d: [0..255]D interface: 32 characters maximum, mandatory for link local addresses ipv6-address x:x:x:x:x:x[-interface] x:x:x:x:x:d.d.d.d[-interface] x: [0..FFFF]H d: [0..255]D interface: 32 characters maximum, mandatory for link local addresses
ipv4-address	a.b.c.d				
ipv6-address	x:x:x:x:x:x[-interface] x:x:x:x:x:d.d.d.d[-interface] x: [0..FFFF]H d: [0..255]D interface: 32 characters maximum, mandatory for link local addresses ipv6-address x:x:x:x:x:x[-interface] x:x:x:x:x:d.d.d.d[-interface] x: [0..FFFF]H d: [0..255]D interface: 32 characters maximum, mandatory for link local addresses				

facility

Syntax	<p>facility <i>syslog-facility</i></p> <p>no facility</p>
Context	config>log>syslog
Description	<p>This command configures the facility code for messages sent to the syslog target host.</p> <p>Multiple syslog IDs can be created with the same target host but each syslog ID can only have one facility code. If multiple facility codes are entered, the last <i>facility-code</i> entered overwrites the previous facility-code.</p>

If multiple facilities need to be generated for a single syslog target host, then multiple **log-id** entries must be created, each with its own filter criteria to select the events to be sent to the syslog target host with a given facility code.

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

Default local7

Parameters *syslog-facility* — The syslog facility name represents a specific numeric facility code. The code should be entered in accordance with the syslog RFC. However, the software does not validate if the facility code configured is appropriate for the event type being sent to the syslog target host.

Values kernel, user, mail, systemd, auth, syslogd, printer, netnews, uucp, cron, authpriv, ftp, ntp, logaudit, logalert, cron2, local0, local1, local2, local3, local4, local5, local6, local7

Valid responses per RFC3164, *The BSD syslog Protocol*, are listed in [Table 74](#).

Table 74 Syslog Protocol Valid Responses

Numerical Code	Facility Code
0	kernel
1	user
2	mail
3	systemd
4	auth
5	syslogd
6	printer
7	net-news
8	uucp
9	cron
10	auth-priv
11	ftp
12	ntp
13	log-audit
14	log-alert
15	cron2
16	local0
17	local1
18	local2
19	local3
20	local4
21	local5
22	local6
23	local7

Values 0 to 23

log-prefix

Syntax	log-prefix <i>log-prefix-string</i> no log-prefix
Context	config>log>syslog
Description	<p>This command adds the string prepended to every syslog message sent to the syslog host.</p> <p>RFC3164, <i>The BSD syslog Protocol</i>, allows an alphanumeric string (tag) to be prepended to the content of every log message sent to the syslog host. This alphanumeric string can, for example, be used to identify the node that generates the log entry. The software appends a colon (:) and a space to the string and it is inserted in the syslog message after the date stamp and before the syslog message content.</p> <p>Only one string can be entered. If multiple strings are entered, the last string overwrites the previous string. The alphanumeric string can contain lowercase (a-z), uppercase (A-Z) and numeric (0-9) characters.</p> <p>The no form of the command removes the log prefix string.</p>
Default	no log-prefix
Parameters	<i>log-prefix-string</i> — Specifies an alphanumeric string of up to 32 characters. Spaces and colons (:) cannot be used in the string.

level

Syntax	level <i>syslog-level</i> no level
Context	config>log>syslog
Description	<p>This command configures the syslog message severity level threshold. All messages with severity level equal to or higher than the threshold are sent to the syslog target host.</p> <p>Only a single threshold level can be specified. If multiple levels are entered, the last level entered will overwrite the previously entered commands.</p> <p>The no form of the command reverts to the default value.</p>
Default	level info
Parameters	<i>value</i> — Specifies the threshold severity level name.
	Values emergency, alert, critical, error, warning, notice, info, debug

Table 75 Level Parameter Value Descriptions

Router severity level	Numerical Severity (highest to lowest)	Configured Severity	Definition
	0	emergency	system is unusable
3	1	alert	action must be taken immediately
4	2	critical	critical condition
5	3	error	error condition
6	4	warning	warning condition
	5	notice	normal but significant condition
1 cleared 2 indeterminate	6	info	informational messages
	7	debug	debug-level messages

port

Syntax **port** *value*
no port

Context config>log>syslog

Description This command configures the UDP port that will be used to send syslog messages to the syslog target host.

The port configuration is needed if the syslog target host uses a port other than the standard UDP syslog port 514.

Only one port can be configured. If multiple **port** commands are entered, the last entered port overwrites the previously entered ports.

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

Default no port

Parameters *value* — The value is the configured UDP port number used when sending syslog messages.

Values 1 to 65535

5.12.2.10 SNMP Trap Groups

snmp-trap-group

Syntax	[no] snmp-trap-group <i>log-id</i>
Context	config>log
Description	<p>This command creates the context to configure a group of SNMP trap receivers and their operational parameters for a given log-id.</p> <p>A group specifies the types of SNMP traps and specifies the log ID which will receive the group of SNMP traps. A trap group must be configured in order for SNMP traps to be sent.</p> <p>To suppress the generation of all alarms and traps see the event-control command. To suppress alarms and traps that are sent to this log-id, see the filter command. Once alarms and traps are generated they can be directed to one or more SNMP trap groups. Logger events that can be forwarded as SNMP traps are always defined on the main event source.</p> <p>The no form of the command deletes the SNMP trap group.</p>
Default	There are no default SNMP trap groups.
Parameters	<p><i>log-id</i> — The log ID value of a log configured in the log-id context. Alarms and traps cannot be sent to the trap receivers until a valid <i>log-id</i> exists.</p> <p>Values 1 to 99</p>

trap-target

Syntax	<p>trap-target <i>name</i> [address <i>ip-address</i>] [port <i>port</i>] [snmpv1 snmpv2c snmpv3] notify-community <i>communityName</i> <i>snmpv3SecurityName</i> [security-level {no-auth-no-privacy auth-no-privacy privacy}] [replay]</p> <p>no trap-target <i>name</i></p>
Context	config>log>snmp-trap-group
Description	<p>This command adds/modifies a trap receiver and configures the operational parameters for the trap receiver. A trap reports significant events that occur on a network device such as errors or failures.</p> <p>Before an SNMP trap can be issued to a trap receiver, the log-id, snmp-trap-group and at least one trap-target must be configured.</p> <p>The trap-target command is used to add/remove a trap receiver from an snmp-trap-group. The operational parameters specified in the command include:</p> <ul style="list-style-type: none"> • The IP address of the trap receiver

- The UDP port used to send the SNMP trap
- SNMP version
- SNMP community name for SNMPv1 and SNMPv2c receivers.
- Security name and level for SNMPv3 trap receivers.

A single **snmp-trap-group** *log-id* can have multiple trap-receivers. Each trap receiver can have different operational parameters.

An address can be configured as a trap receiver more than once as long as a different port is used for each instance.

To prevent resource limitations, only configure a maximum of 10 trap receivers.



Note: If the same **trap-target** *name* **port** *port* parameter value is specified in more than one SNMP trap group, each trap destination should be configured with a different *notify-community* value. This allows a trap receiving an application, such as NMS, to reconcile a separate event sequence number stream for each router event log when multiple event logs are directed to the same IP address and port destination.

The **no** form of the command removes the SNMP trap receiver from the SNMP trap group.

Default No SNMP trap targets are defined.

Parameters *name* — Specifies the name of the trap target up to 28 characters in length.

address *ip-address* — The IP address of the trap receiver in dotted decimal notation. Only one IP address destination can be specified per trap destination group. *ipv6* applies to the 7750 SR only.

Values

ipv4-address	a.b.c.d (host bits must be 0)
ipv6-address	x:x:x:x:x:x[-interface] x:x:x:x:x:d.d.d.d[-interface] x: [0..FFFF]H d: [0..255]D interface: 32 characters maximum, mandatory for link local addresses

port *port* — The destination UDP port used for sending traps to the destination, expressed as a decimal integer. Only one port can be specified per **trap-target** statement. If multiple traps need to be issued to the same address then multiple ports must be configured.

Default 162

Values 1 to 65535

snmpv1 | *snmpv2c* | *snmpv3* — Specifies the SNMP version format to use for traps sent to the trap receiver.

The keyword **snmpv1** selects the SNMP version 1 format. When specifying **snmpv1**, the **notify-community** must be configured for the proper SNMP community string that the trap receiver expects to be present in alarms and traps messages. If the SNMP version is changed from **snmpv3** to **snmpv1**, then the **notify-community** parameter must be changed to reflect the community string rather than the *security-name* that is used by **snmpv3**.

The keyword **snmpv2c** selects the SNMP version 2c format. When specifying **snmpv2c**, the **notify-community** must be configured for the proper SNMP community string that the trap receiver expects to be present in alarms and traps messages. If the SNMP version is changed from **snmpv3** to **snmpv2c**, then the **notify-community** parameter must be changed to reflect the community string rather than the *security-name* that is used by **snmpv3**.

The keyword **snmpv3** selects the SNMP version 3 format. When specifying **snmpv3**, the **notify-community** must be configured for the SNMP *security-name*. If the SNMP version is changed from **snmpv1** or **snmpv2c** to **snmpv3**, then the **notify-community** parameter must be changed to reflect the *security-name* rather than the community string used by **snmpv1** or **snmpv2c**.

Pre-existing conditions are checked before the `snmpv3SecurityName` is accepted. These are:

The user name must be configured.

The v3 access group must be configured.

The v3 notification view must be configured.

Default `snmpv3`

Values `snmpv1, snmpv2c, snmpv3`

notify-community *community* | *security-name* — Specifies the community string for **snmpv1** or **snmpv2c** or the **snmpv3** *security-name*. If the **notify-community** is not configured, then no alarms or traps will be issued for the trap destination. If the SNMP version is modified, the **notify-community** must be changed to the proper form for the SNMP version.

community — The community string as required by the **snmpv1** or **snmpv2c** trap receiver. The community string can be an ASCII string up to 31 characters in length.

security-name — The *security-name* as defined in the `config>system>security>user` context for SNMP v3. The *security-name* can be an ASCII string up to 31 characters in length.

security-level {**no-auth-no-privacy** | **auth-no-privacy** | **privacy**} — Specifies the required authentication and privacy levels required to access the views configured on this node when configuring an **snmpv3** trap receiver.

The keyword **no-auth-no-privacy** specifies no authentication and no privacy (encryption) are required.

The keyword **auth-no-privacy** specifies authentication is required but no privacy (encryption) is required. When this option is configured the *security-name* must be configured for **authentication**.

The keyword **privacy** specifies both authentication and privacy (encryption) is required. When this option is configured the *security-name* must be configured for **authentication** and **privacy**.

Default **no-auth-no-privacy**. This parameter can only be configured if SNMPv3 is also configured.

Values no-auth-no-privacy, auth-no-privacy, privacy

replay — Enable replay of missed events to target. If replay is applied to an SNMP trap target address, the address is monitored for reachability. Reachability is determined by whether or not there is a route in the routing table by which the target address can be reached. Before sending a trap to a target address, the SNMP module asks the PIP module if there is either an in-band or out-of-band route to the target address. If there is no route to the SNMP target address, the SNMP module saves the sequence-id of the first event that will be missed by the trap target. When the routing table changes again so that there is now a route by which the SNMP target address can be reached, the SNMP module replays (for example, retransmits) all events generated to the SNMP notification log while the target address was removed from the route table.



Note: Due to route table change convergence time, it is possible that one or more events may be lost at the beginning or end of a replay sequence. The cold-start-wait and route-recovery-wait timers under the **config>log>app-route-notifications** context can help reduce the probability of lost events.

filter

Syntax **filter** *filter-id*
no filter

Context config>log>log-id log-id

Description This command adds an event filter policy with the log destination.

The **filter** command is optional. If an event filter is not configured, all events, alarms and traps generated by the source stream will be forwarded to the destination.

An event filter policy defines (limits) the events that are forwarded to the destination configured in the log-id. The event filter policy can also be used to select the alarms and traps to be forwarded to a destination **snmp-trap-group**.

The application of filters for debug messages is limited to application and subject only.

Accounting records cannot be filtered using the **filter** command.

Only one filter-id can be configured per log destination.


The **no** form of the command removes the specified event filter from the *log-id*.

Default	no filter
Parameters	<i>filter-id</i> — The event filter policy ID is used to associate the filter with the <i>log-id</i> configuration. The event filter policy ID must already be defined in config>log>filter <i>filter-id</i> .
Values	1 to 1000

from

Syntax	from {[main] [security] [change] [debug-trace]} no from
Context	config>log>log-id log-id
Description	<p>This command selects the source stream to be sent to a log destination.</p> <p>One or more source streams must be specified. The source of the data stream must be identified using the from command before you can configure the destination using the to command. The from command can identify multiple source streams in a single statement (for example: from main change debug-trace).</p> <p>Only one from command may be entered for a single <i>log-id</i>. If multiple from commands are configured, then the last command entered overwrites the previous from command.</p> <p>The no form of the command removes all previously configured source streams.</p>
Default	no from
Parameters	<p>main — Instructs all events in the main event stream to be sent to the destination defined in the to command for this destination <i>log-id</i>. The main event stream contains the events that are not explicitly directed to any other event stream. To limit the events forwarded to the destination, configure filters using the filter command.</p> <p>security — Instructs all events in the security event stream to be sent to the destination defined in the to command for this destination <i>log-id</i>. The security stream contains all events that affect attempts to breach system security such as failed login attempts, attempts to access MIB tables to which the user is not granted access or attempts to enter a branch of the CLI to which access has not been granted. To limit the events forwarded to the destination, configure filters using the filter command.</p> <p>change — Instructs all events in the user activity stream to be sent to the destination configured in the to command for this destination <i>log-id</i>. The change event stream contains all events that directly affect the configuration or operation of this node. To limit the events forwarded to the change stream destination, configure filters using the filter command.</p> <p>debug-trace — Instructs all debug-trace messages in the debug stream to be sent to the destination configured in the to command for this destination <i>log-id</i>. Filters applied to debug messages are limited to application and subject.</p>

log-id

Syntax	[no] log-id <i>log-id</i>
Context	config>log
Description	<p>This command creates a context to configure destinations for event streams.</p> <p>The log-id context is used to direct events, alarms/traps, and debug information to respective destinations.</p> <p>A maximum of 15 logs can be configured.</p> <p>Before an event can be associated with this log-id, the from command identifying the source of the event must be configured.</p> <p>Only one destination can be specified for a <i>log-id</i>. The destination of an event stream can be an in-memory buffer, console, session, snmp-trap-group, syslog, or file.</p> <p>Use the event-control command to suppress the generation of events, alarms, and traps for all log destinations.</p> <p>An event filter policy can be applied in the log-id context to limit which events, alarms, and traps are sent to the specified log-id.</p> <p>Log-IDs 99 and 100 are created by the agent. Log-ID 99 captures all log messages. Log-ID 100 captures log messages with a severity level of major and above.</p> <p> Note: Log-ID 99 provides valuable information for the admin-tech file. Removing or changing the log configuration may hinder debugging capabilities. It is strongly recommended not to alter the configuration for Log-ID 99.</p> <p>The no form of the command deletes the log destination ID from the configuration.</p>
Default	no log-id
Parameters	<i>log-id</i> — The log ID number, expressed as a decimal integer.
Values	1 to 100

to console

Syntax	to console
Context	config>log>log-id
Description	<p>This command specifies a log ID destination. This parameter is mandatory when configuring a log destination. This command instructs the events selected for the log ID to be directed to the console. If the console is not connected, then all the entries are dropped.</p>

The source of the data stream must be specified in the **from** command prior to configuring the destination with the **to** command.

The **to** command cannot be modified or re-entered. If the destination or maximum size of an SNMP or memory log needs to be modified, the log ID must be removed and then re-created.

Default No destination is specified.

to file

Syntax **to file** *log-file-id*

Context config>log>log-id log-id

Description This command specifies a log ID destination. This parameter is mandatory when configuring a log destination. This command instructs the events selected for the log ID to be directed to a specified file.

The source of the data stream must be specified in the **from** command prior to configuring the destination with the **to** command.

The **to** command cannot be modified or re-entered. If the destination or maximum size of an SNMP or memory log needs to be modified, the log ID must be removed and then re-created.

When the **file-id** location parameter is modified, log files are not written to the new location until a rollover occurs or the log is manually cleared. A rollover can be forced by using the **clear>log** command. Subsequent log entries are then written to the new location. If a rollover does not occur or the log not cleared, the old location remains in effect.

Default No destination is specified.

Parameters *log-file-id* — Instructs the events selected for the log ID to be directed to the *log-file-id*. The characteristics of the *log-file-id* referenced here must have already been defined in the **config>log>file log-file-id** context.

Values 1 to 99

to memory

Syntax **to memory** [*size*]

Context config>log>log-id log-id

Description This command specifies a log ID destination. This parameter is mandatory when configuring a log destination. This command instructs the events selected for the log ID to be directed to a memory log. A memory file is a circular buffer. Once the file is full, each new entry replaces the oldest entry in the log.

The source of the data stream must be specified in the **from** command prior to configuring the destination with the **to** command.

The **to** command cannot be modified or re-entered. If the destination or maximum size of an SNMP or memory log needs to be modified, the log ID must be removed and then re-created.

Default	to memory
Parameters	<i>size</i> — The <i>size</i> parameter indicates the number of events that can be stored in the memory.
Default	100
Values	50 to 1024

to session

Syntax	to session
Context	config>log>log-id log-id
Description	This command specifies a log ID destination. This parameter is mandatory when configuring a log destination. This command instructs the events selected for the log ID to be directed to the current console or telnet session. This command is only valid for the duration of the session. When the session is terminated the “to session” configuration is removed. A log ID with a <i>session</i> destination is saved in the configuration file but the “to session” part is not stored. The source of the data stream must be specified in the from command prior to configuring the destination with the to command. The to command cannot be modified or re-entered. If the destination or maximum size of an SNMP or memory log needs to be modified, the log ID must be removed and then re-created.
Default	none

to snmp

Syntax	to snmp [<i>size</i>]
Context	config>log>log-id log-id
Description	This is one of the commands used to specify the log ID destination. This parameter is mandatory when configuring a log destination. This command instructs the alarms and traps to be directed to the snmp-trap-group associated with <i>log-id</i> . A local circular memory log is always maintained for SNMP notifications sent to the specified snmp-trap-group for the <i>log-id</i> .

The source of the data stream must be specified in the **from** command prior to configuring the destination with the **to** command.

The **to** command cannot be modified or re-entered. If the destination or maximum size of an SNMP or memory log needs to be modified, the log ID must be removed and then re-created.

Default	none
Parameters	<i>size</i> — The <i>size</i> parameter defines the number of events stored in this memory log.
	Values 50 to 1024
	Default 100

to syslog

Syntax	to syslog <i>syslog-id</i>
Context	config>log>log-id
Description	This is one of the commands used to specify the log ID destination. This parameter is mandatory when configuring a log destination. This command instructs the alarms and traps to be directed to a specified syslog. To remain consistent with the standards governing syslog, messages to syslog are truncated to 1k bytes. The source of the data stream must be specified in the from command prior to configuring the destination with the to command. The to command cannot be modified or re-entered. If the destination or maximum size of an SNMP or memory log needs to be modified, the log ID must be removed and then re-created.
Default	none
Parameters	<i>syslog-id</i> — Instructs the events selected for the log ID to be directed to the <i>syslog-id</i> . The characteristics of the <i>syslog-id</i> referenced here must have been defined in the config>log>syslog <i>syslog-id</i> context.
	Values 1 to 10

python-policy

Syntax	python-policy <i>policy-name</i> no python-policy
Context	config>log>log-id
Description	<p>This command associates the Python script with the events sent to this log ID. The Python policy can be associated with the log only if the destination in the log ID is set to syslog.</p> <p>For information about Python policy configuration, refer to the Python Script Support for ESM in the <i>7450 ESS, 7750 SR, and 7950 XRS Triple Play Guide</i> guide.</p> <p>The no form of this command disables Python processing of the events in this log ID.</p>
Default	no python-policy
Parameters	<i>policy-name</i> — Specifies a Python policy name up to 32 characters in length

time-format

Syntax	time-format { local utc }
Context	config>log>log-id
Description	This command specifies whether the time should be displayed in local or Coordinated Universal Time (UTC) format.
Default	utc
Parameters	local — Specifies that timestamps are written in the system's local time. utc — Specifies that timestamps are written using the UTC value. This was formerly called Greenwich Mean Time (GMT) and Zulu time.

5.12.2.11 Accounting Policy Commands

accounting-policy

Syntax	accounting-policy <i>policy-id</i> [interval minutes] no accounting-policy <i>policy-id</i>
Context	config>log
Description	This command creates an access or network accounting policy. An accounting policy defines the accounting records that are created.

Access accounting policies are policies that can be applied to one or more SAPs. Changes made to an existing policy, using any of the sub-commands, are applied immediately to all SAPs where this policy is applied.

If an accounting policy is not specified on a SAP, then accounting records are produced in accordance with the access policy designated as the **default**. If a default access policy is not specified, then no accounting records are collected other than the records for the accounting policies that are explicitly configured.

Only one policy can be regarded as the default access policy. If a policy is configured as the default policy, then a **no default** command must be used to allow the data that is currently being collected to be written before a new access default policy can be configured.

Network accounting policies are policies that can be applied to one or more network ports or SONET/SDH channels. Any changes made to an existing policy, using any of the sub-commands, will be applied immediately to all network ports or SONET/SDH channels where this policy is applied.

If no accounting policy is defined on a network port, accounting records will be produced in accordance with the default network policy as designated with the **default** command. If no network default policy is created, then no accounting records will be collected other than the records for the accounting policies explicitly configured. Default accounting policies cannot be explicitly applied. For example, for **accounting-policy 10**, if default is set, then that policy cannot be used:

```
*A:75>config>service>vpls>spoke-sdp# accounting-policy 10
```

Only one policy can be regarded as the default network policy. If a policy is configured as the default policy, then a **no default** command must be used to allow the data that is currently being collected to be written before a new network default policy can be configured.

The **no** form of the command deletes the policy from the configuration. The accounting policy cannot be removed unless it is removed from all the SAPs, network ports or channels where the policy is applied.

Default No default accounting policy is defined.

Parameters *policy-id* — The policy ID that uniquely identifies the accounting policy, expressed as a decimal integer.
Values 1 to 99

collection-interval

Syntax **collection-interval** *minutes*
no collection-interval

Context config>log>acct-policy

Description This command configures the accounting collection interval.

Parameters *minutes* — Specifies the interval between collections, in minutes.
Values 1 to 120
A range of 1 to 4 is only allowed when the record type is set to SAA.

auto-bandwidth

Syntax [**no**] **auto-bandwidth**

Context config>log>accounting-policy

Description In the configuration of an accounting policy this designates the accounting policy as the one used for auto-bandwidth statistics collection.

Default no auto-bandwidth

default

Syntax [**no**] **default**

Context config>log>accounting-policy

Description This command configures the default accounting policy to be used with all SAPs that do not have an accounting policy.

If no access accounting policy is defined on a SAP, accounting records are produced in accordance with the default access policy. If no default access policy is created, then no accounting records will be collected other than the records for the accounting policies that are explicitly configured.

If no network accounting policy is defined on a network port, accounting records will be produced in accordance with the default network policy. If no network default policy is created, then no accounting records will be collected other than the records for the accounting policies explicitly configured.

Only one access accounting policy ID can be designated as the default access policy. Likewise, only one network accounting policy ID can be designated as the default network accounting policy.

The record name must be specified prior to assigning an accounting policy as default.

If a policy is configured as the default policy, then a **no default** command must be issued before a new default policy can be configured.

The **no** form of the command removes the default policy designation from the policy ID. The accounting policy will be removed from all SAPs or network ports that do not have this policy explicitly defined.

include-router-info

Syntax	[no] include-router-info
Context	config>log>accounting-policy
Description	This command allows operator to optionally include router information at the top of each accounting file generated for a given accounting policy. When the no form of this command is selected, the optional router information is not include at the top of the file.
Default	no include-router-info

include-system-info

Syntax	[no] include-system-info
Context	config>log>accounting-policy
Description	This command allows the operator to optionally include router information at the top of each accounting file generated for a given accounting policy. When the no version of this command is selected, optional router information is not include at the top of the file.
Default	no include-system-info

record

Syntax	[no] record <i>record-name</i>
Context	config>log>accounting-policy

Description This command adds the accounting record type to the accounting policy to be forwarded to the configured accounting file. A record name can only be used in one accounting policy. To obtain a list of all record types that can be configured, use the **show log accounting-records** command.



Note: aa, video and subscriber records are not applicable to the 7950 XRS.

```
A:ALA-49# show log accounting-records
=====
Accounting Policy Records
=====
Record # Record Name                               Def. Interval
-----
1      service-ingress-octets                          5
2      service-egress-octets                          5
3      service-ingress-packets                          5
4      service-egress-packets                          5
5      network-ingress-octets                          15
6      network-egress-octets                          15
7      network-ingress-packets                         15
8      network-egress-packets                         15
9      compact-service-ingress-octets                 5
10     combined-service-ingress                        5
11     combined-network-ing-egr-octets                15
12     combined-service-ing-egr-octets                5
13     complete-service-ingress-egress                5
14     combined-sdp-ingress-egress                    5
15     complete-sdp-ingress-egress                    5
16     complete-subscriber-ingress-egress             5
17     aa-protocol                                     15
18     aa-application                                 15
19     aa-app-group                                   15
20     aa-subscriber-protocol                         15
21     aa-subscriber-application                      15
23     custom-record-subscriber                       5
24     custom-record-service                          5
25     custom-record-aa-sub                           15
26     queue-group-octets                             15
27     queue-group-packets                           15
28     combined-queue-group                          15
29     combined-mpls-lsp-ingress                      5
30     combined-mpls-lsp-egress                       5
31     combined-ldp-lsp-egress                        5
32     saa                                             5
33     complete-pm                                    5
34     video                                           10
35     kpi-system                                     5
36     kpi-bearer-mgmt                                5
37     kpi-bearer-traffic                             5
38     kpi-ref-point                                  5
39     kpi-path-mgmt                                  5
40     kci-iom-3                                       5
41     kci-system                                      5
42     kci-bearer-mgmt                                5
```

43	kci-path-mgmt	5
44	complete-kpi	5
45	complete-kci	5
46	kpi-bearer-group	5
47	kpi-ref-path-group	5
48	kpi-kci-bearer-mgmt	5
49	kpi-kci-path-mgmt	5
50	kpi-kci-system	5
51	complete-kpi-kci	5
52	aa-performance	15
53	complete-ethernet-port	15
54	extended-service-ingress-egress	5
55	complete-network-ing-egr	15
56	aa-partition	15
57	complete-pm	5
0	unknown-record-name	0
59	kpi-bearer-traffic-gtp-endpoint	5
60	kpi-ip-reas	5
61	kpi-radius-group	5
62	kpi-ref-pt-failure-cause-code	5
63	kpi-dhcp-group	5
	complete-pm	5

=====

A:ALA-49#

To configure an accounting policy for access ports, select a service record (for example, `service-ingress-octets`). To change the record name to another service record then the record command with the new record name can be entered and it will replace the old record name.

When configuring an accounting policy for network ports, a network record should be selected. When changing the record name to another network record, the record command with the new record name can be entered and it will replace the old record name.

If the change required modifies the record from network to service or from service to network, then the old record name must be removed using the **no** form of this command.

Only one record may be configured in a single accounting policy. For example, if an accounting-policy is configured with a **access-egress-octets** record, in order to change it to **service-ingress-octets**, use the **no record** command under the accounting-policy to remove the old record and then enter the **service-ingress-octets** record.



Note: Collecting excessive statistics can adversely affect the CPU utilization and take up large amounts of storage space.

The **no** form of the command removes the record type from the policy.

Default no record

Parameters *record-name* — The accounting record name. [Table 76](#) lists the accounting record names available and the default collection interval.

Table 76 Default Collection Interval for Accounting Records

Record Type	Accounting Record Name	Default Interval
1	service-ingress-octets	5
2	service-egress-octets	5
3	service-ingress-packets	5
4	service-egress-packets	5
5	network-ingress-octets	15
6	network-egress-octets	15
7	network-ingress-packets	15
8	network-egress-packets	15
9	compact-service-ingress-octets	5
10	combined-service-ingress	5
11	combined-network-ing-egr-octets	15
12	combined-service-ing-egr-octets	5
13	complete-service-ingress-egress	5
14	combined-sdp-ingress-egress	5
15	complete-sdp-ingress-egress	5
16	complete-subscriber-ingress-egress	5
17	aa-protocol	15
18	aa-application	15
19	aa-app-group	15
20	aa-subscriber-protocol	15
21	aa-subscriber-application	15
23	custom-record-subscriber	5
24	custom-record-service	5
25	custom-record-aa-sub	15
26	queue-group-octets	15
27	queue-group-packets	15
28	combined-queue-group	15

Table 76 Default Collection Interval for Accounting Records (Continued)

Record Type	Accounting Record Name	Default Interval
29	combined-mpls-lsp-ingress	5
30	combined-mpls-lsp-egress	5
31	combined-ldp-lsp-egress	5
32	saa	5
33	complete-pm	5
34	video	10
35	kpi-system	5
36	kpi-bearer-mgmt	5
37	kpi-bearer-traffic	5
38	kpi-ref-point	5
39	kpi-path-mgmt	5
40	kpi-iom-3	5
41	kci-system	5
42	kci-bearer-mgmt	5
43	kci-path-mgmt	5
44	complete-kpi	5
45	complete-kci	5
46	kpi-bearer-group	5
47	kpi-ref-path-group	5
48	kpi-kci-bearer-mgmt	5
49	kpi-kci-path-mgmt	5
50	kpi-kci-system	5
51	complete-kpi-kci	5
52	aa-performance	15
53	complete-ethernet-port	15
54	extended-service-ingress-egress	5
55	complete-network-ing-egr	15

to

Syntax	to file <i>file-id</i>
Context	config>log>accounting-policy
Description	This command specifies the destination for the accounting records selected for the accounting policy.
Default	No destination is specified.
Parameters	<p><i>file-id</i> — The <i>file-id</i> option specifies the destination for the accounting records selected for this destination. The characteristics of the file-id must have already been defined in the config>log>file context. A file-id can only be used once.</p> <p>The file is generated when the file policy is referenced. This command identifies the type of accounting file to be created. The file definition defines its characteristics.</p> <p>If the to command is executed while the accounting policy is in operation, then it becomes active during the next collection interval.</p>
Values	1 to 99

5.12.2.11.1 Accounting Policy Custom Record Commands

collection-interval

Syntax	collection-interval <i>minutes</i> no collection-interval
Context	config>log>acct-policy
Description	This command configures the accounting collection interval. The no form of the command returns the value to the default.
Default	60
Parameters	<i>minutes</i> — Specifies the collection interval in minutes. Values 5 to 120

custom-record

Syntax	[no] custom-record
Context	config>log>acct-policy
Description	This command enables the context to configure the layout and setting for a custom accounting record associated with this accounting policy. The no form of the command reverts the configured values to the defaults.

aa-specific

Syntax	[no] aa-specific
Context	config>log>acct-policy>cr
Description	This command enables the context to configure information for this custom record. The no form of the command

aa-sub-counters

Syntax	aa-sub-counters [all] no aa-sub-counters
Context	config>log>acct-policy>cr>aa

Description This command enables the context to configure subscriber counter information. This command only applies to the 7750 SR.

The **no** form of the command

Parameters **all** — Specifies all counters.

long-duration-flow-count

Syntax **long-duration-flow-count**

Context config>log>acct-policy>cr>aa>aa-sub-cntr

Description This command includes the long duration flow count. This command only applies to the 7750 SR.

The **no** form of the command excludes the long duration flow count in the AA subscriber's custom record.

Default no long-duration-flow-count

medium-duration-flow-count

Syntax [**no**] **medium-duration-flow-count**

Context config>log>acct-policy>cr>aa>aa-sub-cntr

Description This command includes the medium duration flow count in the AA subscriber's custom record. This command only applies to the 7750 SR.

The **no** form of the command excludes the medium duration flow count.

Default no medium-duration-flow-count

short-duration-flow-count

Syntax [**no**] **short-duration-flow-count**

Context config>log>acct-policy>cr>aa>aa-sub-cntr

Description This command includes the short duration flow count in the AA subscriber's custom record. This command only applies to the 7750 SR.

The **no** form of the command excludes the short duration flow count.

Default no short-duration-flow-count

total-flow-duration

- Syntax** [no] total-flow-duration
- Context** config>log>acct-policy>cr>aa>aa-sub-cntr
- Description** This command includes the total flow duration flow count in the AA subscriber's custom record. This command only applies to the 7750 SR.
- The **no** form of the command excludes the total flow duration flow count.

total-flows-completed-count

- Syntax** [no] total-flows-completed-count
- Context** config>log>acct-policy>cr>aa>aa-sub-cntr
- Description** This command includes the total flows completed count in the AA subscriber's custom record. This command only applies to the 7750 SR.
- The **no** form of the command excludes the total flow duration flow count.

from-aa-sub-counters

- Syntax** [no] from-aa-sub-counters
- Context** config>log>acct-policy>cr>aa
- Description** This command enables the context to configure Application Assurance "from subscriber" counter parameters. This command only applies to the 7750 SR.
- The **no** form of the command excludes the "from subscriber" count.

all

- Syntax** all
- Context** config>log>acct-policy>cr>aa>aa-from-sub-cntr
config>log>acct-policy>cr>aa>aa-to-sub-cntr
- Description** This command include all counters and only applies to the 7750 SR.

flows-active-count

- Syntax** [no] flows-active-count

Context	config>log>acct-policy>cr>aa>aa-from-sub-cntr config>log>acct-policy>cr>aa>aa-to-sub-cntr
Description	This command includes the active flow count and only applies to the 7750 SR. The no form of the command excludes the active flow count in the AA subscriber's custom record.
Default	no flows-active-count

flows-admitted-count

Syntax	[no] flows-admitted-count
Context	config>log>acct-policy>cr>aa>aa-from-sub-cntr config>log>acct-policy>cr>aa>aa-to-sub-cntr
Description	This command includes the admitted flow count and only applies to the 7750 SR. The no form of the command excludes the flow's admitted count in the AA subscriber's custom record.
Default	no flows-admitted-count

flows-denied-count

Syntax	[no] flows-denied-count
Context	config>log>acct-policy>cr>aa>aa-from-sub-cntr config>log>acct-policy>cr>aa>aa-to-sub-cntr
Description	This command includes the flow's denied count in the AA subscriber's custom record and only applies to the 7750 SR. The no form of the command excludes the flow's denied count.
Default	no flows-denied-count

forwarding-class

Syntax	[no] forwarding-class
Context	config>log>acct-policy>cr>aa>aa-from-sub-cntr config>log>acct-policy>cr>aa>aa-to-sub-cntr
Description	This command enables the collection of a Forwarding Class bitmap information added to the XML aa-sub and router level accounting records, and only applies to the 7750 SR.

Default no forwarding-class

max-throughput-octet-count

Syntax [no] max-throughput-octet-count

Context config>log>acct-policy>cr>aa>aa-from-sub-cntr
config>log>acct-policy>cr>aa>aa-to-sub-cntr

Description This command includes the maximum throughput as measured in the octet count. This command only applies to the 7750 SR.

The **no** form of the command excludes the maximum throughput octet count.

max-throughput-packet-count

Syntax [no] max-throughput-packet-count

Context config>log>acct-policy>cr>aa>aa-from-sub-cntr
config>log>acct-policy>cr>aa>aa-to-sub-cntr

Description This command includes the maximum throughput as measured in the packet count. This command only applies to the 7750 SR.

The **no** form of the command excludes the maximum throughput packet count.

max-throughput-timestamp

Syntax [no] max-throughput-timestamp

Context config>log>acct-policy>cr>aa>aa-from-sub-cntr
config>log>acct-policy>cr>aa>aa-to-sub-cntr

Description This command includes the timestamp of the maximum throughput. This command only applies to the 7750 SR.

The **no** form of the command excludes the timestamp.

octets-admitted-count

Syntax [no] octets-admitted-count

Context config>log>acct-policy>cr>aa>aa-from-sub-cntr
config>log>acct-policy>cr>aa>aa-to-sub-cntr

Description This command includes the admitted octet count in the AA subscriber's custom record and only applies to the 7750 SR.

The **no** form of the command excludes the admitted octet count.

Default no octets-admitted-count

octets-denied-count

Syntax [no] octets-denied-count

Context config>log>acct-policy>cr>aa>aa-from-sub-cntr
config>log>acct-policy>cr>aa>aa-to-sub-cntr

Description This command includes the denied octet count in the AA subscriber's custom record and only applies to the 7750 SR.

The **no** form of the command excludes the denied octet count.

Default no octets-denied-count

packets-admitted-count

Syntax [no] packets-admitted-count

Context config>log>acct-policy>cr>aa>aa-from-sub-cntr
config>log>acct-policy>cr>aa>aa-to-sub-cntr

Description This command includes the admitted packet count in the AA subscriber's custom record and only applies to the 7750 SR.

The **no** form of the command excludes the admitted packet count.

Default no packets-admitted-count

packets-denied-count

Syntax [no] packets-denied-count

Context config>log>acct-policy>cr>aa>aa-from-sub-cntr
config>log>acct-policy>cr>aa>aa-to-sub-cntr

Description This command includes the denied packet count in the AA subscriber's custom record and only applies to the 7750 SR.

The **no** form of the command excludes the denied packet count.

Default no packets-denied-count

to-aa-sub-counters

- Syntax** **to-aa-sub-counters**
no to-aa-sub-counters
- Context** config>log>acct-policy>cr>aa
- Description** This command enables the context to configure Application Assurance “to subscriber” counter parameters and only applies to the 7750 SR.
- The **no** form of the command excludes the “to subscriber” count.

override-counter

- Syntax** [**no**] **override-counter** *override-counter-id*
- Context** config>log>acct-policy>cr
- Description** This command enables the context to configure override counter (HSMDA) parameters. This command only applies to the 7750 SR.
- The **no** form of the command removes the ID from the configuration.
- Parameters** *override-counter-id* — Specifies the override counter ID.
- Values** 1 to 8

queue

- Syntax** [**no**] **queue** *queue-id*
- Context** config>log>acct-policy>cr
- Description** This command specifies the queue-id for which counters will be collected in this custom record. The counters that will be collected are defined in egress and ingress counters.
- The **no** form of the command reverts to the default value.
- Parameters** *queue-id* — Specifies the queue-id for which counters will be collected in this custom record.

e-counters

- Syntax** [**no**] **e-counters**
- Context** config>log>acct-policy>cr>override-cntr
config>log>acct-policy>cr>queue


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

Description This command configures egress counter parameters for this custom record.
The **no** form of the command reverts to the default value.

i-counters

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

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

Description This command configures ingress counter parameters for this custom record.
The **no** form of the command

Parameters **all** — Specifies all ingress counters should be included.

in-profile-octets-discarded-count

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

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

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

in-profile-octets-forwarded-count

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

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

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

in-profile-packets-discarded-count

- Syntax** [no] in-profile-packets-discarded-count
- Context** config>log>acct-policy>cr>oc>e-count
config>log>acct-policy>cr>roc>e-count
config>log>acct-policy>cr>queue>e-count
config>log>acct-policy>cr>ref-queue>e-count
- Description** This command includes the in-profile packets discarded count.
The **no** form of the command excludes the in-profile packets discarded count.

in-profile-packets-forwarded-count

- Syntax** [no] in-profile-packets-forwarded-count
- Context** config>log>acct-policy>cr>oc>e-count
config>log>acct-policy>cr>roc>e-count
config>log>acct-policy>cr>queue>e-count
config>log>acct-policy>cr>ref-queue>e-count
- Description** This command includes the in-profile packets forwarded count.
The **no** form of the command excludes the in-profile packets forwarded count.

out-profile-octets-discarded-count

- Syntax** [no] out-profile-octets-discarded-count
- Context** config>log>acct-policy>cr>oc>e-count
config>log>acct-policy>cr>roc>e-count
config>log>acct-policy>cr>queue>e-count
config>log>acct-policy>cr>ref-queue>e-count
- Description** This command includes the out of profile packets discarded count.
The **no** form of the command excludes the out of profile packets discarded count.

out-profile-octets-forwarded-count

- Syntax** [no] out-profile-octets-forwarded-count
- Context** config>log>acct-policy>cr>oc>e-count
config>log>acct-policy>cr>roc>e-count
config>log>acct-policy>cr>queue>e-count
config>log>acct-policy>cr>ref-queue>e-count

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

out-profile-packets-discarded-count

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

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

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

out-profile-packets-forwarded-count

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

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

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

all-octets-offered-count

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

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

Description This command includes all octets offered in the count.
The **no** form of the command excludes the octets offered in the count.

Default no all-octets-offered-count

all-packets-offered-count

Syntax	[no] all-packets-offered-count
Context	config>log>acct-policy>cr>oc>i-count config>log>acct-policy>cr>roc>i-count config>log>acct-policy>cr>queue>i-count config>log>acct-policy>cr>ref-queue>i-count
Description	This command includes all packets offered in the count. The no form of the command excludes the packets offered in the count.
Default	no all-packets-offered-count

high-octets-discarded-count

Syntax	[no] high-octets-discarded-count
Context	config>log>acct-policy>cr>oc>i-count config>log>acct-policy>cr>roc>i-count config>log>acct-policy>cr>queue>i-count config>log>acct-policy>cr>ref-queue>i-count
Description	This command includes the high octets discarded count. The no form of the command excludes the high octets discarded count.
Default	no high-octets-discarded-count

high-octets-offered-count

Syntax	[no] high-octets-offered-count
Context	config>log>acct-policy>cr>oc>i-count config>log>acct-policy>cr>roc>i-count config>log>acct-policy>cr>queue>i-count config>log>acct-policy>cr>ref-queue>i-count
Description	This command includes the high octets offered count. The no form of the command excludes the high octets offered count.

high-packets-discarded-count

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

Context	config>log>acct-policy>cr>oc>i-count config>log>acct-policy>cr>roc>i-count config>log>acct-policy>cr>queue>i-count config>log>acct-policy>cr>ref-queue>i-count
Description	This command includes the high packets discarded count. The no form of the command excludes the high packets discarded count.
Default	no high-packets-discarded-count

high-packets-offered-count

Syntax	[no] high-packets-offered-count
Context	config>log>acct-policy>cr>oc>i-count config>log>acct-policy>cr>roc>i-count config>log>acct-policy>cr>queue>i-count config>log>acct-policy>cr>ref-queue>i-count
Description	This command includes the high packets offered count. The no form of the command excludes the high packets offered count.
Default	no high-packets-offered -count

in-profile-octets-forwarded-count

Syntax	[no] in-profile-octets-forwarded-count
Context	config>log>acct-policy>cr>oc>i-count config>log>acct-policy>cr>roc>i-count config>log>acct-policy>cr>queue>i-count config>log>acct-policy>cr>ref-queue>i-count
Description	This command includes the in profile octets forwarded count. The no form of the command excludes the in profile octets forwarded count.
Default	no in-profile-octets-forwarded-count

in-profile-packets-forwarded-count

Syntax	[no] in-profile-packets-forwarded-count
Context	config>log>acct-policy>cr>oc>i-count config>log>acct-policy>cr>roc>i-count

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

- Description** This command includes the in profile packets forwarded count.
- The **no** form of the command excludes the in profile packets forwarded count.
- Default** no in-profile-packets-forwarded-count

low-octets-discarded-count

- Syntax** [no] low-octets-discarded-count
- Context** config>log>acct-policy>cr>oc>i-count
config>log>acct-policy>cr>roc>i-count
config>log>acct-policy>cr>queue>i-count
config>log>acct-policy>cr>ref-queue>i-count
- Description** This command includes the low octets discarded count.
- The **no** form of the command excludes the low octets discarded count.
- Default** no low-octets-discarded-count

low-packets-discarded-count

- Syntax** [no] low-packets-discarded-count
- Context** config>log>acct-policy>cr>oc>i-count
config>log>acct-policy>cr>roc>i-count
config>log>acct-policy>cr>queue>i-count
config>log>acct-policy>cr>ref-queue>i-count
- Description** This command includes the low packets discarded count.
- The **no** form of the command excludes the low packets discarded count.
- Default** no low-packets-discarded-count

low-octets-offered-count

- Syntax** [no] low-octets-offered-count
- Context** config>log>acct-policy>cr>oc>i-count
config>log>acct-policy>cr>roc>i-count
config>log>acct-policy>cr>queue>i-count
config>log>acct-policy>cr>ref-queue>i-count

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

low-packets-offered-count

Syntax [no] low-packets-offered-count

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

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

out-profile-octets-forwarded-count

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

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

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

Default no out-profile-octets-forwarded-count

out-profile-packets-forwarded-count

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

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

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

Default no out-profile-packets-forwarded-count

uncoloured-octets-offered-count

- Syntax** `[no] uncoloured-packets-offered-count`
- Context** `config>log>acct-policy>cr>queue>i-count`
`config>log>acct-policy>cr>ref-queue>i-count`
- Description** This command includes the uncoloured octets offered in the count.
The **no** form of the command excludes the uncoloured octets offered in the count.

uncoloured-packets-offered-count

- Syntax** `[no] uncoloured-packets-offered-count`
- Context** `config>log>acct-policy>cr>queue>i-count`
`config>log>acct-policy>cr>ref-queue>i-count`
- Description** This command includes the uncolored packets offered count.
The **no** form of the command excludes the uncoloured packets offered count.

ref-aa-specific-counter

- Syntax** `ref-aa-specific-counter any`
`no ref-aa-specific-counter`
- Context** `config>log>acct-policy>cr`
- Description** This command enables the use of significant-change so only those aa-specific records which have changed in the last accounting interval are written.
The **no** form of the command disables the use of significant-change so all aa-specific records are written whether or not they have changed within the last accounting interval.
- Parameters** **any** — Indicates that a record is collected as long as any field records activity when non-zero significant-change value is configured.

ref-override-counter

- Syntax** `ref-override-counter ref-override-counter-id`
`ref-override-counter all`
`no ref-override-counter`
- Context** `config>log>acct-policy>cr`
- Description** This command configures a reference override counter.

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

Default no ref-override-counter

ref-queue

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

Context config>log>acct-policy>cr

Description This command configures a reference queue.

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

Default no ref-queue

significant-change

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

Context config>log>acct-policy>cr

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

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

Values 0 to 4294967295 (For custom-record-aa-sub only values 0 or 1 are supported.)

5.13 Log Command Reference

5.13.1 Command Hierarchies

- [Show Commands](#)
- [Clear Command](#)

5.13.1.1 Show Commands

Refer to the SR OS Services Guide for information about log show routines for VPRN services.

```
show
  — log
    — accounting-policy [acct-policy-id] [access | network]
    — accounting-records
    — applications
    — event-control [application-id] [event-name | event-number]
    — event-handling
      — handler [handler-name]
      — handler detail
      — information
      — scripts
    — event-parameters [application-id] [event-name | event-number]
    — file-id [log-file-id]
    — filter-id [filter-id]
    — log-collector
    — log-id [log-id] [severity severity-level] [application application] [sequence from-seq
      [to-seq]] [count count] [router router-instance] [expression] [subject subject
      [regex]] [ascending | descending] [message format] [msg-regex]]
    — snmp-trap-group [log-id]
    — syslog [syslog-id]
```

5.13.1.2 Clear Command

```
clear
  — log log-id
  — log
    — log-id log-id
    — event-handling
      — handler event-handler-name
      — information
```

5.13.2 Command Descriptions

- [Show Commands](#)
- [Clear Commands](#)

5.13.2.1 Show Commands

The command output in the following section are examples only; actual displays may differ depending on supported functionality and user configuration.

accounting-policy

- Syntax** `accounting-policy [acct-policy-id] [access | network]`
- Context** `show>log`
- Description** This command displays accounting policy information.
- Parameters** *policy-id* — The policy ID that uniquely identifies the accounting policy, expressed as a decimal integer.
- Values** 1 to 99
- access** — Only displays access accounting policies.
- network** — Only displays network accounting policies.
- Output** Accounting Policy Output
- [Table 77](#) describes accounting policy output fields.

Table 77 Show Accounting Policy Output Fields

Label	Description
Policy ID	The identifying value assigned to a specific policy.

Table 77 Show Accounting Policy Output Fields (Continued)

Label	Description
Type	Identifies accounting record type forwarded to the configured accounting file.
	access Indicates that the policy is an access accounting policy.
	network Indicates that the policy is a network accounting policy.
	none Indicates no accounting record types assigned.
Def	Yes Indicates that the policy is a default access or network policy.
	No Indicates that the policy is not a default access or network policy.
Admin State	Displays the administrative state of the policy.
	Up Indicates that the policy is administratively enabled.
	Down Indicates that the policy is administratively disabled.
Oper State	Displays the operational state of the policy.
	Up Indicates that the policy is operationally up.
	Down Indicates that the policy is operationally down.
Intvl	Displays the interval, in minutes, in which statistics are collected and written to their destination. The default depends on the record name type.
File ID	The log destination.
Record Name	The accounting record name which represents the configured record type.
This policy is applied to	Specifies the entity where the accounting policy is applied.

Sample Output

```
A:ALA-1# show log accounting-policy
=====
Accounting Policies
=====
Policy Type   Def Admin Oper  Intvl   File Record Name
Id           State State
-----
1      network No  Up    Up    15      1  network-ingress-packets
2      network Yes Up    Up    15      2  network-ingress-octets
10     access  Yes Up    Up     5      3  complete-service-ingress-egress
=====
A:ALA-1#
```

```
A:ALA-1# show log accounting-policy 10
=====
Accounting Policies
=====
Policy Type   Def Admin Oper  Intvl   File Record Name
Id           State State
-----
10     access  Yes Up    Up     5      3  complete-service-ingress-egress
```

Description : (Not Specified)

```
This policy is applied to:
  Svc Id: 100  SAP : 1/1/8:0  Collect-Stats
  Svc Id: 101  SAP : 1/1/8:1  Collect-Stats
  Svc Id: 102  SAP : 1/1/8:2  Collect-Stats
  Svc Id: 103  SAP : 1/1/8:3  Collect-Stats
  Svc Id: 104  SAP : 1/1/8:4  Collect-Stats
  Svc Id: 105  SAP : 1/1/8:5  Collect-Stats
  Svc Id: 106  SAP : 1/1/8:6  Collect-Stats
  Svc Id: 107  SAP : 1/1/8:7  Collect-Stats
  Svc Id: 108  SAP : 1/1/8:8  Collect-Stats
  Svc Id: 109  SAP : 1/1/8:9  Collect-Stats
...
=====
```

```
A:ALA-1#
A:ALA-1# show log accounting-policy network
=====
Accounting Policies
=====
Policy Type   Def Admin Oper  Intvl   File Record Name
Id           State State
-----
1      network No  Up    Up    15      1  network-ingress-packets
2      network Yes Up    Up    15      2  network-ingress-octets
=====
A:ALA-1#
```

```
A:ALA-1# show log accounting-policy access
=====
Accounting Policies
=====
Policy Type   Def Admin Oper  Intvl   File Record Name
Id           State State
-----
```

```
-----
10    access  Yes Up    Up    5          3 complete-service-ingress
-----
A:ALA-1#
```

accounting-records

- Syntax** accounting-records
- Context** show>log
- Description** This command displays accounting policy record names.
- Output** Accounting Records Output

[Table 78](#) describes accounting records output fields.

Table 78 Accounting Policy Output Fields

Label	Description
Record #	The record ID that uniquely identifies the accounting policy, expressed as a decimal integer.
Record Name	The accounting record name.
Def. Interval	The default interval, in minutes, in which statistics are collected and written to their destination.

Sample Output



Note: aa, video and subscriber records are not applicable to the 7950 XRS.

```
A:ALA-1# show log accounting-records
=====
Accounting Policy Records
=====
Record # Record Name                               Def. Interval
-----
1      service-ingress-octets                          5
2      service-egress-octets                          5
3      service-ingress-packets                        5
4      service-egress-packets                         5
5      network-ingress-octets                         15
6      network-egress-octets                         15
7      network-ingress-packets                       15
8      network-egress-packets                        15
9      compact-service-ingress-octets                5
```

10	combined-service-ingress	5
11	combined-network-ing-egr-octets	15
12	combined-service-ing-egr-octets	5
13	complete-service-ingress-egress	5
14	combined-sdp-ingress-egress	5
15	complete-sdp-ingress-egress	5
16	complete-subscriber-ingress-egress	5
17	aa-protocol	15
18	aa-application	15
19	aa-app-group	15
20	aa-subscriber-protocol	15
21	aa-subscriber-application	15
22	aa-subscriber-app-group	15

=====
A:ALA-1#

applications

Syntax applications

Context show>log

Description This command displays a list of all application names that can be used in event-control and filter commands.

Output

Sample Output

```
*A:7950 XRS-20# show log applications
```

```
=====  
Log Event Application Names  
=====  
Application Name  
-----  
BGP  
...  
CHASSIS  
...  
IGMP  
...  
LDP  
LI  
...  
MIRROR  
...  
MPLS  
...  
OSPF  
PIM  
...  
PORT  
...
```

```

SYSTEM
...
USER
...
VRTR
...
=====
A:ALA-1#
    
```

event-control

- Syntax** `event-control [application [event-name | event-number]]`
- Context** `show>log`
- Description** This command displays event control settings for events including whether the event is suppressed or generated and the severity level for the event.

If no options are specified all events, alarms and traps are listed.
- Parameters**
 - application-id** — Only displays event control for the specified application.
 - Default** All applications.
 - The following are some sample applications:
 - Values** bgp, cflowd, chassis, debug, igmp, lldp, mirror, ospf, pim, port, snmp, system, user, vrtr
 - event-name** — Only displays event control for the named application event.
 - Default** All events for the application.
 - event-number** — Only displays event control for the specified application event number.
 - Default** All events for the application.
- Output** Show Event Control Output

[Table 79](#) describes the output fields for the event control.

Table 79 Event-Control Output Field Descriptions

Label	Description
Application	The application name.
ID#	The event ID number within the application. L ID# An “L” in front of an ID represents event types that do not generate an associated SNMP notification. Most events do generate a notification, only the exceptions are marked with a preceding “L”.

Table 79 Event-Control Output Field Descriptions (Continued)

Label	Description (Continued)
Event Name	The event name.
P	CL The event has a cleared severity/priority.
	CR The event has critical severity/priority.
	IN The event has indeterminate severity/priority.
	MA The event has major severity/priority.
	MI The event has minor severity/priority.
	WA The event has warning severity/priority.
g/s	gen The event will be generated/logged by event control.
	sup The event will be suppressed/dropped by event control.
	thr Specifies that throttling is enabled.
Logged	The number of events logged/generated.
Dropped	The number of events dropped/suppressed.

Sample Output

The following is a sample output:

```
A:gal171# show log event-control
=====
Log Events
=====
Application
ID#      Event Name                P  g/s  Logged  Dropped
-----
BGP:
 2001  bgpEstablished             MI  gen    0       0
 2002  bgpBackwardTransition     WA  gen    0       0
 2003  tBgpMaxPrefix90           WA  gen    0       0
 2004  tBgpMaxPrefix100         CR  gen    0       0
```

L	2005	sendNotification	WA	gen	0	0
L	2006	receiveNotification	WA	gen	0	0
L	2007	bgpInterfaceDown	WA	gen	0	0
L	2008	bgpConnNoKA	WA	gen	0	0
L	2009	bgpConnNoOpenRcvd	WA	gen	0	0
L	2010	bgpRejectConnBadLocAddr	WA	gen	0	0
L	2011	bgpRemoteEndClosedConn	WA	gen	0	0
L	2012	bgpPeerNotFound	WA	gen	0	0
L	2013	bgpConnMgrTerminated	WA	gen	0	0
L	2014	bgpTerminated	WA	gen	0	0
L	2015	bgpNoMemoryPeer	CR	gen	0	0
L	2016	bgpVariableRangeViolation	WA	gen	0	0
L	2017	bgpCfgViol	WA	gen	0	0
CFLOWD:						
	2001	cflowdCreated	MI	gen	0	0
	2002	cflowdCreateFailure	MA	gen	0	0
	2003	cflowdDeleted	MI	gen	0	0
	2004	cflowdStateChanged	MI	gen	0	0
	2005	cflowdCleared	MI	gen	0	0
	2006	cflowdFlowCreateFailure	MI	gen	0	0
	2007	cflowdFlowFlushFailure	MI	gen	0	0
	2008	cflowdFlowUnsuppProto	MI	sup	0	0
CCAG:						
CHASSIS:						
	2001	cardFailure	MA	gen	0	0
	2002	cardInserted	MI	gen	4	0
	2003	cardRemoved	MI	gen	0	0
	2004	cardWrong	MI	gen	0	0
	2005	EnvTemperatureTooHigh	MA	gen	0	0
...						
DEBUG:						
L	2001	traceEvent	MI	gen	0	0
DOT1X:						
FILTER:						
	2001	filterPBRPacketsDropped	MI	gen	0	0
IGMP:						
	2001	vRtrIgmpIfRxQueryVerMismatch	WA	gen	0	0
	2002	vRtrIgmpIfCModeRxQueryMismatch	WA	gen	0	0
IGMP_SNOOPING:						
IP:						
L	2001	clearRTMError	MI	gen	0	0
L	2002	ipEtherBroadcast	MI	gen	0	0
L	2003	ipDuplicateAddress	MI	gen	0	0
L	2004	ipArpInfoOverwritten	MI	gen	0	0
L	2005	fibAddFailed	MA	gen	0	0
L	2006	qosNetworkPolicyMallocFailed	MA	gen	0	0
L	2007	ipArpBadInterface	MI	gen	0	0
L	2008	ipArpDuplicateIpAddress	MI	gen	0	0
L	2009	ipArpDuplicateMacAddress	MI	gen	0	0
ISIS:						
	2001	vRtrIsisDatabaseOverload	WA	gen	0	0
	2002	vRtrIsisManualAddressDrops	WA	gen	0	0
	2003	vRtrIsisCorruptedLSPDetected	WA	gen	0	0
	2004	vRtrIsisMaxSeqExceedAttempt	WA	gen	0	0
	2005	vRtrIsisIDLLenMismatch	WA	gen	0	0
	2006	vRtrIsisMaxAreaAdrrsMismatch	WA	gen	0	0
....						
USER:						

L	2001	cli_user_login	MI	gen	2	0
L	2002	cli_user_logout	MI	gen	1	0
L	2003	cli_user_login_failed	MI	gen	0	0
L	2004	cli_user_login_max_attempts	MI	gen	0	0
L	2005	ftp_user_login	MI	gen	0	0
L	2006	ftp_user_logout	MI	gen	0	0
L	2007	ftp_user_login_failed	MI	gen	0	0
L	2008	ftp_user_login_max_attempts	MI	gen	0	0
L	2009	cli_user_io	MI	sup	0	48
L	2010	snmp_user_set	MI	sup	0	0
L	2011	cli_config_io	MI	gen	4357	0
VRRP:						
	2001	vrpTrapNewMaster	MI	gen	0	0
	2002	vrpTrapAuthFailure	MI	gen	0	0
	2003	tmnxVrrpIPListMismatch	MI	gen	0	0
	2004	tmnxVrrpIPListMismatchClear	MI	gen	0	0
	2005	tmnxVrrpMultipleOwners	MI	gen	0	0
	2006	tmnxVrrpBecameBackup	MI	gen	0	0
L	2007	vrpPacketDiscarded	MI	gen	0	0
VRTR:						
	2001	tmnxVRtrMidRouteTCA	MI	gen	0	0
	2002	tmnxVRtrHighRouteTCA	MI	gen	0	0
	2003	tmnxVRtrHighRouteCleared	MI	gen	0	0
	2004	tmnxVRtrIllegalLabelTCA	MA	gen	0	0
	2005	tmnxVRtrMcastMidRouteTCA	MI	gen	0	0
	2006	tmnxVRtrMcastMaxRoutesTCA	MI	gen	0	0
	2007	tmnxVRtrMcastMaxRoutesCleared	MI	gen	0	0
	2008	tmnxVRtrMaxArpEntriesTCA	MA	gen	0	0
	2009	tmnxVRtrMaxArpEntriesCleared	MI	gen	0	0
	2011	tmnxVRtrMaxRoutes	MI	gen	0	0

=====
A:ALA-1#

A:ALA-1# show log event-control ospf

=====
Log Events

=====
Application

ID#	Event Name	P	g/s	Logged	Dropped
2001	ospfVirtIfStateChange	WA	gen	0	0
2002	ospfNbrStateChange	WA	gen	1	0
2003	ospfVirtNbrStateChange	WA	gen	0	0
2004	ospfIfConfigError	WA	gen	0	0
2005	ospfVirtIfConfigError	WA	gen	0	0
2006	ospfIfAuthFailure	WA	gen	0	0
2007	ospfVirtIfAuthFailure	WA	gen	0	0
2008	ospfIfRxBadPacket	WA	gen	0	0
2009	ospfVirtIfRxBadPacket	WA	gen	0	0
2010	ospfTxRetransmit	WA	sup	0	0
2011	ospfVirtIfTxRetransmit	WA	sup	0	0
2012	ospfOriginateLsa	WA	sup	0	404
2013	ospfMaxAgeLsa	WA	gen	3	0
2014	ospfLsdbOverflow	WA	gen	0	0
2015	ospfLsdbApproachingOverflow	WA	gen	0	0
2016	ospfIfStateChange	WA	gen	2	0
2017	ospfNssaTranslatorStatusChange	WA	gen	0	0
2018	vRtrOspfSpfRunsStopped	WA	gen	0	0
2019	vRtrOspfSpfRunsRestarted	WA	gen	0	0

```

2020 vRtrOspfOverloadEntered      WA gen      1      0
2021 vRtrOspfOverloadExited      WA gen      0      0
2022 ospfRestartStatusChange     WA gen      0      0
2023 ospfNbrRestartHelperStatusChange WA gen      0      0
2024 ospfVirtNbrRestartHelperStsChg WA gen      0      0
=====
A:ALA-1#

A:ALA-1# show log event-control ospf ospfVirtIfStateChange
=====
Log Events
=====
Application
ID#      Event Name                P   g/s      Logged      Dropped
-----
2001 ospfVirtIfStateChange     WA gen      0      0
=====
A:ALA-1#

```

event-handling

- Syntax** `event-handling`
- Context** `show>log`
- Description** This command enables the context to display Event Handling System (EHS) information.

handler

- Syntax** `handler [handler-name]`
`handler detail`
- Context** `show>log>event-handling`
- Description** This command enters the context to display EHS handler information.
- Parameters** *handler-name* — Specifies the name of a specific handler. 32 characters maximum.
detail — Keyword to list details of all handlers.
- Output** Show Handler Output

[Table 80](#) describes handler output fields.

Table 80 Handler Output Field Descriptions

Label	Description
Handler	The name of the handler.
Description	The handler description string.

Table 80 Handler Output Field Descriptions (Continued)

Label	Description (Continued)
Admin State	The administrative state of the handler.
Oper State	The operational state of the handler.
Handler Action-List Entry	
Entry-id	The action-list entry identifier.
Description	The action-list entry description string.
Admin State	The administrative state of the action-list entry.
Policy Name	The name of the related script policy.
Policy Owner	The owner of the related script policy.
Last Exec	The timestamp of the last successful execution of the action-list entry.
Handler Action-List Entry Execution Statistics	
Enqueued	The number of times the action-list entry was successfully passed on to the SR OS sub-system or module that will attempt to process and execute the action. For a script-policy entry, this indicates that the script request has been enqueued but does not necessarily indicate that the script has successfully launched or completed. For status and information about the script, use the show>system>script-control command.
Err Launch	The number of times the action-list entry was not successfully handed over to the next SR OS sub-system or module in the processing chain. This can be caused by a variety of conditions including a full script request input queue.
Err Adm Status	The number of times the action-list entry was not executed because the entry was administratively disabled.
Total	The total number of times that the action-list entry attempted execution.

Sample Output

```
A:node1>show>log>event-handling# handler
=====
Event Handling System - Handler List
=====
Handler      Admin  Oper  Description
Name         State State
-----
h-sample           up    up
```

```

h-main                up      up
h-backup              down    down
=====

*A:7950 XRS-20# show log event-handling handler "h-sample"

=====
Event Handling System - Handlers
=====

Handler               : h-sample
=====
Description            : (Not Specified)
Admin State            : up                      Oper State : up

-----
Handler Action-List Entry
-----
Entry-id               : 10
Description             : (Not Specified)
Admin State            : up                      Oper State : up
Script
  Policy Name          : sp-sample
  Policy Owner         : TiMOS CLI
Min Delay              : 0
Last Exec              : 05/24/2015 19:03:31
-----
Handler Action-List Entry Execution Statistics
  Enqueued              : 4
  Err Launch           : 0
  Err Adm Status       : 0
Total                  : 4
=====

```

information

- Syntax** **information**
- Context** show>log>event-handling
- Description** This command displays general information about EHS, as well as handler and trigger statistics.
- Output** Show Information Output

Sample output

```

=====
Event Handling System - Event Trigger Statistics
=====

Application Name
Event Id                Total      Success  ErrNoEntry  AdmStatus
-----

```



```

OAM
2001
-----
Entry FilMatch  Trigger  Debounce  FilFail  ErrAdmSta  ErrFilter  ErrHandler
-----
1      0         0         0         0         0         0         0
10     0         0         0         0         0         0         0
-----
SUM    0         0         0         0         0         0         0
-----
Application Name
Event Id          Total      Success    ErrNoEntry  AdmStatus
-----
OAM
2004
-----
Entry FilMatch  Trigger  Debounce  FilFail  ErrAdmSta  ErrFilter  ErrHandler
-----
1      0         0         0         0         0         0         0
-----
SUM    0         0         0         0         0         0         0
=====
EVENTS PROCESSED          Total      Success    ErrNoEntry  AdmStatus
-----
                                0         0         0         0
=====
Event Handling System - Event Handler Statistics
=====
Handler          Total      Success    ErrNoEntry  AdmStatus
my-handler-1
-----
Entry Id        Launch    MinDelay  ErrLaunch  ErrAdmSta
-----
1                0         0         0         0
-----
SUMMARY         0         0         0         0
=====
HANDLERS SUMMARY          Total      Success    ErrNoEntry  AdmStatus
-----
                                0         0         0         0
=====

```

scripts

- Syntax** **scripts**
- Context** show>log>event-handling
- Description** This command displays handler configuration and script run queue information.
- Output** Show Scripts Output

Sample output

```

=====
Event Handling System - Script Policy Association
=====
-----
No Matching Entries Found
=====
Event Handling System - Script Association
=====
-----
No Matching Entries Found
=====
Event Handling System - Script Launched List
=====
Run #      Script owner      Script name      Script state
-----
No Matching Entries
=====

```

event-parameters

- Syntax** **event-parameters** [*application-id* [*event-name* | *event-number*]]
- Context** show>log
- Description** This command displays an event's (or all events) common parameters and specific parameters. This allows a user to know what parameters can be passed from a triggering event to the triggered EHS script.
- Parameters** **application-id** — Only displays event parameters for the specified application.
 - Default** All applications.
 - The following are some sample applications:
 - Values** bgp, cflowd, chassis, debug, igmp, lldp, mirror, ospf, pim, port, snmp, system, user, vrtr
- event-name** — Only displays event parameters for the named application event.
 - Default** All events for the application.
- event-number** — Only displays event parameters for the specified application event number.
 - Default** All events for the application.
- Output** show event-parameters output

Sample output

```

# show log event-parameters "oam" 2001
=====
Common Event Parameters
      appid

```

```

name
eventid
severity
subject
gentime
Event Specific Parameters
tmnxOamPingCtlOwnerIndex
tmnxOamPingCtlTestIndex
tmnxOamPingCtlTgtAddrType
tmnxOamPingCtlTgtAddress
tmnxOamPingResultsTestRunIndex
tmnxOamPingResultsOperStatus
tmnxOamPingResultsMinRtt
tmnxOamPingResultsMaxRtt
tmnxOamPingResultsAverageRtt
tmnxOamPingResultsRttSumOfSquares
tmnxOamPingResultsRttOFSumSquares
tmnxOamPingResultsMtuResponseSize
tmnxOamPingResultsSvcPing
tmnxOamPingResultsProbeResponses
tmnxOamPingResultsSentProbes
tmnxOamPingResultsLastGoodProbe
tmnxOamPingCtlTestMode
tmnxOamPingHistoryIndex
=====

```

file-id

- Syntax** `file-id [log-file-id]`
- Context** `show>log`
- Description** This command displays event file log information.

If no command line parameters are specified, a summary output of all event log files is displayed.

Specifying a file ID displays detailed information on the event file log.
- Parameters** *log-file-id* — Displays detailed information on the specified event file log.
- Output** Log File Output

[Table 81](#) describes the output fields for a log file summary.

Table 81 Log File Summary Output Fields

Label	Description
file-id	The log file ID.
rollover	The rollover time for the log file which is how long in between partitioning of the file into a new file.

Table 81 Log File Summary Output Fields (Continued)

Label	Description (Continued)
retention	The retention time for the file in the system which is how long the file should be retained in the file system.
admin location	The primary flash device specified for the file location.
	none indicates no specific flash device was specified.
backup location	The secondary flash device specified for the file location if the admin location is not available.
	none Indicates that no backup flash device was specified.
oper location	The actual flash device on which the log file exists.
file-id	The log file ID.
rollover	The rollover time for the log file which is how long in between partitioning of the file into a new file.
retention	The retention time for the file in the system which is how long the file should be retained in the file system.
file name	The complete pathname of the file associated with the log ID.
expired	Indicates whether or not the retention period for this file has passed.
state	in progress Indicates the current open log file.
	complete Indicates the old log file.

Sample Output

```
A:ALA-1# show log file-id
=====
File Id List
=====
file-id  rollover  retention  admin    backup    oper
          location  location  location location  location
-----
1         60        4         cf1:    cf2:    cf1:
2         60        3         cf1:    cf3:    cf1:
3        1440     12        cf1:    none    cf1:
10       1440     12        cf1:    none    none
11       1440     12        cf1:    none    none
15       1440     12        cf1:    none    none
```

```

20          1440          12          cf1:          none          none
=====
A:ALA-1#

A:ALA-1# show log file-id 10
=====
File Id List
=====
file-id  rollover  retention  admin    backup    oper
          location  location  location
-----
10 1440      12        cf3:     cf2:     cf1:
Description : Main
=====
File Id 10 Location cf1:
=====
file name                                expired  state
-----
cf1:\log\log0302-20060501-012205        yes     complete
cf1:\log\log0302-20060501-014049        yes     complete
cf1:\log\log0302-20060501-015344        yes     complete
cf1:\log\log0302-20060501-015547        yes     in progress
=====
A:ALA-1#

```

filter-id

- Syntax** **filter-id** [*filter-id*]
- Context** show>log
- Description** This command displays event log filter policy information.
- Parameters** *filter-id* — Displays detailed information on the specified event filter policy ID.
- Output** Event Log Filter Summary Output

[Table 82](#) describes the output fields for event log filter summary information.

Table 82 Event Log Filter Summary Output Fields

Label	Description
Filter Id	The event log filter ID.
Applied	no The event log filter is not currently in use by a log ID.
	yes The event log filter is currently in use by a log ID.

Table 82 Event Log Filter Summary Output Fields (Continued)

Label	Description
Default Action	drop The default action for the event log filter is to drop events not matching filter entries.
	forward The default action for the event log filter is to forward events not matching filter entries.
Description	The description string for the filter ID.

Sample Output

```
*A:ALA-48>config>log# show log filter-id
=====
Log Filters
=====
Filter Applied Default Description
Id           Action
-----
1           no       forward
5           no       forward
10          no       forward
1001        yes      drop      Collect events for Serious Errors Log
=====
*A:ALA-48>config>log#
```

Event Log Filter Detailed Output

[Table 83](#) describes the output fields for detailed event log filter information.

Table 83 Event Log Filter Detail Output Fields

Label	Description
Filter-id	The event log filter ID.
Applied	no The event log filter is not currently in use by a log ID.
	yes The event log filter is currently in use by a log ID.

Table 83 Event Log Filter Detail Output Fields (Continued)

Label	Description
Default Action	drop The default action for the event log filter is to drop events not matching filter entries.
	forward The default action for the event log filter is to forward events not matching filter entries.
Description (Filter-id)	The description string for the filter ID.

Table 84 describes the output fields for log filter match criteria information.

Table 84 Log Filter Match Criteria Output Fields

Label	Description
Entry-id	The event log filter entry ID.
Action	default There is no explicit action for the event log filter entry and the filter's default action is used on matching events.
	drop The action for the event log filter entry is to drop matching events.
	forward The action for the event log filter entry is to forward matching events.
Description (Entry-id)	The description string for the event log filter entry.
Application	The event log filter entry application match criterion.
Event Number	The event log filter entry application event ID match criterion.

Table 84 Log Filter Match Criteria Output Fields (Continued)

Label	Description
Severity	cleared The log event filter entry application event severity cleared match criterion.
	indeterminate The log event filter entry application event severity indeterminate match criterion.
	critical The log event filter entry application event severity critical match criterion.
	major The log event filter entry application event severity cleared match criterion.
	minor The log event filter entry application event severity minor match criterion.
	warning The log event filter entry application event severity warning match criterion.
Subject	Displays the event log filter entry application event ID subject string match criterion.
Router	Displays the event log filter entry application event ID router <i>router-instance</i> string match criterion.

Table 84 Log Filter Match Criteria Output Fields (Continued)

Label	Description
Operator	There is an operator field for each match criteria: application, event number, severity, and subject.
equal	Matches when equal to the match criterion.
greaterThan	Matches when greater than the match criterion.
greaterThanOrEqual	Matches when greater than or equal to the match criterion.
lessThan	Matches when less than the match criterion.
lessThanOrEqual	Matches when less than or equal to the match criterion.
notEqual	Matches when not equal to the match criterion.
off	No operator specified for the match criterion.

Sample Output

```
*A:ALA-48>config>log# show log filter-id 1001
=====
Log Filter
=====
Filter-id      : 1001      Applied      : yes      Default Action: drop
Description    : Collect events for Serious Errors Log
-----
Log Filter Match Criteria
-----
Entry-id      : 10              Action       : forward
Application   :                  Operator     : off
Event Number  : 0              Operator     : off
Severity      : major          Operator     : greaterThanOrEqual
Subject       :                  Operator     : off
Match Type    : exact string      :
Router        :                  Operator     : off
Match Type    : exact string      :
Description   : Collect only events of major severity or higher
-----
*A:ALA-48>config>log#
```

log-collector

- Syntax** **log-collector**
- Context** show>log
- Description** Show log collector statistics for the main, security, change and debug log collectors.
- Output** Log-Collector Output

[Table 85](#) describes log-collector output fields.

Table 85 Show Log-Collector Output Fields

Label	Description
<Collector Name>	Main The main event stream contains the events that are not explicitly directed to any other event stream.
	Security The security stream contains all events that affect attempts to breach system security such as failed login attempts, attempts to access MIB tables to which the user is not granted access or attempts to enter a branch of the CLI to which access has not been granted.
	Change The change event stream contains all events that directly affect the configuration or operation of this node.
	Debug The debug-trace stream contains all messages in the debug stream.
Dest. Log ID	Specifies the event log stream destination.
Filter ID	The value is the index to the entry which defines the filter to be applied to this log's source event stream to limit the events output to this log's destination. If the value is 0, then all events in the source log are forwarded to the destination.
Status	Enabled Logging is enabled.
	Disabled Logging is disabled.

Table 85 Show Log-Collector Output Fields (Continued)

Label	Description
Dest. Type	<p>Console</p> <p>A log created with the console type destination displays events to the physical console device.</p> <p>Events are displayed to the console screen whether a user is logged in to the console or not.</p>
	<p>Session</p> <p>A user logged in to the console device or connected to the CLI via a remote telnet or SSH session can also create a log with a destination type of 'session'. Events are displayed to the session device until the user logs off.</p>
	<p>Syslog</p> <p>Log events are sent to a syslog receiver.</p>
	<p>SNMP traps</p> <p>Events defined as SNMP traps are sent to the configured SNMP trap destinations and are logged in NOTIFICATION-LOG-MIB tables.</p>
	<p>File</p> <p>All selected log events will be directed to a file on one of the compact flash disks.</p>
	<p>Memory</p> <p>All selected log events will be directed to an in-memory storage area.</p>

Sample Output

```
A:ALA-1# show log log-collector
=====
Log Collectors
=====
Main          Logged   : 1224          Dropped   : 0
  Dest Log Id: 99   Filter Id: 0      Status: enabled   Dest Type: memory
  Dest Log Id: 100 Filter Id: 1001   Status: enabled   Dest Type: memory

Security      Logged   : 3           Dropped   : 0

Change       Logged   : 3896        Dropped   : 0

Debug        Logged   : 0           Dropped   : 0

=====
A:ALA-1#
```

log-id

Syntax	log-id [<i>log-id</i>] [severity <i>severity-level</i>] [application <i>application</i>] [sequence <i>from-seq</i> [<i>to-seq</i>]] [count <i>count</i>] [router <i>router-instance</i> [expression]] [message <i>message</i> [regular-expression]] [subject <i>subject</i> [regex]] [ascending descending] [message <i>format</i> [msg-regex]]
Context	show>log
Description	<p>This command displays an event log summary with settings and statistics or the contents of a specific log file, SNMP log, or memory log.</p> <p>If the command is specified with no command line options, a summary of the defined system logs is displayed. The summary includes log settings and statistics.</p> <p>If the log ID of a memory, SNMP, or file event log is specified, the command displays the contents of the log. Additional command line options control what and how the contents are displayed.</p> <p>Contents of logs with console, session or syslog destinations cannot be displayed. The actual events can only be viewed on the receiving syslog or console device.</p>
Parameters	<p>log-id — Displays the contents of the specified file log or memory log ID. The log ID must have a destination of an SNMP or file log or a memory log for this parameter to be used.</p> <p>Default Displays the event log summary</p> <p>Values 1 to 99</p> <p>severity <i>severity-level</i> — Displays only events with the specified and higher severity.</p> <p>Default All severity levels</p> <p>Values cleared, indeterminate, critical, major, minor, warning</p> <p>application <i>application</i> — Displays only events generated by the specified application.</p> <p>Default All applications</p> <p>The following values are examples of applications:</p> <p>Values bgp, cflowd, chassis, dhcp, debug, filter, igmp, ip, isis, lag, ldp, lldp, logger, mirror, mpls, oam, ospf, pim, port, ppp, rip, route_policy, rsvp, security, snmp, stp, svcmgr, system, user, vrrp, vrtr, ospf_ng, ntp</p> <p>expression — Specifies to use a regular expression as match criteria for the router instance string.</p> <p>sequence <i>from-seq</i> [<i>to-seq</i>] — Displays the log entry numbers from a particular entry sequence number (<i>from-seq</i>) to another sequence number (<i>to-seq</i>). The <i>to-seq</i> value must be larger than the <i>from-seq</i> value.</p>

If the *to-seq* number is not provided, the log contents to the end of the log is displayed unless the **count** parameter is present in which case the number of entries displayed is limited by the **count**.

Default All sequence numbers

Values 1 to 4294967295

count *count* — Limits the number of log entries displayed to the *number* specified.

Default All log entries

Values 1 to 4294967295

router-instance — Specifies a router name up to 32 characters to be used in the display criteria.

message *format* — Specifies a message string up to 400 characters to be used in the display criteria.

msg-regexp — Specifies to use a regular expression as parameters with the specified *message* string.

subject *subject* — Displays only log entries matching the specified text *subject* string. The subject is the object affected by the event, for example the port-id would be the subject for a link-up or link-down event.

regexp — Specifies to use a regular expression as parameters with the specified *subject* string..

ascending | **descending** — Specifies sort direction. Logs are normally shown from the newest entry to the oldest in **descending** sequence number order on the screen. When using the **ascending** parameter, the log will be shown from the oldest to the newest entry.

Default Descending

Output Show Log-ID Output

[Table 86](#) describes the log ID field output.

Table 86 Log-Id Output Field Descriptions

Label	Description
Log Id	An event log destination.
Source	no The event log filter is not currently in use by a log ID.
	yes The event log filter is currently in use by a log ID.

Table 86 Log-Id Output Field Descriptions (Continued)

Label	Description (Continued)
Filter ID	The value is the index to the entry which defines the filter to be applied to this log's source event stream to limit the events output to this log's destination. If the value is 0, then all events in the source log are forwarded to the destination.
Admin State	Up Indicates that the administrative state is up.
	Down Indicates that the administrative state is down.
Oper State	Up Indicates that the operational state is up.
	Down Indicates that the operational state is down.
Logged	The number of events that have been sent to the log source(s) that were forwarded to the log destination.
Dropped	The number of events that have been sent to the log source(s) that were not forwarded to the log destination because they were filtered out by the log filter.
Dest. Type	Console All selected log events are directed to the system console. If the console is not connected, then all entries are dropped.
	Syslog All selected log events are sent to the syslog address.
	SNMP traps Events defined as SNMP traps are sent to the configured SNMP trap destinations and are logged in NOTIFICATION-LOG-MIB tables.
	File All selected log events will be directed to a file on one of the CPM's compact flash disks.
	Memory All selected log events will be directed to an in-memory storage area.
Dest ID	The event log stream destination.
Size	The allocated memory size for the log.

Table 86 Log-Id Output Field Descriptions (Continued)

Label	Description (Continued)
Time format	The time format specifies the type of timestamp format for events sent to logs where log ID destination is either syslog or file. When the time format is UTC, timestamps are written using the Coordinated Universal Time value. When the time format is local, timestamps are written in the system's local time.

Sample Output

```
A:ALA-1# show log log-id
=====
Event Logs
=====
Log Source      Filter Admin Oper  Logged  Dropped Dest      Dest  Size
Id              Id      State State          Type          Id
-----
1  none         none   up   down   52      0      file     10    N/A
2  C            none   up   up     41      0      syslog   1     N/A
99 M           none   up   up    2135    0      memory   500
=====
A:ALA-1#
```

Sample Memory or File Event Log Contents Output

```
A:gal171# show log log-id 99
=====
Event Log 99
=====
Description : Default System Log
Memory Log contents [size=500 next event=70 (not wrapped)]

69 2007/01/25 18:20:40.00 UTC CRITICAL: SYSTEM #2029 Base Redundancy
"The active CPM card A is operating in singleton mode. There is no standby CPM card
."

68 2007/01/25 17:48:38.16 UTC WARNING: SYSTEM #2006 Base LOGGER
"New event throttle interval 10, configuration modified"

67 2007/01/25 00:34:53.97 UTC CRITICAL: SYSTEM #2029 Base Redundancy
"The active CPM card A is operating in singleton mode. There is no standby CPM card
."

66 2007/01/24 22:59:22.00 UTC CRITICAL: SYSTEM #2029 Base Redundancy
"The active CPM card A is operating in singleton mode. There is no standby CPM card
."

65 2007/01/24 02:08:47.92 UTC CRITICAL: SYSTEM #2029 Base Redundancy
"The active CPM card A is operating in singleton mode. There is no standby CPM card
."
...
=====
```

```
A:gal171

A:NS061550532>config>log>snmp-trap-group# show log log-id 1
=====
Event Log 1
=====
SNMP Log contents [size=100 next event=3 (not wrapped)]
Cannot send to SNMP target address 10.1.1.1.
Waiting to replay starting from event #2

14 2000/01/05 00:54:09.11 UTC WARNING: MPLS #2007 Base VR 1:
"Instance is in administrative state: inService, operational state: inService"

13 2000/01/05 00:54:09.11 UTC WARNING: MPLS #2008 Base VR 1:
"Interface linkToIxia is in administrative state: inService, operational state:
inService"
....
=====
A:NS061550532>config>log>snmp-trap-group#
```

snmp-trap-group

- Syntax** `snmp-trap-group [log-id]`
- Context** `show>log`
- Description** This command displays SNMP trap group configuration information.
- Parameters** *log-id* — Displays only SNMP trap group information for the specified trap group log ID.
Values 1 to 99
- Output** SNMP Trap Group Output

[Table 87](#) describes SNMP trap group output fields.

Table 87 SNMP Trap Group Output Fields

Label	Description
Log-ID	The log destination ID for an event stream.
Address	The IP address of the trap receiver,
Port	The destination UDP port used for sending traps to the destination, expressed as a decimal integer.
Version	Specifies the SNMP version format to use for traps sent to the trap receiver. Valid values are <code>snmpv1</code> , <code>snmpv2c</code> , <code>snmpv3</code> .
Community	The community string required by snmpv1 or snmpv2c trap receivers.

Table 87 SNMP Trap Group Output Fields (Continued)

Label	Description
Security-Level	The required authentication and privacy levels required to access the views on this node.
Replay	Indicates whether or not the replay parameter has been configured, enabled or disabled, for the trap-target address.
Replay from	Indicates the sequence ID of the first missed notification that will be replayed when a route is added to the routing table by which trap-target address can be reached. If no notifications are waiting to be replayed this field shows n/a.
Last Replay	Indicates the last time missed events were replayed to the trap-target address. If no events have ever been replayed this field shows never.

Sample Output

```
A:SetupCLI>config>log>snmp-trap-group# show log snmp-trap-group 44
=====
SNMP Trap Group 44
=====
Description : none
-----
Name       : ntt-test
Address    : 10.10.10.3
Port       : 162
Version    : v2c
Community  : ntttesting
Sec. Level : none
Replay     : disabled
Replay from : n/a
Last replay : never
-----
Name       : test2
Address    : 20.20.20.5
Port       : 162
Version    : v2c
Community  : ntttesting
Sec. Level : none
Replay     : disabled
Replay from : n/a
Last replay : never
=====
A:SetupCLI>config>log>snmp-trap-group#
```

syslog

Syntax `syslog [syslog-id]`

- Context** show>log
- Description** This command displays syslog event log destination summary information or detailed information on a specific syslog destination.
- Parameters** *syslog-id* — Displays detailed information on the specified syslog event log destination.
 Values 1 to 10
- Output** Syslog Event Log Destination Summary Output

 Table 88 describes the syslog output fields.

Table 88 Show Log Syslog Output Fields

Label	Description
Syslog ID	The syslog ID number for the syslog destination.
IP Address	The IP address of the syslog target host.
Port	The configured UDP port number used when sending syslog messages.
Facility	The facility code for messages sent to the syslog target host.
Severity Level	The syslog message severity level threshold.
Below Level Dropped	A count of messages not sent to the syslog collector target because the severity level of the message was above the configured severity. The higher the level, the lower the severity.
Prefix Present	Yes A log prefix was prepended to the syslog message sent to the syslog host.
	No A log prefix was not prepended to the syslog message sent to the syslog host.
Description	A text description stored in the configuration file for a configuration context.
LogPrefix	The prefix string prepended to the syslog message.
Log-id	Events are directed to this <i>destination</i> .

Sample Output

```
*A:ALA-48>config>log# show log syslog
=====
Syslog Target Hosts
=====
Id      Ip Address                               Port      Sev Level
=====
```

```

                Below Level Drop                Facility    Pfx Level
-----
2      unknown                514            info
        0                      local7         yes
3      unknown                514            info
        0                      local7         yes
5      unknown                514            info
        0                      local7         yes
10     unknown                514            info
        0                      local7         yes
=====
*A:ALA-48>config>log#

*A:MV-SR>config>log# show log syslog 1
=====
Syslog Target 1
=====
IP Address      : 192.168.15.22
Port            : 514
Log-ids         : none
Prefix         : Sr12
Facility       : local1
Severity Level  : info
Prefix Level    : yes
Below Level Drop : 0
Description     : Linux Station Springsteen
=====
*A:MV-SR>config>log#

```

5.13.2.2 Clear Commands

log

- Syntax** `log log-id`
- Context** clear
- Description** The **clear log log-id** command has been deprecated and replaced by the **clear log log-id log-id** command. The **clear log log-id** command continues to be supported, but it is recommended to use the **clear log log-id log-id** command instead.
- Parameters** *log-id* — Specifies the event log ID to be initialized/rolled over.
 - Values** 1 to 100

log-id

- Syntax** `log-id log-id`

Context	clear>log
Description	Reinitializes/rolls over the specified memory/file event log ID. Memory logs are reinitialized and cleared of contents. File logs are manually rolled over by this command. This command is only applicable to event logs that are directed to file destinations and memory destinations. SNMP, syslog and console/session logs are not affected by this command.
Parameters	<i>log-id</i> — Specifies the event log ID to be initialized/rolled over. Values 1 to 100

event-handling

Syntax	event-handling
Context	clear>log
Description	This command enables the context to clear Event Handling System (EHS) information.

handler

Syntax	handler <i>event-handler-name</i>
Context	clear>log>event-handling
Description	This command clears the counters in the show log event-handling handler event-handler-name output. It does affect the global or aggregate counters shown using the information command.
Parameters	<i>event-handler-name</i> — Specifies the name of the event handler, up to 32 characters.

information

Syntax	information
Context	clear>log>event-handling
Description	This command clears handler statistics in the show log event-handling information output.

6 sFlow

6.1 In This Chapter

This chapter provides information to configure sFlow and applies to the 7950 XRS, 7750 SR-7/12 and 7750 SR-12e platforms.

Topics in this chapter include:

- [sFlow Overview](#)
- [sFlow Features](#)
 - [sFlow Counter Polling Architecture](#)
 - [sFlow Support on Logical Ethernet Ports](#)
 - [sFlow SAP Counter Map](#)
- [sFlow Record Formats](#)

6.2 sFlow Overview

Some Layer 2 network deployments collect statistics on physical Ethernet ports and on Layer 2 interfaces at a high-frequency using a push model to, among others, monitor traffic, diagnose network issues, and/or provide billing. SR OS supports cflowd and XML accounting; however, those mechanisms are either Layer-3 specific, or focus on providing statistics at extremely large scale (thus use a pull model and cannot support high-frequency counter updates). To meet the statistics collection requirements of such Layer 2 deployments, SROS supports sFlow statistics export using sFlow version 5.

The following list gives the main caveats for sFlow support:

- sFlow data sources require multi-core line cards (IOM3 and later), enabling sFlow on a card that is not a multi-core is not blocked and can be detected by SNMP trap/log generated by sFlow
- To meet high-frequency export of counters, sFlow implementation is targeted for low per-port VLL/VPLS SAP scale only. The configuration is blocked if the per-port VLL/VPLS SAP limit exceeds sFlow limit. Contact your Nokia representative for per-platform scaling limits applicable.

6.3 sFlow Features

This section describes sFlow functionality supported in SR OS.

6.3.1 sFlow Counter Polling Architecture

When sFlow is enabled on an SROS router, the system takes upon a role of an sFlow network device as described in sFlow protocol version 5. A single sFlow agent can be configured for counter polling (flow sampling is not supported). There is no support for sub-agents.

The sFlow agent sends sFlow data to an operator-configured sFlow receiver. A single receiver is supported with configurable primary and backup IPv4 or IPv6 UDP destination sockets for redundancy (each sFlow packet exported is duplicated to both sockets when both are configured). The receiver's UDP sockets can be reachable either in-band or out-of-band (default) and must both be IPv4 or IPv6. An operator can also set the maximum size of the sFlow datagrams. Operators are expected to set this value to avoid IP fragmentation (Datagrams exceeding the specified size are fragmented before handed to IP layer).

The sFlow agent manages all sFlow data sources in the system. SROS supports sFlow data that are physical ports. When a port is configured as an sFlow data source, counters for that port and all VPLS and ePipe SAPs on that port are collected and exported using sFlow (see later on section for record format). Flow data sources can only be configured when an sFlow receiver is configured. To remove the sFlow receiver, all sFlow data sources must first be deconfigured at the port level.

Each data source is processed at a 15-second, non-configurable interval. If multiple data sources exist on a line card, the line card distributes the processing of each data source within a 15 second interval to avoid sFlow storms. When a timer expires to trigger a data source processing, data is collected for the physical port and for all VLL and VPLS SAPs on that port and exported using sFlow version 5 records as described in later subsections of this document. Each port and all SAP records for a given data source for a given interval are collected and sent with the counter sequence number and the timestamp value (the time value corresponds to the time counters were actually collected by a line card). The timestamp value uses line card's sysUptime value, which is synchronized with CPM time automatically by the system. A line card sends the counters to a CPM card, where sFlow UDP datagrams are created, sequenced with the CPM sequence number and sent to the receiver. If no UDP sockets are configured, no errors are generated because data is not sent. If no UDP sockets are reachable, the created UDP sFlow datagrams are dropped.



Note: Line cards will reset the counter record sequence numbers if, as a result of configuration or operational change, the return statistics no longer provide continuity with the previous interval. This may occur when:

- The card hard or soft resets
- The MDA resets
- The sFlow agent counter map changes



Note: The CPM will reset the sFlow datagram sequence numbers if, as a result of configuration or operational change, the sFlow datagram to be sent no longer provides continuity with the previous datagram. The following lists examples of when this takes place:

- HA switch
- CTL reboot
- Creation of an sFlow receiver

6.3.2 sFlow Support on Logical Ethernet Ports

sFlow data sources operate in a context of physical Ethernet port. To enable sFlow on Ethernet logical ports and their SAPs, an operator must explicitly enable sFlow on every physical Ethernet port that is a member of the given logical port. Currently only LAG logical ports are supported (including MC-LAG).



Note: sFlow configuration does not change automatically when a port is added or removed to or from a LAG.

For SAPs on a LAG, egress statistics will increment based on ports used by each SAP on LAG egress while ingress statistics will increment based on ports used by each SAP on LAG ingress unless LAG features like, for example, per-fp-ingress-queuing or per-fp-sap-optimization result in SAP statistics collection against a single LAG port.

If logical-level view is required, for example, per LAG statistics, a receiver is expected to perform data correlation based on per-physical port interface and SAP records exported for the given logical port's physical ports and their SAPs. sFlow data records contain information that allows physical ports/SAP records correlation to a logical port. See [sFlow Record Formats](#).



Note: Correlation of records must allow for small difference in timestamp values returned for member ports or SAP on a LAG because all ports run independent timestamps.

6.3.3 sFlow SAP Counter Map

To allow per SAP sFlow statistics export, operators must configure ingress and egress sFlow counter maps. The counter maps are required, because SROS systems support more granular per policer/queue counters and not IF-MIB counters per VLL/VPLS SAPs. In an absence of a map configured, 0's will be returned in corresponding statistics records.

A single ingress and a single egress counter map are supported. The maps specify which ingress and which egress SAP QoS policy queue/policer statistics map to sFlow unicast, multicast, and broadcast counters returned in an sFlow SAP record. Multiple queues and/or policers can map to each of unicast, multicast, broadcast counters. A single queue/policer can only map to one type of traffic. Queues, policers configured in a SAP QoS policy but not configured in an sFlow map or vice-versa are ignored when sFlow statistics are collected.

6.3.4 sFlow Record Formats

89 describes sFlow record used and exported:

Table 89 sFlow Record Fields

Record	Field	Value
sFlow Datagram Header (SAP and port)	Datagram version	5
	Agent Address	Active CPM IPv4 address (from BoF)
	Sub-agent ID	0
	Sequence number	CPM inserted sFlow datagram sequence number
	SysUptime	sysUptime when the counters for records included in the datagram were collected by the line card
	NumSamples	Number of counter records in the datagram

Table 89 sFlow Record Fields (Continued)

Record	Field	Value
Counter header (SAP and Port)	Enterprise	0 (standard sFlow)
	sFlow Sample Type	4 (Expanded counter sample)
	Sample Length	sFlow packet size excluding header
	Sequence number	Line card-inserted sequence number
	Source ID Type	0
	Source ID Index	tmnxPortId of the physical port (sFlow data source)
	Counter records	Count of counter records in the datagram
Ethernet Interface Counters (EIC) – port (Ethernet Layer)	Enterprise	Statistics returned are based on dot3StatsEntry in EtherLike-MIB.mib. Statistics support may depend on hardware type.
	Format	
	Flow data length	
	Alignment Errors	
	FCS Errors	
	Single Collision Frames	
	Multiple Collision Frames	
	SQE Test Errors	
	Deferred Transmissions	
	Late Collisions	
	Excessive Collisions	
	Internal Mac Transmit Errors	
	Carrier Sense Errors	
	Frame Too Longs	
Internal Mac Receive Errors		
Symbol Errors		

Table 89 sFlow Record Fields (Continued)

Record	Field	Value
Generic Interface Counters (GIC) – port/ SAP	Enterprise	0 (standard sFlow)
	Format	1 (GIC)
	Flow data length	88
	ifIndex	Port: ifIndex (tmnxPortId) of phys port SAP: SapEncapValue - part of SAP SNMP key
	ifType	Port: 6 (EthernetCsmacd) SAP: 1 (Other)
	ifSpeed	Port: Port speed value SAP: <ul style="list-style-type: none"> • top 32 bits: svclId for SAP (TIMETRA-SAP.mib) • lower 32 bits: sapPortId (TIMETRA-SAP.mib) The values plus ifIndex in the record are SAP SNMP key. SapPortId is LAG's tmnxPortId for SAPs on a LAG and port's tmnxPortId for SAPs on physical port
	ifDirection	Derived from MAU MIB (0 = unknown, 1 = full duplex, 2 = half duplex, 3 = in, 4 = out)
	ifAdminStatus	0 (down) 1 (up)
	ifOperStatus	0 (down) 1 (up)
	Input Octets	Statistics return for port are based on ifEntry or ifXEntry in IF-MIB.mib as applicable. Statistics returned for SAPs are sum of counters based on the sFlow ingress/egress counter map configured.
	Input Packets	
	Input Multicast packets	
	Input Broadcast packets	
Input Discarded packets		

Table 89 sFlow Record Fields (Continued)

Record	Field	Value
Generic Interface Counters (GIC) – port/ SAP (Continued)	Input Errors	Statistics return for port are based on ifEntry or ifXEntry in IF-MIB.mib as applicable. Statistics returned for SAPs are sum of counters based on the sFlow ingress/ egress counter map configured.
	Input Unknown Protocol Packets	
	Output Octets	
	Output Packets	
	Output Multicast packets	
	Output Broadcast packets	
	Output Discarded packets	
	Output Errors	
	Promiscuous Mode	0 (FALSE)

Notes:

- 0 is returned for statistics that are not supported by a given hardware type.
- If required, CPM executes rollover logic to convert internal 64-bit counters to a 32-bit sFlowd counter returned.

6.4 sFlow Command Reference

The commands listed in this section apply to the 7950 XRS, 7750 SR-12e, and 7750 SR-7/12 platforms.

6.4.1 Command Hierarchies

- [System Commands](#)
- [Show Commands](#)

To enable sFlow collection, an operator must enable sFlow on physical Ethernet ports in addition to the following configuration. Refer to the Ethernet Port Commands section in the SR OS Interface Configuration Guide for the CLI required to enable sFlow on physical ports.

6.4.1.1 System Commands

```

config
  — sflow
    — egress-counter-map {policer policer-id | queue queue-id} traffic-type {unicast |
      multicast | broadcast} [create]
    — no egress-counter-map {policer policer-id | queue queue-id }
    — ingress-counter-map { policer policer-id | queue queue-id } traffic-type { unicast |
      multicast | broadcast } [create]
    — no ingress-counter-map { policer policer-id | queue queue-id }
    — receiver receiver-name [create]
    — no receiver
      — ip-addr-primary ip-address[:port]
      — no ip-addr-primary
      — ip-addr-backup ip-address[:port]
      — no ip-addr-backup
      — max-data-size bytes
  
```

6.4.1.2 Show Commands

```

show
  — sflow
  
```

6.5 sFlow Configuration Command Descriptions

This section provides the sFlow configuration command descriptions.

6.5.1 Command Descriptions

The topics in this section include:

- [System Commands](#)
- [Show Commands](#)

6.5.1.1 System Commands

The following commands apply to the 7950 XRS, 7750 SR-12e, and 7750 SR-7/12 platforms.

sflow

Syntax	sflow
Context	config>sflow
Description	This command enables context to configured sflow agent parameters.

egress-counter-map

Syntax	egress-counter-map policer <i>policer-id</i> traffic-type {unicast multicast broadcast} [create] egress-counter-map queue <i>queue-id</i> traffic-type {unicast multicast broadcast} [create] no egress-counter-map policer <i>policer-id</i> no egress-counter-map queue <i>queue-id</i>
Context	config>sflow
Description	This command configures the egress counter map for sFlow. The map must be configured so sFlow agent understands how to interpret data collected against SAP queues and policers. Multiple queues and policers can be mapped to the same traffic-type using separate line entries.

The **no** form of this command deletes a SAP policy queue/policer from the map.

- Default** No mapping is created by default.
- Parameters** *policer-id* — Specifies the policer ID in a SAP egress QoS policy. If the SAP policy does not have a policer with the specified ID, the map entry will be ignored for this SAP.
Values 1 to 8
- queue-id* — Specifies the queue ID in a SAP egress QoS policy. If the SAP policy does not have a queue with the specified ID, the map entry will be ignored for this SAP.
Values 1 to 8

ingress-counter-map

- Syntax** **ingress-counter-map policer *policer-id* traffic-type {unicast | multicast | broadcast} [create]**
ingress-counter-map queue *queue-id* traffic-type {unicast | multicast | broadcast} [create]
no ingress-counter-map policer *policer-id*
no ingress-counter-map queue *queue-id*
- Context** config>sflow
- Description** This command configures the ingress counter map for sFlow. The map must be configured so sFlow agent understands how to interpret data collected against SAP queues and policers. Multiple queues/policers can be mapped to the same **traffic-type** using separate line entries.
- The **no** form of this command deletes a SAP policy queue/policer from the map.
- Default** No mapping is created by default.
- Parameters** *policer-id* — Specifies the policer ID in a SAP ingress QoS policy. If the SAP policy does not have a policer with the specified ID, the map entry will be ignored for this SAP.
Values 1 to 32
- queue-id* — Specifies the queue ID in a SAP ingress QoS policy. If the SAP policy does not have a queue with the specified ID, the map entry will be ignored for this SAP.
Values 1 to 32

receiver

- Syntax** **receiver *receiver-name* [create]**
no receiver
- Context** config>sflow

- Description** This command creates an sFlow receiver context or enters existing sFlow receiver context for the sFlow agent.
- The **no** form of this command deletes an existing sFlow receiver context.
- Default** No receivers are created by default.
- Parameters** *receiver-names* — String of up to 127 characters.

ip-addr-primary

- Syntax** **ip-addr-primary** *ip-address[:port]*
no ip-addr-primary
- Context** config>sflow>receiver
- Description** This command configures primary IPv4 or IPv6 destination address for the sFlow agent to send sFlow datagrams to. Optionally a destination port can also be configured (by default port 6343 is used).
- The **no** form of this command deletes primary sFlow receiver destination.
- Default** no ip-addr-primary
- Parameters** *ip-address* — Specifies the IPv4 or IPv6 address to send the sFlow datagrams to.

Values

a.b.c.d	(IPv4)
x:x:x:x:x:x:x	(IPv6)
[x:x:x:x:x:x]	(IPv6)
x - [0..FFFF]H	

port — Specifies the UDP destination port to send the sFlow datagrams to.

Values 1 to 65535

ip-addr-backup

- Syntax** **ip-addr-backup** *ip-address[:port]*
no ip-addr-backup
- Context** config>sflow>receiver
- Description** This command configures back-up IPv4 or IPv6 destination address for the sFlow agent to send sFlow datagrams to. Optionally a destination port can also be configured (by default port 6343 is used).

The **no** form of this command deletes backup sFlow receiver destination.

Default no ip-addr-backup

Parameters *ip-address* — Specifies the IPv4 or IPv6 address to send the sFlow datagrams to.

Values

a.b.c.d (IPv4)

x:x:x:x:x:x:x (IPv6)

[x:x:x:x:x:x:x] (IPv6)

x - [0..FFFF]H

port — Specifies the UDP destination port to send the sFlow datagrams to.

Values 1 to 65535

max-data-size

Syntax max-data-size *bytes*

Context config>sflow>receiver

Description This configures maximum data size for sFlow UDP datagrams sent to the collector.

To restore default configuration, execute max-data-size 1400.

Default 1400

Parameters *bytes* — An integer

Values 200 to 1500

6.6 sFlow Show Command Descriptions

This section provides the sFlow show command descriptions.

6.6.1 Command Descriptions

The commands described in this section apply to the 7950 XRS, 7750 SR-12e, and 7750 SR-7/12 platforms.

The command outputs in this section are examples only; actual displays may differ depending on supported functionality and user configuration.

6.6.1.1 Show Commands

sflow

Syntax	sflow
Context	show>sflow
Description	This command displays the primary and backup receiver statistics, the mapping configuration and a summary of how many ports and SAPs have sFlow enabled.

[Table 90](#) describes the show sflow output fields.

Table 90 Show Sflow Output Fields

Label	Description
sFlow Status	
Receiver	Displays the configured name for the sFlow receiver.
Max Data Size	The configured maximum data size for sFlow UDP packets.
IP Addr Primary	The primary IP address and destination port for sFlow receiver.
IP Addr Backup	The backup IP address and destination port for sFlow receiver.

Table 90 Show Sflow Output Fields (Continued)

Label	Description
Packets Sent	The number of packets sent successfully to the primary or backup receiver destination, since the destination was configured, CPM card HA switchover, or system reboot.
Packet Errors	The number of packets that could not be sent to the primary or backup receiver destination because of an error, since the destination was configured, CPM card HA switchover, or system reboot. An example of an error is destination IP not reachable.
Last Packet Sent	Displays the date and time of the last packet sent.
Counter Pollers	
Port	Displays the port on which sFlow is enabled.
No. of SAPs	The number of SAPs on the port with sFlow enabled.
No. of sFlow counter pollers	The number of sFlow counter pollers.
Counter Mappings	
Direction	Displays the direction of traffic (ingress or egress) the map entry applies to.
Policer/Queue	Displays the policer or queue instance being mapped by sFlow map.
Traffic type	Displays the type of sFlow traffic statistics (unicast, multicast or broadcast) that the policer/queue maps to.
No. of sFlow counter mappings	The number of entries in the sFlow ingress and egress counter map.

Output

Sample Output

```
*B:bkvm10# show sflow
=====
sFlow Status
=====
Receiver           : pat
Max Data Size     : 312

IP Addr Primary   : 138.120.142.163:6343
Packets Sent      : 2572
Packet Errors     : 2
Last Packet Sent  : 07/08/2014 22:23:57nt
```

IP Addr Backup : N/A
Packets Sent : 0
Packet Errors : 0
Last Packet Sent : No Pkts sent

Counter Pollers

Port	No. of SAPs
1/1/2	3
1/2/1	0

No. of sFlow counter pollers: 2

Counter Mappings

Direction	Policer/Queue	Traffic Type
egress	queue 1	unicast
egress	queue 5	multicast
egress	queue 8	broadcast
ingress	policer 1	unicast
ingress	policer 6	multicast
ingress	policer 12	broadcast

No. of sFlow counter mappings: 6

=====

7 Telemetry

7.1 In This Chapter

This chapter provides information to configure Telemetry.


Topics in this chapter include:

- [Telemetry Overview](#)
- [About Telemetry](#)
- [Telemetry Examples](#)

7.2 Telemetry Overview

Telemetry is a network monitoring and fault management framework. It is driven by natural network growth (network volume increases) and the need to use fresh data obtained from the network to make fast networking decisions such as traffic optimization and preventive troubleshooting.

Unlike legacy monitoring platforms such as SNMP, Telemetry does not only rely on collectors to continuously pull data from the network elements. Instead, network devices push and stream data (such as statistics) continuously to collectors based on subscriptions. Collectors can then filter, analyze, store, and make decisions using the collected data from the network devices. [Figure 21](#) illustrates this process.

Figure 21 Telemetry Application

7.3 About Telemetry

Telemetry uses the proprietary NOKIA SR OS YANG data models to stream data that is encoded as Google Protocol Buffers (gPB) messages. Google Remote Procedure Call (gRPC) is the transport used to subscribe to the SR OS device and receive streamed telemetry data. SR OS supports gPB version 3.0.0-b2.

7.3.1 gRPC in Telemetry

The gRPC transport method uses HTTP/2 bidirectional streaming between the gRPC client (the collector) and the gRPC server (the SR OS device). A gRPC session is a single connection from the gRPC client to the gRPC server over the TCP/TLS port. A gRPC session can be used by:

- a gRPC client to send a telemetry subscription request to the gRPC server
- a gRPC server to send asynchronous telemetry data to the gRPC collector

A gRPC channel is a single RPC call.

The gRPC version supported on the SR OS gRPC server is 1.0.1.

The SR OS gRPC encryption and authentication follows the basic conventions described in the OpenConfig `gnmi-authentication.md` published on github.com (version 0.1.0 from Oct 5, 2016).

TLS encryption is used for added security. The following summarizes the process of encryption and authentication:

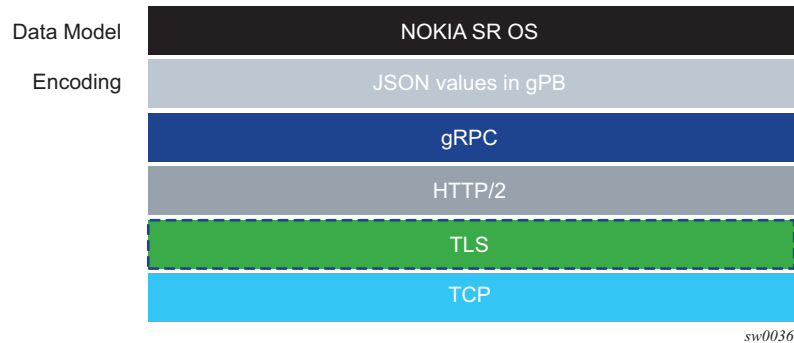
- SR OS device authentication
 - The gRPC clients do not share gRPC sessions. Each gRPC client should initially start a separate gRPC session.
 - When a gRPC session is established, the gRPC server certificates are verified by the gRPC client to ensure every gRPC server is authenticated by the gRPC client.
 - When gRPC is shutdown on the gRPC server and a gRPC client is trying to establish a gRPC session, the gRPC client will get an error for every RPC.
 - When a gRPC session is established, gRPC is shutdown on the gRPC server, all active RPCs are gracefully terminated, and an error is returned for every RPC.
- TLS encryption
 - The gRPC session should be in an encrypted state before it can be used.
 - If the gRPC client and gRPC server are unable to negotiate an encrypted gRPC session, the gRPC session fails and the gRPC server sends an error.
 - Fallback from an encrypted to an unencrypted gRPC session is not allowed.

For information on how to configure TLS with gRPC, see the [TLS](#) chapter.

- User authentication
 - Each RPC sent by the gRPC client carries a user/password.
 - For the first RPC on the gRPC session, the gRPC server tries to authenticate the user via the specified authentication order; for instance, local user database, RADIUS, or TACACS+.
For example, if TACACS+ is first in the authentication order, the gRPC server sends a request to the TACACS+ server to authenticate the gRPC user.
 - For the subsequent RPCs on that same authenticated gRPC session, the user/password are re-authenticated only if changed.
 - When there is no user/password provided with the RPC, the gRPC server returns an error.
 - If the RPC user is changed, then any active subscriber RPCs on that same gRPC session are terminated by the gRPC server.
 - If the RPC password is changed, then the active gRPC session will continue to exist until a different user/password is sent in a subsequent RPC, or the gRPC session is terminated.
 - Each telemetry message is carried over an encrypted gRPC session which was previously encrypted; the session is not re-encrypted.

Figure 22 shows the telemetry protocol stack.

Figure 22 Telemetry Stack



The gRPC service runs on port 57400 by default on the SR OS. The service is not configurable.

A single gRPC server supports concurrent gRPC sessions and channels.

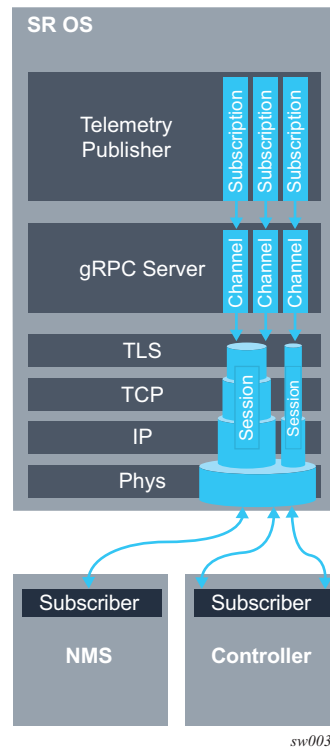
- There are a maximum of eight concurrent gRPC sessions for all of the gRPC clients.
- There are a maximum of 225 concurrent gRPC channels for all of the gRPC clients. Since each RPC is a unique channel, the maximum number of subscriptions for all the gRPC clients on a single SR OS device is 225.

Closing a gRPC channel terminates an active Telemetry subscription. A gRPC session that is used by the disconnected subscription is not be terminated. Closing the entire gRPC session terminates all active Telemetry subscriptions on the disconnected gRPC session.

A Telemetry subscription can be administratively terminated from the CLI. An active gRPC session that is used by the terminated subscription is not terminated. See the CLI section for command details.

Figure 23 shows a gRPC service using the TLS architecture.

Figure 23 gRPC Using TLS Architecture



7.3.2 Operations Layer

This section summarizes support for subscription requests and subscription responses.

SR OS Telemetry follows the OpenConfig gnmi.proto published on github.com (version 0.2.0, from Nov 8, 2016). This model defines the relationship and behavior between the gRPC client and server.

SR OS Telemetry follows the basic conventions described in the OpenConfig gnmi-specification.pdf published on github.com (version 0.2.1 from Nov 10, 2016).

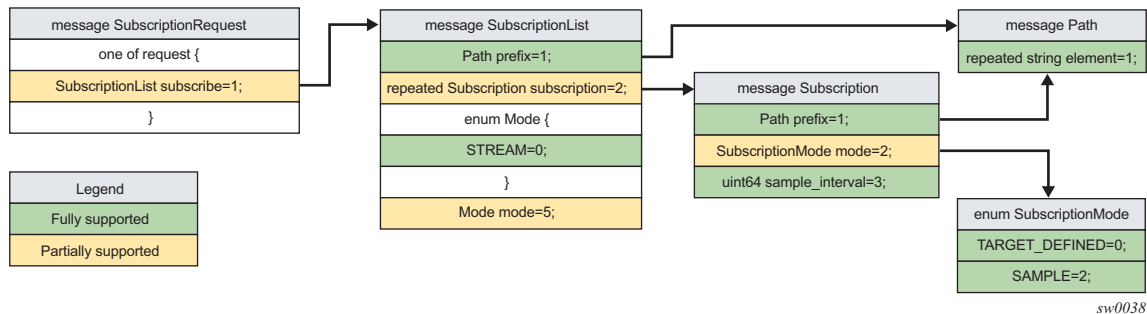
A subscription is initiated from the gRPC client by sending a "Subscribe" RPC that contains a "SubscribeRequest" message to the gRPC server. A "prefix" can be specified to be used with all paths specified in the "SubscribeRequest". If a "prefix" is present then it is logically appended to the start of every "path" to provide a full "path".

A "subscription" contains:

- A "path" list of one or more paths:
 - A path represents the data tree as a series of repeated strings/elements. Each element represents a data tree node name and its associated attributes.
 - A path should be syntactically valid within the set of schema modules that the gRPC server supports.
 - That list cannot be modified throughout the lifetime of the subscription.
 - If the subscription path is to a container node, then all children leaves of that container node are considered to be subscribed to.
 - Any specified path must be unique within the list (paths cannot be repeated within the list). An error is returned upon using the same path more than once in a single subscription.
 - A specified path does not need to pre-exist within the current data tree on the gRPC server. In the case that a particular path does not exist, the gRPC server continues to monitor for the existence of the path, and transmits telemetry updates if the path exist in the future.
 - The gRPC server does not send any data for a non-existing path. For instance, if a path is non-existing at the time of subscription creation or if the path was deleted after the subscription is established.
 - The maximum number of explicit paths per a single subscription that can be specified is 64. This means that the maximum number of explicit paths per all subscriptions on a single SR OS device is 14400 (225 subscriptions multiplied by 64 paths). Upon receiving a SubscribeRequest message that is trying to subscribe to more than 64 explicit paths, an error is returned by the SR OS device. A path using a wildcard is still considered a single explicit path.
- A subscription mode:
 - "SAMPLE" mode is supported for each path, where the gRPC server sends notifications at the specified sampling interval.
 - Using "TARGET_DEFINED" mode still means "SAMPLE" mode.
- A sample interval:
 - A "sample_interval" is supported for each path. A sample interval of 0 means 10 seconds by default. If a "sample_interval" of less than 10 seconds is specified, the gRPC server returns an error. A sample interval is specified in nano seconds.

Figure 24 illustrates the SR OS support of a subscription request.

Figure 24 Subscription Request



When a subscription is successfully initiated on the gRPC server, “SubscribeReponse” message are sent from the gRPC server to the gRPC client. One set of messages is sent with every "sample_interval". The "SubscribeResponse" message contains "update" notifications as per the subscription's path list.

A "sync_response" notification is sent every time the gRPC server sends all of the updates for the subscribed-to paths. The "sync_response" must be set to true for the gRPC client to consider the stream has synced. A "sync_response" is used to signal the gRPC client that it has a full view of the subscribed-to data.

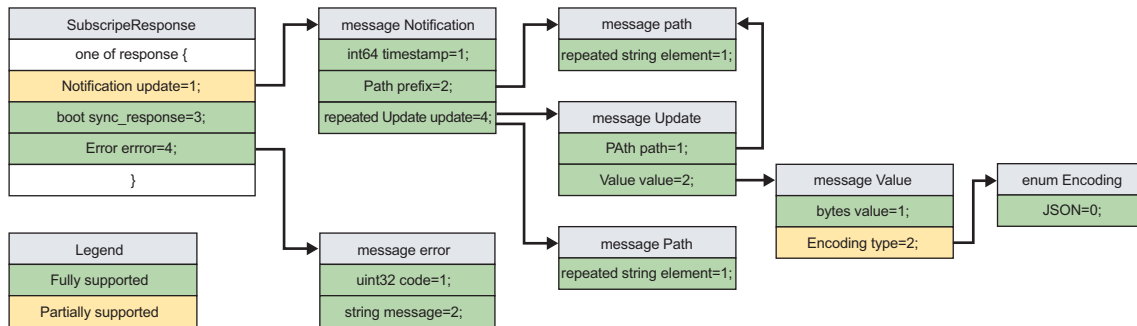
The gRPC server sends an error if required. The error contains a description of the context of the error.

An "update" notification contains:

- A "timestamp" of the statistics collection time. Time stamps are always represented as nanoseconds.
- A "prefix":
 - If a "prefix" is present, then it is logically appended to the start of every "path" to provide the full path.
 - The presence of a "prefix" in the "SubscriptionResponse" message is not related to the presence of a "prefix" in the original "SubscriptionRequest" message. The "prefix" in the "SubscriptionResponse" message is optimized by the gRPC server.
- A list of "update" path and value pairs.
 - A path represents the data tree path as a series of repeated strings/elements. Where each element represents a data tree node name and its associated attributes. See the [Schema Paths](#) section for more information.
 - The “Value” message represents the data tree node’s “value” and “encoding” which is always “JSON”.

[Figure 25](#) illustrates the SR OS support of a subscription response.

Figure 25 Subscription Response



sw0039

7.3.3 Schema Paths

Telemetry subscriptions include a set of schema paths used to identify which data nodes are of interest to the collector.

The paths in Telemetry 'Subscribe' RPC requests follow the basic conventions described in the OpenConfig gnmi-path-conventions.md published on github.com (version 0.1.0 from Nov 11, 2016).

A path consists of a set of path segments often shown with a '/' character as a delimiter. For example: `configure/router[router-instance=Base]/interface[interface-name=my-interface1]/description`.

These paths are encoded as a set of individual string segments in gnmi.proto (without any '/' characters). For example, `["configure", "router[router-instance=Base]", "interface[interface-name=my-interface1]", "description"]`

A path selects an entire subtree of the data model and includes all descendants of the node indicated in the path. The following table summarizes the types of paths that are supported in SR OS telemetry:

Table 91 Schema Paths

Path example	Description
<code>/configure/router[router-instance=Base]/interface[interface-name=abc]</code>	Selects all config leafs of interface abc and all descendants.
<code>/configure/router[router-instance=Base]/interface[interface-name=abc]/description</code>	Selects only the description leaf of interface abc.

Table 91 Schema Paths (Continued)

Path example	Description
/state/router[router-instance=Base]/interface[interface-name=*]	Selects all state information for all Base router interfaces. Wildcard in a single segment of a path.
/configure/router[router-instance=Base]/interface[interface-name=*]/description	Selects the description leaf for all Base router interfaces. Wildcard in a single segment of a path.
/	The root path. This selects all config and state data from all models (in all namespaces) supported on the router. Encoded as "" in gRPC/gPB.

The following items describe types of telemetry paths that are not supported in SR OS:

- Wildcards for entire path segments are not supported.
For example, /state/service/*/oper-status
- If a wildcard is used for any key of a list, then a wildcard must be used for all the keys of that list. In a single path segment, all the keys must either have specific values or all the keys must have wildcards. A mix of wildcards and specific values for different parts of a list key is not supported.
For example: /state/cflowd/collector[ip-address=138.120.44.45][port=*]/oper-status
- Functions such as 'current()', 'last()' and mathematical operators, such as stat<5 or octets>3 are not supported in paths. The '|' (OR operator, used to select multiple paths) is not supported.
- Wildcards in multiple segments of a path are not supported.
For example: /state/card[slot-number=*]/mda[mda-slot=*]
- The '//' wildcard pattern is not supported.
For example: /state//oper-status

7.4 Telemetry Examples

This section contains examples of Telemetry subscription requests and responses. The following examples are dumps of protobuf messages from a python API. Format may vary across different implementations.

Example 1 — Subscribe to a single path

```

2017-01-20 12:34:51,594 - SENT::SubscribeRequest
subscribe {
  subscription {
    path {
      element: "state"
      element: "router[router-instance=Base]"
      element: "interface[interface-name=test]"
      element: "statistics"
      element: "ip"
      element: "in-packets"
    }
    mode: SAMPLE
    sample_interval: 30000000000
  }
}

```

```

2017-01-20 12:34:51,605 - RCVD::SubscribeResponse
2017-01-20 12:35:21,611 - RCVD::Subscribe
update {
  timestamp: 1484912121607764002
  prefix {
    element: "state"
    element: "router[router-instance=Base]"
    element: "interface[interface-name=test]"
    element: "statistics"
    element: "ip"
  }
  update {
    path {
      element: "in-packets"
    }
    value {
      value: "0"
      type: JSON
    }
  }
}
2017-01-20 12:35:21,650 - RCVD::Subscribe
sync_response: true

```

```

2017-01-20 12:35:51,612 - RCVD::Subscribe
update {
  timestamp: 1484912151608586530
  prefix {
    element: "state"
    element: "router[router-instance=Base]"
    element: "interface[interface-name=test]"
    element: "statistics"
    element: "ip"
  }
  update {
    path {
      element: "in-packets"
    }
    value {
      value: "16"
    }
  }
}

```



```

        type: JSON
      }
    }
  }
}
2017-01-20 12:35:51,614 - RCVD::Subscribe
sync_response: true
....
....

```

Example 2 — Subscribe to a single path with wild card

```

2017-01-24 08:58:06,175 - SENT::SubscribeRequest
subscribe {
  subscription {
    path {
      element: "state"
      element: "router[router-instance=Base]"
      element: "interface[interface-name=*)"
      element: "statistics"
      element: "ip"
      element: "in-packets"
    }
    mode: SAMPLE
    sample_interval: 30000000000
  }
}

2017-01-24 08:58:06,181 - RCVD::SubscribeResponse
2017-01-24 08:58:36,191 - RCVD::Subscribe
update {
  timestamp: 1485244716188240643
  prefix {
    element: "state"
    element: "router[router-instance=Base]"
    element: "interface[interface-name=system]"
    element: "statistics"
    element: "ip"
  }
  update {
    path {
      element: "in-packets"
    }
    value {
      value: "0"
      type: JSON
    }
  }
}

2017-01-24 08:58:36,231 - RCVD::Subscribe
update {
  timestamp: 1485244716192259548
  prefix {
    element: "state"
    element: "router[router-instance=Base]"
    element: "interface[interface-name=to_node_B]"
    element: "statistics"
    element: "ip"
  }
}

```

```
update {
  path {
    element: "in-packets"
  }
  value {
    value: "0"
    type: JSON
  }
}
}
2017-01-24 08:58:36,233 - RCVD::Subscribe
update {
  timestamp: 1485244716194644789
  prefix {
    element: "state"
    element: "router[router-instance=Base]"
    element: "interface[interface-name=to_node_D]"
    element: "statistics"
    element: "ip"
  }
  update {
    path {
      element: "in-packets"
    }
    value {
      value: "0"
      type: JSON
    }
  }
}
}
2017-01-24 08:58:36,235 - RCVD::Subscribe
sync_response: true

2017-01-24 08:59:06,192 - RCVD::Subscribe
update {
  timestamp: 1485244746189318112
  prefix {
    element: "state"
    element: "router[router-instance=Base]"
    element: "interface[interface-name=system]"
    element: "statistics"
    element: "ip"
  }
  update {
    path {
      element: "in-packets"
    }
    value {
      value: "0"
      type: JSON
    }
  }
}
}
2017-01-24 08:59:06,196 - RCVD::Subscribe
update {
  timestamp: 1485244746193708158
  prefix {
    element: "state"
    element: "router[router-instance=Base]"
```

```
        element: "interface[interface-name=to_node_B]"
        element: "statistics"
        element: "ip"
    }
    update {
        path {
            element: "in-packets"
        }
        value {
            value: "0"
            type: JSON
        }
    }
}
2017-01-24 08:59:06,199 - RCVD::Subscribe
update {
    timestamp: 1485244746196077911
    prefix {
        element: "state"
        element: "router[router-instance=Base]"
        element: "interface[interface-name=to_node_D]"
        element: "statistics"
        element: "ip"
    }
    update {
        path {
            element: "in-packets"
        }
        value {
            value: "0"
            type: JSON
        }
    }
}
2017-01-24 08:59:06,200 - RCVD::Subscribe
sync_response: true
....
....
```

Example 3: Subscribe to more than one path

```
2017-01-24 12:54:18,228 - SENT::SubscribeRequest
subscribe {
    subscription {
        path {
            element: "state"
            element: "router[router-instance=Base]"
            element: "interface[interface-name=to_node_B]"
        }
        mode: SAMPLE
        sample_interval: 30000000000
    }
    subscription {
        path {
            element: "state"
            element: "router[router-instance=Base]"
            element: "mpls"
            element: "statistics"
        }
    }
}
```

```

        element: "lsp-egress-stats[lsp-name=lsp_to_dest_f]"
    }
    mode: SAMPLE
    sample_interval: 30000000000
}
}

```

Example 4: Subscribe to a list with wild card

```

2017-01-24 13:45:30,947 - SENT::SubscribeRequest
subscribe {
  subscription {
    path {
      element: "state"
      element: "router[router-instance=Base]"
      element: "interface[interface-name=*"
    }
    mode: SAMPLE
    sample_interval: 30000000000
  }
}

```

Example 5: Subscribe to path where the object did not exist before subscription

```

2017-01-24 13:53:50,165 - SENT::SubscribeRequest
subscribe {
  subscription {
    path {
      element: "state"
      element: "router[router-instance=Base]"
      element: "interface[interface-name=to_node_B]"
    }
    mode: SAMPLE
    sample_interval: 30000000000
  }
}

```

```

2017-01-24 13:53:50,166 - RCVD::SubscribeResponse
2017-01-24 13:54:20,169 - RCVD::Subscribe
sync_response: true

```

```

2017-01-24 13:54:50,174 - RCVD::Subscribe
update {
  timestamp: 1485262490169309451
  prefix {
    element: "state"
    element: "router[router-instance=Base]"
    element: "interface[interface-name=to_node_B]"
  }
  update {
...
...

```

Example 6: Subscribe to a path where the object existed before subscription then got deleted after subscription

```
2017-01-24 14:00:41,292 - SENT::SubscribeRequest
subscribe {
  subscription {
    path {
      element: "state"
      element: "router[router-instance=Base]"
      element: "interface[interface-name=to_node_B]"
    }
    mode: SAMPLE
    sample_interval: 30000000000
  }
}

2017-01-24 14:00:41,294 - RCVD::SubscribeResponse
2017-01-24 14:01:11,295 - RCVD::Subscribe
update {
  timestamp: 1485262871290064704
  prefix {
    element: "state"
    element: "router[router-instance=Base]"
    element: "interface[interface-name=to_node_B]"
  }
  update {
    ...
    ...
  }
}
2017-01-24 14:01:11,359 - RCVD::Subscribe
sync_response: true

2017-01-24 14:01:41,293 - RCVD::Subscribe
sync_response: true

2017-01-24 14:02:11,296 - RCVD::Subscribe
sync_response: true
```

7.5 gRPC Command Reference

The commands listed in this section apply to the 7950 XRS, 7750 SR-12e, and 7750 SR-7/12 platforms.

7.5.1 Command Hierarchies

7.5.1.1 System Commands

```
config
  — system
    — grpc
      — max-msg-size number
      — no max-msg-size
      — no shutdown
      — tls-server-profile name
      — no tls-server-profile
```

7.5.1.2 QoS Commands

```
config
  — router
    — sgt-qos
      — application
        — grpc
          — dscp dscp-value
```

7.6 Telemetry Configuration Command Descriptions

This section provides Telemetry configuration command descriptions.

7.6.1 Command Descriptions

The topics in this section include:

- [System Commands](#)
- [QoS Commands](#)

7.6.1.1 System Commands

grpc

Syntax	<code>grpc</code>
Context	<code>config>system</code> <code>config>router>sgt-qos>application</code>
Description	This command enables the context to configure gRPC parameters.

max-msg-size

Syntax	<code>max-msg-size <i>number</i></code> <code>no max-msg-size</code>
Context	<code>config>system>grpc</code>
Description	This command configures the maximum gRPC rx message size
Parameters	<i>number</i> — Specifies the maximum message size in MB.
Values	1 to 1024
Default	512

shutdown

Syntax	no shutdown
Context	config>system>grpc
Description	This command disables the gRPC server. The shutdown command is not blocked if there are active gRPC sessions. Shutting down gRPC will terminate all active gRPC sessions.

tls-server-profile

Syntax	tls-server-profile <i>name</i> no tls-server-profile
Context	config>system>grpc
Description	This command provides the TLS profile name to use for the gRPC server.
Parameters	<i>name</i> — The tls-server profile name. Values 32 characters maximum

7.6.1.2 QoS Commands

dscp

Syntax	dscp { <i>dscp-value</i> <i>dscp-name</i> }
Context	config>router>sgt-qos>application>grpc
Description	This command configures a DiffServ Code Point (DSCP) name to be used for gRPC.
Parameters	<i>dscp-value</i> — Represents the gRPC traffic class. <i>dscp-name</i> — Represents the gRPC traffic class.

7.7 gRPC Show, Admin Command Reference

This section provides the gRPC show and admin command descriptions.

7.7.1 Command Hierarchies

- [Show Commands](#)
- [Admin Commands](#)

7.7.1.1 Show Commands

```
show
  — system
    — grpc

show
  — router
    — sgt-qos
      — application
        — grpc
          — dscp
```

7.7.1.2 Admin Commands

```
admin
  — system
    — telemetry
      — grpc
        — subscription subscription-id cancel
        — subscription cancel-all

admin
  — disconnect {grpc}
```

7.7.2 Command Descriptions

- [Show Commands](#)

- [Admin Commands](#)

7.7.2.1 Show Commands

grpc

- Syntax** **grpc**
- Context** show>system
- Description** This command displays the gRPC server status.
- Output** SHH options output describes the gRPC Output fields.

Table 92 Show System gRPC output Fields

Labels	Description
gRPC Server	
Administrative State	Enabled Displays that gRPC is enabled. Disabled Displays that gRPC is disabled.
Operational State	Up Displays that gRPC is operational. Down Displays that gRPC is not operational.

Sample Output

```

=====
gRPC Server
=====
Administrative State      : Disabled
Operational State       : Down
=====
    
```

grpc

- Syntax** **grpc**
- Context** show>router>sgt-qos>application

Description This command displays the gRPC router status.

dscp

Syntax **dscp**

Context show>router>sgt-qos>application>grpc

Description This command shows the configured DiffServ Code Point (DSCP) name/value for gRPC.

7.7.2.2 Admin Commands

subscription

Syntax **subscription *subscription-id* cancel**

Context admin>system>telemetry>grpc

Description This command cancels an active Telemetry subscription.

Parameters *subscription-id* — The ID of the Telemetry subscription to cancel.

Values 0 to 4294967295

Default none

subscription cancel-all

Syntax **subscription cancel-all**

Context admin>system>telemetry>grpc

Description This command cancels all active Telemetry subscriptions.

disconnect

Syntax **disconnect {grpc}**

Context admin

Description This command disconnects all active gRPC sessions.

8 TLS

8.1 In This Chapter

This chapter provides information to configure Transport Layer Security (TLS).

8.2 TLS Overview

Transport Layer Security (TLS) is used for two primary purposes:

- authentication of an end device (client or server) using a digital signature (DS)

TLS uses PKI for device authentication. DSs are used to authenticate the client or the server. The server typically sends a certificate with a DS to the client.

In certain situations, the server can request a certificate from the client to authenticate it. The client has a certificate (called a Trust Anchor) from the certificate authority (CA) which is used to authenticate server certificate and its DS. After the client provides a digitally signed certificate to the server and both parties are authenticated, the encryption PDUs can then be transmitted.

When SR OS is acting as a server and it requests a certificate from the client, the client must provide the certificate. If the client fails to provide a certificate for authentication, SR OS will terminate the TLS session. The server TLS settings can be configured to not request certificates, in which case the client is not obligated to send the server a certificate for authentication.

- encryption and authentication of application PDUs

After the clients and server have been successfully authenticated, the cipher suite is negotiated between the server and clients, and the PDUs will be encrypted based on the agreed cipher protocol.

8.3 TLS Server Interaction with Applications

TLS is a standalone configuration. The user must configure TLS server profiles with certificates and trust anchors, and then assign the TLS server profiles to the appropriate applications. When a TLS server profile is assigned to an application, the application should not send any clear text PDUs until the TLS handshake has been successfully completed and the encryption ciphers have been negotiated between the TLS server and the TLS client.

After successful negotiation and handshake, the TLS will be operationally up, and the TLS will notify the application which will begin transmitting PDUs. These PDUs will be encrypted using TLS based on the agreed ciphers. If, at any point, the TLS becomes operationally down, the application should stop transmitting PDUs.

For example, a TLS connection with the gMI application would operate as follows:

1. A TLS server profile is assigned to the gMI application.
2. gMI stops sending clear text PDUs because a TLS server profile has been assigned and TLS is not ready to encrypt.
3. The TLS server begins the handshake.
4. Authentication occurs at the TLS layer.
5. The TLS server and TLS client negotiate ciphers.
6. SALTs are negotiated for the symmetric key. A SALT is a seed for creating AES encryption keys.
7. When negotiations are successfully completed, the handshake finishes and gMI is notified.
8. TLS becomes operationally up, and gMI can resume transmitting PDUs. Until TLS becomes operationally up, gMI PDUs arriving from the client are dropped on ingress.

8.3.1 TLS Application Support

[Table 93](#) lists the applications that support TLS.

Table 93 TLS Application Support

Application	TLS Server Supported	TLS Client Supported
LDAP	NO	YES
GRPC	YES	NO

Table 93 TLS Application Support (Continued)

Application	TLS Server Supported	TLS Client Supported
OPEN-FLOW	YES	YES

8.4 TLS Handshake

Figure 26 shows the TLS handshake.

Figure 26 TLS Handshake

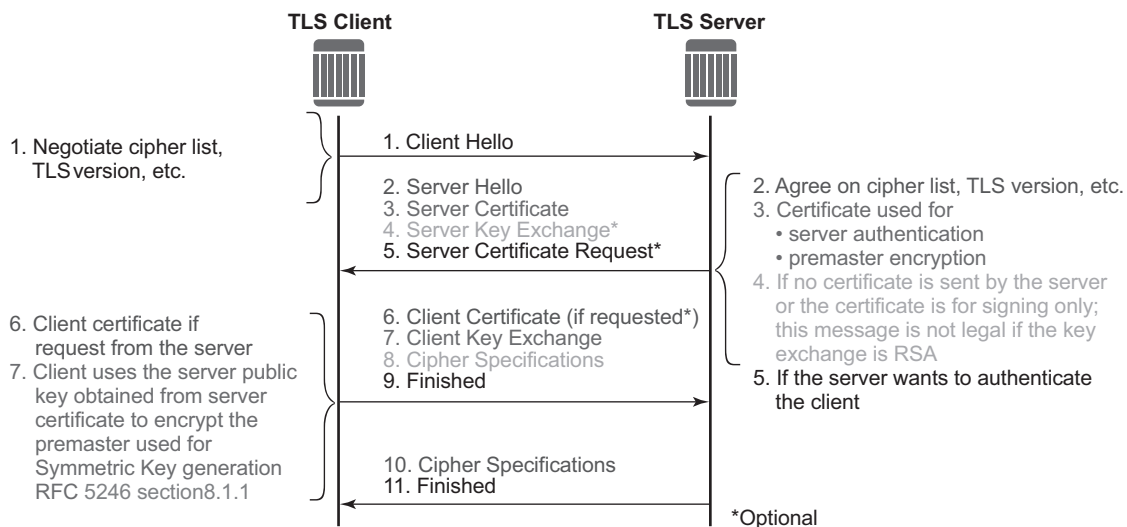


Table 94 further describes the steps in the TLS handshake.

Table 94 TLS Handshake Step Descriptions

Step	Description
1	The TLS handshake begins with the client Hello message. This message includes the cipher list that the client wishes to use and negotiate, among other information.
2	The TLS server sends back a server Hello message, along with the first common cipher found on both the client cipher list and the server cipher list. This agreed cipher will be used for data encryption.

Table 94 TLS Handshake Step Descriptions (Continued)

Step	Description
3	The TLS server continues by sending a server certificate message, where the server provides a certificate to the client so that the client can authenticate the server identity. The public key of this certificate (RSA key) can also be used for encryption of the symmetric key seed that will be used by the client and server to create the symmetric encryption key. This occurs only if the PKI is using RSA for asymmetric encryption.
4	Server key exchange is not supported by SR OS. SR OS only uses RSA keys; Diffie-Hellman key exchange is not supported.
5	The server can optionally be configured to request a certificate from the client to authenticate the client.
6	If the server has requested a certificate, the client should provide a certificate using a client certificate message. If the client does not provide a certificate, the server will drop the TLS session.
7	The client uses the server public RSA key that was included in the server certificate to encrypt a seed used for creating the symmetric key. This seed is used by the client and server to create the identical symmetric key for encrypting and decrypting the data plane traffic.
8	The client sends a cipher spec to switch encryption to this symmetric key.
9	The client successfully finishes the handshake.
10	The server sends a cipher spec to switch encryption to this symmetric key.
11	The server successfully finishes the handshake.

After a successful handshake, TLS will be operationally up, and applications can then use it for application encryption.

8.5 TLS Client Certificate

TLS protocol is used for authentication, and as such, the server can ask to authenticate the client via PKI. If the server requests authentication from the client, the client must provide an X.509v3 certificate to the server so that it can be authenticated via the digital signature of its client. SR OS allows the configuration of an X.509v3 certificate for TLS clients. When the server requests a certificate via the server's Hello message, the client will transmit its certificate to the server using a client certificate message.

8.6 TLS Symmetric Key Rollover

SR OS supports key rollover via HelloRequest messages as detailed in RFC 5246, section 7.4.1.1. Some applications have a longer live time than other applications, in which case SR OS can use a timer that prompts the HelloRequest negotiation for the symmetric key rollover. This timer is configurable using CLI.

If an application does not support the HelloRequest message, the **no tls-re-negotiate-timer** command should be configured under the **config>system>security>tls** context. For example, the GRPC application does not support HelloRequest messages.

When **no tls-re-negotiate-timer** is configured, the HelloRequest message is not generated, and symmetric keys are not renegotiated.

8.7 Supported TLS Ciphers

As shown in [Figure 26](#), TLS negotiates the supported ciphers between the client and the server.

The client sends the supported cipher suites in the client Hello message and the server compares them with the server cipher list. The top protocol on both lists is chosen and returned from the server within the server Hello message.

The 7750 SR supports the following ciphers as a TLS client or TLS server:

- `tls-rsa-with-null-md5`
- `tls-rsa-with-null-sha`
- `tls-rsa-with-null-sha256`
- `tls-rsa-with3des-ede-cbc-sha`
- `tls-rsa-with-aes128-cbc-sha`
- `tls-rsa-with-aes256-cbc-sha`
- `tls-rsa-with-aes128-cbc-sha256`
- `tls-rsa-with-aes256-cbc-sha256`

8.8 SR OS Certificate Management

SR OS implements a centralized certificate management protocol that can be used by TLS and IPsec. Refer to the *7450 ESS and 7750 SR Multiservice Integrated Service Adapter Guide* for information about the configuration of the certificates and the corresponding protocols, such as OCSP, CMPv2, and CRL.

The main certificate configurations are:

- certificate configuration and management, configured using the **admin>certificate** commands
- PKI configuration (including creating a CA profile), configured using the **config>system>security>pki** commands

The two main configuration sub-trees for certificates are displayed below. See [Public Key Infrastructure \(PKI\) Commands](#) for more information.

```
CLI Syntax:  admin>certificate
              clear-ocsp-cache
              cmpv2
              crl-update
              display
              export
              gen-keypair
              gen-local-cert-req
              import
              reload

              config>system>security>pki
                [no] ca-profile
                certificate-display-format
                [no] certificate-expiration-warning
                [no] crl-expiration-warning
                [no] maximum-cert-chain-depth
```

8.8.1 Certificate Profile

The certificate profile is available for both the TLS server and the TLS client. The **cert-profile** command is configured for the server or client to transmit the provider certificate and its DS to the peer so that the peer can authenticate it via the **trust-anchor** and CA certificate.

Multiple provider certificates can be configured on SR OS; however, SR OS currently uses the smallest index as the active provider certificate, and will only send the certificate to the peer.

8.8.2 TLS Server Authentication of the Client Certificate CN Field

If the client provides a certificate upon request by the server, SR OS checks the certificate's common name (CN) field against local CN configurations. The CN is validated via the client IPv4/IPv6 address or FQDN.

If **cn-authentication** is not enabled, SR OS will not authenticate via the CN field and will only rely on certificate signature authentication.

8.8.3 CN Regexp Format

CN entries are configured by using the **config>system>security>pki>common-name-list** command. Entries should use regular expression (regexp), FQDN, or the IP address.

For information about regexp, refer to the *7450 ESS, 7750 SR, and 7950 XRS Basic System Configuration Guide*, "CLI Usage".

8.9 Operational Guidelines

8.9.1 Server Authentication Behavior

Following the Hello messages, the server sends its certificate in a certificate message if it is to be authenticated. If required, a ServerKeyExchange message may also be sent. Refer to RFC 5246, section 7.3, for more information about the authentication behavior on the LDAP server.

The **trust-anchor-profile** command determines whether or not the server must be authenticated by the client.

CLI Syntax: `config>system>security>tls`

```
client-tls-profile ldap create  
[no] trust-anchor-profile
```



Note: If the **trust-anchor-profile** is configured and the **ca-certificate** or **ca-profile** is missing from this **trust-anchor-profile**, the TLS connection will fail and an “unknown_ca” error will be generated, as per RFC 5246 section 7.2.2.

One of the following two configurations can be used to establish server connectivity.

- a. If **trust-anchor-profile** is configured under the TLS **client-tls-profile** context, the server must be authenticated via the **trust-anchor-profile** command before a trusted connection is established between the server and the client.
- b. If there is no **trust-anchor-profile** under the **client-tls-profile** context, the trusted connection can be established without server authentication. The RSA key of the certificate will be used for public key encryption, requiring basic certificate checks to validate the certificate. These basic checks are:
 - time validity—the certificate is checked to ensure that it is neither expired nor not yet valid
 - certificate type—the certificate is not a CA certificate
 - keyUsage extension—if present, this must contain a digital signature and key encryption
 - host verification—the IP address or DNS name of the server is looked up, if available (for LDAP, only the IP address is used), in the common name (cn) or subjectAltName extension. This is to verify that the certificate was issued to that server and not to another.

8.9.2 Client TLS Profile and Trust Anchor Behavior and Scale

SR OS allows the creation of client TLS profiles, which can be assigned to applications such as LDAP to encrypt the application layer.

The **client-tls-profiles** command is used for negotiating and authenticating the server. After the server is authenticated via the trust anchor profile (configured using the **trust-anchor-profile** command) of a client TLS profile, it negotiates the ciphers and authentication algorithms to be used for encryption of the data.

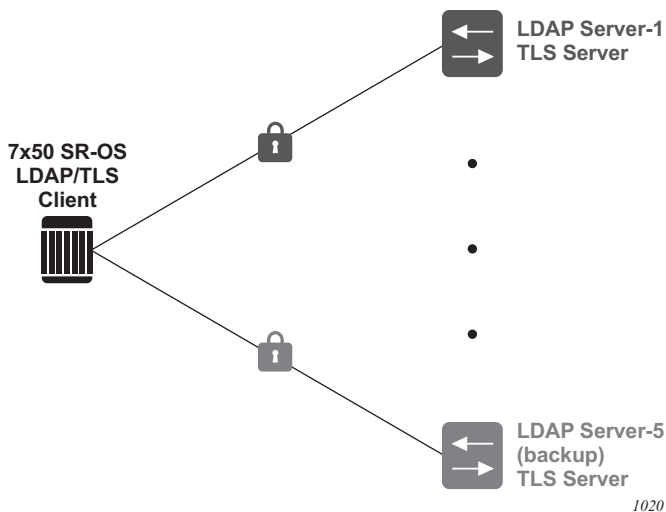
The client TLS profile must be assigned to an application for it to start encrypting. Up to 16 client TLS profiles can be configured. Because each of these client TLS profiles needs a trust anchor profile to authenticate the server, up to 16 trust anchor profiles can be configured. A trust anchor profile holds up to 8 trust anchors (configured using the **trust-anchor** command), which each hold a CA profile (**ca-profile**).

A CA profile is a container for installing CA certificates (**ca-certificates**). These CA certificates are used to authenticate the server certificate. When the client receives the server certificate, it reads through the trust anchor profile CA certificates and tries to authenticate the server certificate against each CA certificate. The first CA certificate that authenticates the server is used.

8.10 LDAP Redundancy and TLS

LDAP supports up to five redundant (backup) servers, as shown in [Figure 27](#) and the configuration examples below. Depending on the **timeout** and **retry** configurations, if an LDAP server is determined to be out of service or operationally down, SR OS will switch to the redundant servers. SR OS will select the LDAP server with the next largest configured server index.

Figure 27 LDAP and TLS Redundancy



Configuration of Server-1:

```
A*:SwSim14>config>system>security>ldap# info
public-key-authentication
server 1 create
address 1.1.1.1
```

```
        ldap-server "active-server"
        tls-profile "server-1-profile"

A*:SwSim14>config>system>security>tls# info
  client-tls-profile "server-1-profile" create
  cipher-list "to-active-server"
  trust-anchor-profile "server-1-ca"
  no shutdown
  exit
```

Configuration of Server-5 (backup):

```
A*:SwSim14>config>system>security>ldap# info
  public-key-authentication
  server 5 create
    address 5.5.5.1
    ldap-server "backup-server-5"
    tls-profile "server-5-profile"

A*:SwSim14>config>system>security>tls# info
  client-tls-profile "server-5-profile" create
  cipher-list "to-backup-server-5"
  trust-anchor-profile "server-5-ca"
  no shutdown
  exit
```

Each LDAP server can have its own TLS profile, each of which can have its own configuration of **trust-anchor** and **cipher-list**. For security reasons, the LDAP servers may be in different geographical areas and, as such, each will be assigned its own server certificate and trust anchor. The design is open to allow the user to mix and match all components.

8.11 Basic TLS Configuration

Basic TLS server configuration must have the following:

- a cipher list created using the **config>system>security>tls>server-cipher-list** command, and assigned to the TLS server profile using the **config>system>security>tls>server-tls-profile>cipher-list** command
- a certificate profile created using the **config>system>security>tls>cert-profile** command, and assigned to the TLS server profile using the **config>system>security>tls>server-tls-profile>cert-profile** command

Basic TLS client configuration must have a cipher list created using the **config>system>security>tls>client-cipher-list** command, and assigned to the TLS client profile using the **config>system>security>tls>client-tls-profile>cipher-list** command.

TLS imports the trust anchor certificate for (TLS) peer certificate authentication and public key retrieval.

The following displays the CLI syntax for TLS:

```

CLI Syntax:  config>system>security>tls
                cert-profile profile-name [create]
                no cert-profile profile-name
                client-cipher-list name [create]
                no client-cipher-list name
                client-tls-profile name [create]
                no client-tls-profile name
                server-cipher-list name [create]
                no server-cipher-list name
                server-tls-profile name [create]
                no server-tls-profile name
                trust-anchor-profile name [create]
                no trust-anchor-profile name
    
```

The following displays a TLS configuration example.

```

config>system>security>tls# info
-----
trust-anchor-profile "server-1-ca" create
trust-anchor "tls-server-1-ca"
exit
client-cipher-list "to-active-server" create
cipher 1 name tls-rsa-with-aes256-cbc-sha256
cipher 2 name tls-rsa-with-aes128-cbc-sha256
cipher 3 name tls-rsa-with-aes256-cbc-sha
exit
client-tls-profile "server-1-profile" create
    
```

```

cipher-list "to-active-server"
trust-anchor-profile "server-1-ca"
no shutdown
exit
-----

```

8.12 Common Configuration Tasks

The following sections are basic TLS configuration tasks that can be performed.

- [Configuring a Server TLS Profile](#)
- [Configuring a Client TLS Profile](#)
- [Configuring a TLS Client or TLS Server Certificate](#)
- [Configuring a TLS Trust Anchor](#)

8.12.1 Configuring a Server TLS Profile

The following displays the CLI syntax for a server TLS profile.

CLI Syntax:

```

config>system>security>tls
server-tls-profile name [create]
no server-tls-profile name
authenticate-client
trust-anchor-profile ca-profile-name
no trust-anchor-profile
cert-profile name
no cert-profile
cipher-list name
no cipher-list
[no] shutdown
tls-re-negotiate-timer [0 to 65000]
no tls-re-negotiate-timer

```

8.12.2 Configuring a Client TLS Profile

The following displays the CLI syntax for a client TLS profile, which also configures the server authentication behavior:

CLI Syntax:

```

config>system>security>tls
client-tls-profile name [create]

```

```
no client-tls-profile name
trust-anchor-profile name
no trust-anchor-profile
```

8.12.3 Configuring a TLS Client or TLS Server Certificate

The following displays the CLI syntax for TLS certificate management:

```
CLI Syntax:  config>system>security>tls
                cert-profile profile-name [create]
                no cert-profile profile-name
                entry entry-id [create]
                no entry entry-id
                   cert cert-filename
                   no cert
                   key key-filename
                   no key
                   [no] send-chain
                   [no] ca-profile name
                   [no] shutdown
                client-tls-profile name [create]
                no client-tls-profile name
                   cert-profile name
                   no cert-profile
                server-tls-profile name [create]
                no server-tls-profile name
                   cert-profile name
                   no cert-profile
```

8.12.4 Configuring a TLS Trust Anchor

The following displays the CLI syntax for a TLS trust anchor:

```
CLI Syntax:  config>system>security>pki
                [no] ca-profile
                certificate-display-format
                [no] certificate-expiration-warning hours
                [no] crl-expiration-warning
                [no] maximum-cert-chain-depth

                config>system>security>tls
                [no] trust-anchor-profile
                [no] client-tls-profile
                [no] cipher-list
```

```
[no] shutdown
[no] trust-anchor-profile-profile
```

The following displays a TLS trust anchor configuration example:

```
*B:SeGW-1>config>system>security>pki# info
-----
ca-profile "tls-server-1-ca" create
  cert-file "tls-1-Root-CERT"
  crl-file "tls-1-CRL-CERT"
  no shutdown
exit
-----
*A:SwSim8>config>system>security>tls# info
-----
trust-anchor-profile "server-1-ca" create
  trust-anchor "tls-server-1-ca"
exit
client-tls-profile "server-1-profile" create
  cipher-list "to-active-server"
  trust-anchor-profile "server-1-ca"
  no shutdown
exit
```

8.13 TLS Command Reference

8.13.1 Command Hierarchies

- [Security TLS Commands](#)
- [LDAP TLS Profile Commands](#)
- [Admin Commands](#)

8.13.1.1 Security TLS Commands

```

config
  -- system
    -- security
      -- tls
        -- cert-profile profile-name [create]
        -- no cert-profile profile-name
          -- entry entry-id [create]
          -- no entry entry-id
            -- cert cert-filename
            -- no cert
            -- key key-filename
            -- no key
            -- [no] send-chain
              -- [no] ca-profile name
              -- [no] shutdown
        -- client-cipher-list name [create]
        -- no client-cipher-list name
          -- cipher index name cipher-suite-code
          -- no cipher index
        -- client-tls-profile name [create]
        -- no client-tls-profile name
          -- cert-profile name
          -- no cert-profile
          -- cipher-list name
          -- no cipher-list
          -- [no] shutdown
          -- trust-anchor-profile name
          -- no trust-anchor-profile
        -- server-cipher-list name [create]
        -- no server-cipher-list name
          -- cipher index name cipher-suite-code
          -- no cipher index
        -- server-tls-profile name [create]
        -- no server-tls-profile name
          -- authenticate-client

```

- **trust-anchor-profile** *name*
- **no trust-anchor-profile**
- **cert-profile** *name*
- **no cert-profile**
- **cipher-list** *name*
- **no cipher-list**
- **[no] shutdown**
- **tls-re-negotiate-timer** *timer-min*
- **no tls-re-negotiate-timer**
- **trust-anchor-profile** *name* [**create**]
- **no trust-anchor-profile** *name*
- **[no] trust-anchor** *ca-profile-name*

8.13.1.2 LDAP TLS Profile Commands

- ```

config
 — system
 — security
 — ldap
 — server server-index [create]
 — no server server-index
 — tls-profile tls-profile-name
 — no tls-profile

```

### 8.13.1.3 Admin Commands

- ```

admin
  — certificate
    — reload type {cert | key | cert-key-pair} filename protocol protocol [key-file filename]
    
```

8.13.2 Command Descriptions

This section provides the CLI command descriptions. Topics include:

- [Security TLS Commands](#)
- [LDAP TLS Profile Commands](#)
- [Admin Commands](#)

8.13.2.1 Security TLS Commands

tls

Syntax	tls
Context	config>system>security
Description	This command configures TLS parameters.

cert-profile

Syntax	cert-profile <i>profile-name</i> [create] no cert-profile <i>profile-name</i>
Context	config>system>security>tls
Description	This command configures TLS certificate profile information. The certificate profile contains the certificates that are sent to the TLS peer (server or client) to authenticate itself. It is mandatory for the TLS server to send this information. The TLS client may optionally send this information upon request from the TLS server. The no form of the command deletes the specified TLS certificate profile.
Parameters	<i>profile-name</i> — Specifies the name of the TLS certificate profile, up to 32 characters maximum. create — Keyword used to create the TLS certificate profile.

entry

Syntax	entry <i>entry-id</i> [create] no entry <i>entry-id</i>
---------------	---

Context config>system>security>tls>cert-profile

Description This command configures an entry for the TLS certificate profile. A certificate profile may have up to eight entries. Currently, TLS uses the entry with the smallest ID number when responding to server requests.

The **no** form of the command deletes the specified entry.

Parameters *entry-id* — Specifies the identification number of the TLS certificate profile entry.

Values 1 to 8

create — Keyword used to create the TLS certificate profile entry.

cert

Syntax **cert** *cert-filename*
no cert

Context config>system>security>tls>cert-profile>entry

Description This command specifies the file name of an imported certificate for the **cert-profile** entry.

The **no** form of the command removes the certificate.

Default no cert

Parameters *cert-filename* — Specifies the file name of the TLS certificate, up to 95 characters maximum.

key

Syntax **key** *key-filename*
no key

Context config>system>security>tls>cert-profile>entry

Description This command specifies the file name of an imported key for the **cert-profile** entry.

The **no** form of the command removes the key.

Default no key

Parameters *key-filename* — Specifies the file name of the key.

send-chain

Syntax [**no**] **send-chain**

Context	config>system>security>tls>cert-profile>entry
Description	<p>This command enables the sending of certificate authority (CA) certificates, and enters the context to configure send-chain information.</p> <p>By default, the system only sends the TLS server certificate or TLS client certificate specified by the cert command. If CA certificates are to be sent using send-chain, they must be in the chain of certificates specified by the config>system>security>pki>ca-profile command. The specification of the send-chain is not necessary for a working TLS profile if the TLS peer has the CA certificate used to sign the server or client certificate in its own trust anchor.</p> <p>For example, given a TLS client running on SR OS, the ROOT CA certificate resides on the TLS server, but the subsequent SUB-CA certificate needed to complete the chain resides within SR OS. The send-chain command allows these SUB-CA certificates to be sent from SR OS to the peer to be authenticated using the ROOT CA certificate that resides on the peer.</p> <p>The no form of the command disables the send-chain.</p>
Default	no send-chain

ca-profile

Syntax	[no] ca-profile <i>name</i>
Context	config>system>security>tls>cert-profile>entry>send-chain
Description	<p>This command enables a certificate authority (CA) certificate in the specified CA profile to be sent to the peer. Up to seven configurations of this command are permitted in the same entry.</p> <p>The no form of the command disables the transmission of a CA certificate from the specified CA profile.</p>
Parameters	<i>name</i> — Specifies the name of the certificate authority profile, up to 32 characters maximum.

shutdown

Syntax	[no] shutdown
Context	config>system>security>tls>cert-profile
Description	<p>This command disables the certificate profile. When the certificate profile is disabled, it will not be sent to the TLS server.</p> <p>The no form of the command enables the certificate profile and allows it to be sent to the TLS server.</p>
Default	shutdown

client-cipher-list

- Syntax** `client-cipher-list name [create]`
`no client-cipher-list name`
- Context** config>system>security>tls
- Description** This command creates a cipher list that the client sends to the server in the client Hello message. It is a list of ciphers that are supported and preferred by the SR OS to be used in the TLS session. The server matches this list against the server cipher list. The most preferred cipher found in both lists is chosen.
- Parameters** *name* — Specifies the name of the client cipher list, up to 32 characters maximum.
create — Keyword used to create the client cipher list.

cipher

- Syntax** `cipher index name cipher-suite-code`
`no cipher index`
- Context** config>system>security>tls>client-cipher-list
config>system>security>tls>server-cipher-list
- Description** This command configures the cipher suite to be negotiated by the server and client.
- Parameters** *index* — Specifies the index number. The index number provides the location of the cipher in the negotiation list, with the lower index numbers being higher in the negotiation list and the higher index numbers being at the bottom of the list.
- Values** 1 to 255
- cipher-suite-code* — Specifies the cipher suite code.
- Values** tls-rsa-with-null-md5
tls-rsa-with-null-sha
tls-rsa-with-null-sha256
tls-rsa-with-3des-ede-cbc-sha
tls-rsa-with-aes128-cbc-sha
tls-rsa-with-aes256-cbc-sha
tls-rsa-with-aes128-cbc-sha256
tls-rsa-with-aes256-cbc-sha256

client-tls-profile

- Syntax** `client-tls-profile name [create]`
`no client-tls-profile name`

Context	config>system>security>tls
Description	This command configures the TLS client profile to be assigned to applications for encryption.
Parameters	<i>name</i> — Specifies the name of the client TLS profile, up to 32 characters maximum. create — Keyword used to create the client TLS profile.

cipher-list

Syntax	cipher-list <i>name</i> no cipher-list
Context	config>system>security>tls>client-tls-profile
Description	This command assigns the cipher list to be used by the TLS client profile for negotiation in the client Hello message.
Parameters	<i>name</i> — Specifies the name of the cipher list.

shutdown

Syntax	[no] shutdown
Context	config>system>security>tls>client-tls-profile config>system>security>tls>server-tls-profile
Description	This command administratively enables or disables the TLS profile. If the TLS profile is shut down, the TLS operational status will be down. Therefore, if the TLS profile is shut down, any application using TLS should not attempt to send any PDUs.

trust-anchor-profile

Syntax	trust-anchor-profile <i>name</i> no trust-anchor-profile
Context	config>system>security>tls>client-tls-profile config>system>security>tls>server-tls-profile>authenticate-client
Description	This command assigns the trust anchor used by this TLS profile to authenticate the server or client. The no form of the command removes the configured trust anchor profile.
Parameters	<i>name</i> — Specifies the name of the trust anchor profile.

server-cipher-list

- Syntax** **server-cipher-list** *name* [**create**]
no server-cipher-list *name*
- Context** config>system>security>tls
- Description** This command creates the cipher list that is compared against cipher lists sent by the client to the server in the client hello message. The list contains all ciphers that are supported and desired by SR OS for use in the TLS session. The first common cipher found in both the server and client cipher lists will be chosen. As such, the most desired ciphers should be added at the top of the list.
- The **no** form of the command removes the cipher list.
- Parameters** *name* — Specifies the name of the server cipher list, up to 32 characters maximum.
create — Keyword used to create the server cipher list.

server-tls-profile

- Syntax** **server-tls-profile** *name* [**create**]
no server-tls-profile *name*
- Context** config>system>security>tls
- Description** This command creates a TLS server profile. This profile can be used by applications that support TLS for encryption. The applications should not send any PDUs until the TLS handshake has been successful.
- The **no** form of the command removes the TLS server profile.
- Parameters** *name* — Specifies the name of the TLS server profile, up to 32 characters maximum.
create — Keyword used to create the TLS server profile.

authenticate-client

- Syntax** **authenticate-client**
- Context** config>system>security>tls>server-tls-profile
- Description** This command enters the context to configure client authentication parameters.

cert-profile

- Syntax** **cert-profile** *name*

no cert-profile

- Context** config>system>security>tls>client-tls-profile
- Description** This command assigns a TLS certificate profile to be used by the TLS client profile. This certificate is sent to the server for authentication of the client and public key.
- The **no** form of the command removes the TLS certificate profile assignment.
- Parameters** *name* — Specifies the name of the TLS certificate profile, up to 32 characters maximum.

cert-profile

- Syntax** **cert-profile** *name*
no cert-profile
- Context** config>system>security>tls>server-tls-profile
- Description** This command assigns a TLS certificate profile to be used by the TLS server profile. This certificate is sent to the client for authentication of the server and public key.
- The **no** form of the command removes the TLS certificate profile assignment.
- Parameters** *name* — Specifies the name of the TLS certificate profile, up to 32 characters maximum.

cipher-list

- Syntax** **cipher-list** *name*
no cipher-list
- Context** config>system>security>tls>server-tls-profile
- Description** This command assigns a cipher list to be used by the TLS server profile. This cipher list is used to find matching ciphers with the cipher list that is received from the client.
- The **no** form of the command removes the cipher list.
- Parameters** *name* — Specifies the name of the cipher list, up to 32 characters maximum.

tls-re-negotiate-timer

- Syntax** **tls-re-negotiate-timer** *timer-min*
no tls-re-negotiate-timer
- Context** config>system>security>tls>server-tls-profile

- Description** This command configures the timed interval after which the server is triggered to send a Hello request message to all clients and force a renegotiation of the symmetric encryption key. When an interval of 0 is configured, the server will never send a hello request message.
- Default** `tls-re-negotiate-timer 0`
- Parameters** *timer-min* — Specifies the interval, in minutes, after which the server is triggered to send a Hello request message.
Values 0 to 65000

trust-anchor-profile

- Syntax** `trust-anchor-profile name [create]`
`no trust-anchor-profile name`
- Context** `config>system>security>tls`
- Description** This command configures a trust anchor profile to be used in the TLS profile. The trust anchor is used for authentication of the server certificate.
- Parameters** *name* — Specifies the name of the trust anchor profile, up to 32 characters maximum.
create — Keyword used to create the trust anchor profile.

trust-anchor

- Syntax** `[no] trust-anchor ca-profile-name`
- Context** `config>system>security>tls>trust-anchor-profile`
- Description** This command configures a trust anchor with a CA profile used by the TLS profile. Up to eight CA profiles can be configured under the trust anchor. TLS will read the CA profiles one by one to try to authenticate the server certificate.
- Parameters** *ca-profile-name* — Specifies the name of the TLS trust anchor, up to 32 characters maximum.

8.13.2.2 LDAP TLS Profile Commands

server

- Syntax** `server server-index [create]`
`no server server-index`

Context	config>system>security>ldap
Description	This command adds or removes an LDAP server.
Parameters	<i>server-index</i> — Specifies the server index. Values 1 to 5 create — Keyword used to create the server index.

tls-profile

Syntax	tls-profile <i>tls-profile-name</i> no tls-profile
Context	config>system>security>ldap>server
Description	This command assigns a TLS profile to the LDAP application. When a TLS profile is assigned, the LDAP application will send encrypted PDUs from the client to the LDAP server. If TLS is operationally down, the LDAP application should not send any PDUs.
Parameters	<i>tls-profile-name</i> — Specifies the name of the TLS client transport profile.

8.13.2.3 Admin Commands

reload

Syntax	reload type { cert key cert-key-pair } <i>filename protocol protocol</i> [key-file <i>filename</i>]
Context	admin>certificate
Description	This command manually reloads the certificate or key cache.
Parameters	type — Specifies what item will be reloaded. cert — Specifies that a certificate cache will be reloaded. key — Specifies that a key cache will be reloaded. cert-key-pair — Specifies that a paired certificate and key cache will be reloaded. <i>filename</i> — Up to 95 characters. <i>protocol</i> — Specifies which protocol the certificate will be reloaded for. Values ipsec, tls

8.14 TLS Show Command Reference

8.14.1 Command Hierarchies

- [Show Commands](#)

8.14.1.1 Show Commands

```
show
  -- system
    -- security
      -- tls
        -- cert-profile name association
        -- cert-profile [name]
        -- cert-profile name entry entry
        -- client-tls-profile [client-tls-profile]
        -- client-tls-profile client-tls-profile association
        -- server-tls-profile [server-tls-profile]
        -- server-tls-profile server-tls-profile association
        -- trust-anchor-profile [trust-anchor-profile]
        -- trust-anchor-profile trust-anchor-profile association
```

8.14.2 Command Descriptions

- [Show Commands](#)

8.14.2.1 Show Commands

The command outputs in the following section are examples only; actual displays may differ depending on supported functionality and user configuration.

tls

Syntax	tls
Context	show>system>security
Description	This command enables the context to display TLS-related information.

cert-profile

Syntax	cert-profile [<i>name</i>] cert-profile <i>name</i> association cert-profile <i>name</i> entry <i>entry</i>
Context	show>system>security>tls
Description	This command displays information about server and client profiles that are using this certificate profile.
Parameters	<i>entry</i> — Specifies a certificate profile entry number for which to display information. Values 1 to 8 <i>name</i> — Specifies the name of a certificate profile for which to display information.

client-tls-profile

Syntax	client-tls-profile [<i>client-tls-profile</i>] client-tls-profile <i>client-tls-profile</i> association
Context	show>system>security>tls
Description	This command displays TLS client profile information

Parameters *client-tls-profile* — Specifies the client TLS profile, up to 32 characters maximum.

Output The following output is an example of TLS client profile information.

Sample Output

```
*A:Dut-C> show system security tls client-tls-profile
=====
Client Profile Information
=====
Name                               AdminState   OperState
-----
ctp                                 up           up
ctp-alt1                            up           up
ctp-alt2                            up           up
=====

*A:Dut-C> show system security tls client-tls-profile "ctp"
=====
Client Profile Entry "ctp"
=====
Cipher List Name                   : cl_all
Trust Anchor Profile Name          : tap
=====
```

server-tls-profile

Syntax **server-tls-profile** [*server-tls-profile*]
server-tls-profile *server-tls-profile* **association**

Context show>system>security>tls

Description This command displays TLS server profile information.

Parameters *server-tls-profile* — Specifies the name of a TLS server profile for which to display information, up to 32 characters maximum.

trust-anchor-profile

Syntax **trust-anchor-profile** [*trust-anchor-profile*]
trust-anchor-profile *trust-anchor-profile* **association**

Context show>system>security>tls

Description This command displays information about server and client profiles that are using the specified TLS trust anchor profile.

Parameters *trust-anchor-profile* — Specifies the trust anchor profile, up to 32 characters maximum.

Output The following output is an example of trust anchor profile information.

Sample Output

```
*A:Dut-C> show system security tls trust-anchor-profile
=====
Trust Anchor Profile Information
=====
Name                                     CA Profiles Down
-----
tap                                     0
tap-alt1                                0
tap-alt2                                0
tap-empty                                0
=====

*A:Dut-C> show system security tls trust-anchor-profile "tap"
=====
CA-profile List for Trust Anchor "tap"
=====
CA Profile Name                         AdminState      OperState
-----
chainA_11                               up              up
revChainA_11                             up              up
=====
*A:Dut-C>show>tls#
```

9 Facility Alarms

9.1 In This Chapter

This chapter provides information about Facility Alarms.

Topics in this chapter include:

- [Facility Alarms Overview](#)
- [Facility Alarms vs. Log Events](#)
- [Facility Alarm Severities and Alarm LED Behavior](#)
- [Facility Alarm Hierarchy](#)
- [Facility Alarm List](#)

9.2 Facility Alarms Overview

Facility Alarms provide a useful tool for operators to easily track and display the basic status of their equipment facilities. Facility Alarm support is intended to cover a focused subset of router states that are likely to indicate service impacts (or imminent service impacts) related to the overall state of hardware assemblies (cards, fans, links, and so on).

In the CLI, for brevity, the keyword or command `alarm` is used for commands related to Facility Alarms. This chapter may occasionally use the term “alarm” as a short form for “facility alarm”.

The CLI display for `show` routines allows the system operator to easily identify current facility alarm conditions and recently cleared facility alarms without searching event logs or monitoring various card and port show commands to determine the health of basic equipment in the system such as cards and ports.

The SR OS alarm model is based on RFC 3877, *Alarm Management Information Base (MIB)*, (which evolved from the IETF DISMAN drafts).

9.3 Facility Alarms vs. Log Events

Facility Alarms are different than log events. Facility Alarms have a state (at least two states: active and clear) and a duration, and can be modeled with state transition events (raised, cleared). A log event occurs when the state of some object in the system changes. Log events notify the operator of a state change (for example, a port going down, an IGP peering session coming up, and so on). Facility alarms show the list of hardware objects that are currently in a bad state. Facility alarms can be examined at any time by an operator, whereas log events can be sent by a router asynchronously when they occur (for example, as an SNMP notification or trap, or a syslog event).

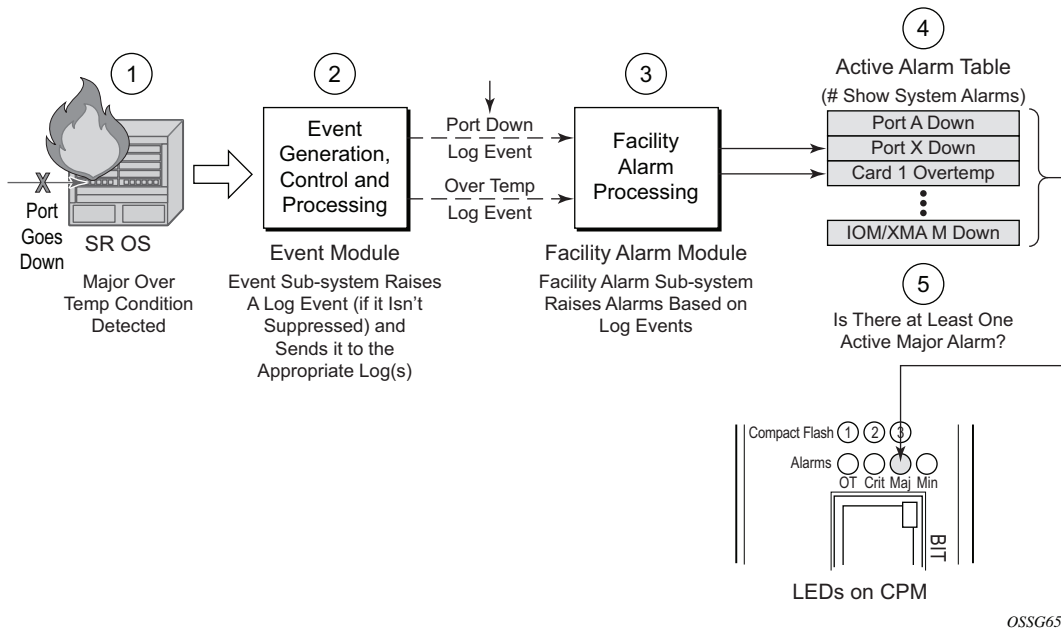
While log events provide notifications about a large number of different types of state changes in SR OS, facility alarms are intended to cover a focused subset of router states that are likely to indicate service impacts (or imminent service impacts) related to the overall state of hardware assemblies (cards, fans, links, and so on).

The facility alarm module processes log events in order to generate the raised and cleared state for the facility alarms. If a raising log event is suppressed under event-control, then the associated facility alarm will not be raised. If a clearing log event is suppressed under event-control, then it is still processed for the purpose of clearing the associated facility alarm. If a log event is a raising event for a Facility Alarm, and the associated Facility Alarm is raised, then changing the log event to **suppress** will clear the associated Facility Alarm.

Log event filtering, throttling and discarding of log events during overload do not affect facility alarm processing. In all cases, non-suppressed log events are processed by the facility alarm module before they are discarded.

[Figure 28](#) illustrates the relationship of log events, facility alarms and the LEDs.

Figure 28 Log Events, Facility Alarms and LEDs



OSSG651

Facility Alarms are different and independent functionality from other uses of the term *alarm* in SR-OS such as:

- Log events that use the term **alarm** (tmnxEqPortSonetAlarm)
- **configure card fp hi-bw-mcast-src [alarm]**
- **configure mcast-management multicast-info-policy bundle channel source-override video analyzer alarms**
- **configure port ethernet report-alarm**
- **configure system thresholds no memory-use-alarm**
- **configure system thresholds rmon no alarm**
- **configure system security cpu-protection policy alarm**

9.4 Facility Alarm Severities and Alarm LED Behavior

The Alarm LEDs on the CPM/CCM reflects the current status of the Facility Alarms:

- The Critical Alarm LED is lit if there is 1 or more active Critical Facility Alarms
- Similarly with the Major and Minor alarm LEDs

- The OT Alarm LED is not controlled by the Facility Alarm module

The supported alarm severities are as follows:

- Critical (with an associated LED on the CPM/CCM)
- Major (with an associated LED on the CPM/CCM)
- Minor (with an associated LED on the CPM/CCM)
- Warning (no LED)

Facility alarms inherit their severity from the raising log event.

A raising log event for a facility alarm configured with a severity of *indeterminate* or *cleared* will result in the facility alarm not being raised. But, a clearing log event is processed in order to clear facility alarms, regardless of the severity of the clearing log event.

Changing the severity of a raising log event only affects subsequent occurrences of that log event and facility alarms. Facility alarms that are already raised when their raising log event severity is changed maintain their original severity.

9.5 Facility Alarm Hierarchy

Facility Alarms for *children* objects is not raised for failure of a *parent* object. For example, when an MDA or XMA fails (or is *shutdown*) there is not a set of port facility alarms raised.

When a parent facility alarm is cleared, children facility alarms that are still in occurrence on the node appears in the active facility alarms list. For example, when a port fails there is a port facility alarm, but if the MDA or XMA is later shutdown the port alarm is cleared (and a card alarm will be active for the MDA or XMA). If the MDA or XMA comes back into service, and the port is still down, then a port alarm becomes active once again.

The supported facility alarm hierarchy is as follows (parent objects that are *down* cause alarms in all children to be masked):

- CPM -> Compact Flash
- CCM -> Compact Flash
- IOM/IMM -> MDA -> Port -> Channel
- XCM -> XMA -> Port
- MCM -> MDA -> Port -> Channel



Note: A *masked* facility alarm is not the same as a *cleared* facility alarm. The cleared facility alarm queue does not display entries for previously raised facility alarms that are currently masked. If the masking event goes away, then the previously raised facility alarms will once again be visible in the active facility alarm queue.

9.6 Facility Alarm List

Table 95 and Table 96 show the supported Facility Alarms.

Table 95 Facility Alarm, Facility Alarm Name/Raising Log Event, Sample Details String and Clearing Log Event

Facility Alarm *1	Facility Alarm Name/Raising Log Event	Sample Details String	Clearing Log Event
7-2001-1	tmnxEqCardFailure	Class MDA Module: failed, reason: Mda 1 failed startup tests	tmnxChassisNotificationClear
7-2003-1	tmnxEqCardRemoved	Class CPM Module: removed	tmnxEqCardInserted
7-2004-1	tmnxEqWrongCard	Class IOM Module: wrong type inserted	tmnxChassisNotificationClear
7-2005-1	tmnxEnvTempTooHigh	Chassis 1: temperature too high	tmnxChassisNotificationClear
7-2006-1	tmnxEqFanFailure	Fan 2 failed	tmnxChassisNotificationClear
7-2007-1	tmnxEqPowerSupplyFailureOvt	Power supply 2 over temperature	tmnxChassisNotificationClear
7-2008-1	tmnxEqPowerSupplyFailureAc	Power supply 1 AC failure	tmnxChassisNotificationClear
7-2009-1	tmnxEqPowerSupplyFailureDc	Power supply 2 DC failure	tmnxChassisNotificationClear
7-2011-1	tmnxEqPowerSupplyRemoved	Power supply 1, power lost	tmnxEqPowerSupplyInserted
7-2017-1	tmnxEqSynclftimingHoldover	Synchronous Timing interface in holdover state	tmnxEqSynclftimingHoldoverClear
7-2019-1	tmnxEqSynclftimingRef1Alarm with attribute tmnxSynclftimingNotifyAlarm == 'los(1)'	Synchronous Timing interface, alarm los on reference 1	tmnxEqSynclftimingRef1AlarmClear

Table 95 Facility Alarm, Facility Alarm Name/Raising Log Event, Sample Details String and Clearing Log Event (Continued)

Facility Alarm *1	Facility Alarm Name/Raising Log Event	Sample Details String	Clearing Log Event
7-2019-2	tmnxEqSynclftimingRef1Alarm with attribute tmnxSynclftimingNotifyAlarm == 'oof(2)'	Synchronous Timing interface, alarm oof on reference 1	same as 7-2019-1
7-2019-3	tmnxEqSynclftimingRef1Alarm with attribute tmnxSynclftimingNotifyAlarm == 'oopir(3)'	Synchronous Timing interface, alarm oopir on reference 1	same as 7-2019-1
7-2021-x	same as 7-2019-x but for ref2	same as 7-2019-x but for ref2	same as 7-2019-x but for ref2
7-2030-x	same as 7-2019-x but for the BITS input	same as 7-2019-x but for the BITS input	same as 7-2019-x but for the BITS input
7-2033-1	tmnxChassisUpgradeInProgress	Class CPM Module: software upgrade in progress	tmnxChassisUpgradeComplete
7-2050-1	tmnxEqPowerSupplyFailureInput	Power supply 1 input failure	tmnxChassisNotificationClear
7-2051-1	tmnxEqPowerSupplyFailureOutput	Power supply 1 output failure	tmnxChassisNotificationClear
7-2073-x	same as 7-2019-x but for the BITS2 input	same as 7-2019-x but for the BITS2 input	same as 7-2019-x but for the BITS2 input
7-2092-1	tmnxEqPowerCapacityExceeded	The system has reached maximum power capacity <x> watts	tmnxEqPowerCapacityExceededClear
7-2094-1	tmnxEqPowerLostCapacity	The system can no longer support configured devices. Power capacity dropped to <x> watts	tmnxEqPowerLostCapacityClear
7-2096-1	tmnxEqPowerOverloadState	The system has reached critical power capacity. Increase available power now	tmnxEqPowerOverloadStateClear
7-4001-1	tmnxInterChassisCommsDown	Control communications disrupted between the Active CPM and the chassis	tmnxInterChassisCommsUp
7-4003-1	tmnxCpmlcPortDown	CPM Interconnect Port is not operational. Error code = invalid-connection	tmnxCpmlcPortUp

Table 95 Facility Alarm, Facility Alarm Name/Raising Log Event, Sample Details String and Clearing Log Event (Continued)

Facility Alarm *1	Facility Alarm Name/Raising Log Event	Sample Details String	Clearing Log Event
7-4007-1	tmnxCpmANoLocalIcPort	CPM A can not reach the chassis using its local CPM interconnect ports	tmnxCpmALocalIcPort Avail
7-4008-1	tmnxCpmBNoLocalIcPort	CPM B can not reach the chassis using its local CPM interconnect ports	tmnxCpmBLocalIcPort Avail
7-4017-1	tmnxSfmlcPortDown	SFM interconnect Port is not operational. Error code = invalid-connection to Fabric 10 IcPort 2	tmnxSfmlcPortUp
59-2004-1	linkDown	Interface intf-towards-node-B22 is not operational	linkUp
64-2091-1	tmnxSysLicenseInvalid	Error - <reason> record. <hw> will reboot the chassis <timeRemaining>	None
64-2092-1	tmnxSysLicenseExpiresSoon	The license installed on <hw> expires <timeRemaining>	None

Table 96 Facility Alarm Name/Raising Log Event, Cause, Effect and Recovery

Facility Alarm *1	Facility Alarm Name/Raising Log Event	Cause	Effect	Recovery
7-2001-1	tmnxEqCardFailure	Generated when one of the cards in a chassis has failed. The card type may be IOM (or XCM), MDA (or XMA), SFM, CCM, CPM, Compact Flash, etc. The reason is indicated in the details of the log event or alarm, and also available in the tmnxChassisNotifyCardFailureReason attribute included in the SNMP notification.	The effect is dependent on the card that has failed. IOM (or XCM) or MDA (or XMA) failure will cause a loss of service for all services running on that card. A fabric failure can impact traffic to/from all cards. 7750 SR/7450 ESS — If the IOM/IMM fails then the two associated MDAs for the slot will also go down. 7950 XRS — If one out of two XMA fails in a XCM slot then the XCM will remain up. If only one remaining operational XMA within a XCM slot fails, then the XCM will go into a booting operational state.	Before taking any recovery steps collect a tech-support file, then try resetting (clear) the card. If that doesn't work then try removing and then re-inserting the card. If that doesn't work then replace the card.
7-2003-1	tmnxEqCardRemoved	Generated when a card is removed from the chassis. The card type may be IOM (or XCM), MDA (or XMA), SFM, CCM, CPM, Compact Flash, etc.	The effect is dependent on the card that has been removed. IOM (or XCM) or MDA (or XMA) removal will cause a loss of service for all services running on that card. A fabric removal can impact traffic to/from all cards.	Before taking any recovery steps collect a tech-support file, then try re-inserting the card. If that doesn't work then replace the card.
7-2004-1	tmnxEqWrongCard	Generated when the wrong type of card is inserted into a slot of the chassis. Even though a card may be physically supported by the slot, it may have been administratively configured to allow only certain card types in a particular slot location. The card type may be IOM (or XCM), MDA (or XMA), SFM, CCM, CPM, Compact Flash, etc.	The effect is dependent on the card that has been incorrectly inserted. Incorrect IOM (or XCM) or MDA (or XMA) insertion will cause a loss of service for all services running on that card.	Insert the correct card into the correct slot, and ensure the slot is configured for the correct type of card.

Table 96 Facility Alarm Name/Raising Log Event, Cause, Effect and Recovery (Continued)

Facility Alarm *1	Facility Alarm Name/Raising Log Event	Cause	Effect	Recovery
7-2005-1	tmnxEnvTempTooHigh	Generated when the temperature sensor reading on an equipment object is greater than its configured threshold.	This could be causing intermittent errors and could also cause permanent damage to components.	Remove or power down the affected cards, or improve the cooling to the node. More powerful fan trays may also be required.
7-2006-1	tmnxEqFanFailure	Generated when one of the fans in a fan tray has failed.	This could be cause temperature to rise and resulting intermittent errors and could also cause permanent damage to components.	Replace the fan tray immediately, improve the cooling to the node, or reduce the heat being generated in the node by removing cards or powering down the node.
7-2007-1	tmnxEqPowerSupplyFailureOvt	Generated when the temperature sensor reading on a power supply module is greater than its configured threshold.	This could be causing intermittent errors and could also cause permanent damage to components.	Remove or power down the affected power supply module or improve the cooling to the node. More powerful fan trays may also be required. The power supply itself may be faulty so replacement may be necessary.
7-2008-1	tmnxEqPowerSupplyFailureAc	Generated when an AC failure is detected on a power supply.	Reduced power can cause intermittent errors and could also cause permanent damage to components.	First try re-inserting the power supply. If that doesn't work, then replace the power supply.
7-2009-1	tmnxEqPowerSupplyFailureDc	Generated when an DC failure is detected on a power supply.	Reduced power can cause intermittent errors and could also cause permanent damage to components.	First try re-inserting the power supply. If that doesn't work, then replace the power supply.

Table 96 Facility Alarm Name/Raising Log Event, Cause, Effect and Recovery (Continued)

Facility Alarm *1	Facility Alarm Name/Raising Log Event	Cause	Effect	Recovery
7-2011-1	tmnxEqPowerSupplyRemoved	Generated when: <ul style="list-style-type: none"> • one of the power supplies is removed from the chassis • low input voltage is detected. The operating voltage range for the 7750 SR-7 and 7750 SR-12 is -40 to -72 VDC. The alarm is raised if the system detects that the voltage of the power supply has dropped to -42.5 VDC. 	Reduced power can cause intermittent errors and could also cause permanent damage to components.	Re-insert the power supply or raise the input voltage to above -42.5 VDC
7-2017-1	tmnxEqSyncIfTimingHoldover	Generated when the synchronous equipment timing subsystem transitions into a holdover state.	Any node-timed ports will have very slow frequency drift limited by the central clock oscillator stability. The oscillator meets the holdover requirements of a Stratum 3 and G.813 Option 1 clock.	Address issues with the central clock input references.
7-2019-1	tmnxEqSyncIfTimingRef1Alarm with attribute tmnxSyncIfTimingNotifyAlarm == 'los(1)'	Generated when an alarm condition on the first timing reference is detected. The type of alarm (los, oof, etc) is indicated in the details of the log event or alarm, and is also available in the tmnxSyncIfTimingNotifyAlarm attribute included in the SNMP notification. The SNMP notification will have the same indices as those of the tmnxCpmCardTable.	Timing reference 1 cannot be used as a source of timing into the central clock.	Address issues with the signal associated with timing reference 1.

Table 96 Facility Alarm Name/Raising Log Event, Cause, Effect and Recovery (Continued)

Facility Alarm *1	Facility Alarm Name/Raising Log Event	Cause	Effect	Recovery
7-2019-2	tmnxEqSynclfTimingRef1Alarm with attribute tmnxSynclfTimingNotifyAlarm == 'oof(2)'	same as 7-2019-1	same as 7-2019-1	same as 7-2019-1
7-2019-3	tmnxEqSynclfTimingRef1Alarm with attribute tmnxSynclfTimingNotifyAlarm == 'oopir(3)'	same as 7-2019-1.	same as 7-2019-1.	same as 7-2019-1.
7-2021-x	same as 7-2019-x but for ref2	same as 7-2019-x but for the second timing reference	same as 7-2019-x but for the second timing reference	same as 7-2019-x but for the second timing reference
7-2030-x	same as 7-2019-x but for the BITS input	same as 7-2019-x but for the BITS timing reference	same as 7-2019-x but for the BITS timing reference	same as 7-2019-x but for the BITS timing reference
7-2033-1	tmnxChassisUpgradeInProgress	The tmnxChassisUpgradeInProgress notification is generated only after a CPM switchover occurs and the new active CPM is running new software, while the IOMs or XCMs are still running old software. This is the start of the upgrade process. The tmnxChassisUpgradeInProgress notification will continue to be generated every 30 minutes while at least one IOM or XCM is still running older software.	A software mismatch between the CPM and IOM or XCM is generally fine for a short duration (during an upgrade) but may not allow for correct long term operation.	Complete the upgrade of all IOMs or XCMs.

Table 96 Facility Alarm Name/Raising Log Event, Cause, Effect and Recovery (Continued)

Facility Alarm *1	Facility Alarm Name/Raising Log Event	Cause	Effect	Recovery
7-2050-1	tmnxEqPowerSupplyFailureInput	Generated when an input failure is detected on a power supply.	Reduced power can cause intermittent errors and could also cause permanent damage to components.	First try re-inserting the power supply. If that doesn't work, then replace the power supply.
7-2051-1	tmnxEqPowerSupplyFailureOutput	Generated when an output failure is detected on a power supply.	Reduced power can cause intermittent errors and could also cause permanent damage to components.	First try re-inserting the power supply. If that doesn't work, then replace the power supply.
7-2073-x	same as 7-2019-x but for the BITS2 input	same as 7-2019-x but for the BITS 2 timing reference	same as 7-2019-x but for the BITS 2 timing reference	same as 7-2019-x but for the BITS 2 timing reference
7-2092-1	tmnxEqPowerCapacityExceeded	Generated when a device needs power to boot, but there is not enough power capacity to support the device.	A non-powered device will not boot until the power capacity is increased to support the device.	Add a new power supply to the system, or change the faulty power supply with a working one.
7-2094-1	tmnxEqPowerLostCapacity	Generated when a power supply fails or is removed which puts the system in an overloaded situation.	Devices are powered off in order of lowest power priority until the available power capacity can support the powered devices.	Add a new power supply to the system, or change the faulty power supply with a working one.
7-2096-1	tmnxEqPowerOverloadState	Generated when the overloaded power capacity can not support the power requirements and there are no further devices that can be powered off.	The system runs a risk of experiencing brownouts while the available power capacity does not meet the required power consumption.	Add power capacity or manually shutdown devices until the power capacity meets the power needs.
7-4001-1	tmnxInterChassisCommsDown	The tmnxInterChassisComms Down alarm is generated when the active CPM cannot reach the far-end chassis.	The resources on the far-end chassis are not available. This event for the far-end chassis means that the CPM, SFM, and XCM cards in the far-end chassis will reboot and remain operationally down until communications are re-established.	Ensure that all CPM interconnect ports in the system are properly cabled together with working cables.

Table 96 Facility Alarm Name/Raising Log Event, Cause, Effect and Recovery (Continued)

Facility Alarm *1	Facility Alarm Name/Raising Log Event	Cause	Effect	Recovery
7-4003-1	tmnxCpmlcPort Down	The tmnxCpmlcPortDown alarm is generated when the CPM interconnect port is not operational. The reason may be a cable connected incorrectly, a disconnected cable, a faulty cable, or a misbehaving CPM interconnect port or card.	At least one of the control plane paths used for inter-chassis CPM communication is not operational. Other paths may be available.	A manual verification and testing of each CPM interconnect port is required to ensure fully functional operation. Physical replacement of cabling may be required.
7-4007-1	tmnxCpmANoLocallcPort	The tmnxCpmANoLocallcPort alarm is generated when the CPM cannot reach the other chassis using its local CPM interconnect ports.	<p>Another control communications path may still be available between the CPM and the other chassis via the mate CPM in the same chassis. If that alternative path is not available then complete disruption of control communications to the other chassis will occur and the tmnxInterChassisCommsDown alarm is raised.</p> <p>A tmnxCpmANoLocallcPort alarm on the active CPM indicates that a further failure of the local CPM interconnect ports on the standby CPM will cause complete disruption of control communications to the other chassis and the tmnxInterChassisCommsDown alarm is raised.</p> <p>A tmnxCpmANoLocallcPort alarm on the standby CPM indicates that a CPM switchover may cause temporary disruption of control communications to the other chassis while the rebooting CPM comes back into service.</p>	Ensure that all CPM interconnect ports in the system are properly cabled together with working cables.
7-4008-1	tmnxCpmBNoLocallcPort	Same as 7-4007-1.	Same as 7-4007-1.	Same as 7-4007-1.

Table 96 Facility Alarm Name/Raising Log Event, Cause, Effect and Recovery (Continued)

Facility Alarm *1	Facility Alarm Name/Raising Log Event	Cause	Effect	Recovery
7-4009-1	tmnxCpmALocalIcPortAvail	The tmnxCpmALocalIcPortAvail notification is generated when the CPM re-establishes communication with the other chassis using its local CPM interconnect ports.	A new control communications path is now available between the CPM_A and the other chassis,	
7-4010-1	tmnxCpmBLocalIcPortAvail	Same as 7-4009-1.	Same as 7-4009-1.	Same as 7-4009-1.
7-4017-1	tmnxSfmlcPortDown	The tmnxSfmlcPortDown alarm is generated when the SFM interconnect port is not operational. The reason may be a cable connected incorrectly, a disconnected cable, a faulty cable, or a misbehaving SFM interconnect port or SFM card.	This port can no longer be used as part of the user plane fabric between chassis. Other fabric paths may be available resulting in no loss of capacity.	A manual verification and testing of each SFM interconnect port is required to ensure fully functional operation. Physical replacement of cabling may be required.
59-2004-1	linkDown	A linkDown trap signifies that the SNMP entity, acting in an agent role, has detected that the ifOperStatus object for one of its communication links is about to enter the down state from some other state (but not from the notPresent state).	The indicated interface is taken down.	If the ifAdminStatus is down then the interface state is deliberate and there is no recovery. If the ifAdminStatus is up then try to determine that cause of the interface going down: cable cut, distal end went down, etc.
64-2091-1	tmnxSysLicenseInvalid	Generated when the license becomes invalid for the reason specified in the log event/alarm.	The system will reboot at the end of the time remaining.	Configure a valid license file location and file name.

Table 96 Facility Alarm Name/Raising Log Event, Cause, Effect and Recovery (Continued)

Facility Alarm *1	Facility Alarm Name/Raising Log Event	Cause	Effect	Recovery
64-2092-1	tmnxSysLicenseExpiresSoon	Generated when the license is due to expire soon.	The system will reboot at the end of the time remaining.	Configure a valid license file location and file name.

The linkDown Facility Alarm is supported for the objects listed in [Table 97](#) (note that all objects may not be supported on all platforms):

Table 97 linkDown Facility Alarm Support

Object	Supported?
Ethernet Ports	Yes
Sonet Section, Line and Path (POS)	Yes
TDM Ports (E1, T1, DS3) including CES MDAs/CMAs	Yes
TDM Channels (DS3 channel configured in an STM-1 port)	Yes
ATM Ports	Yes
Ethernet LAGs	No
APS groups	No
Bundles (MLPPP, IMA, etc)	No
ATM channels, Ethernet VLANs, Frame Relay DLCIs	No

9.7 Configuring Logging with CLI

This section provides information to configure logging using the command line interface.

Topics in this section include:

- [Basic Facility Alarm Configuration](#)
- [Common Configuration Tasks](#)

9.8 Basic Facility Alarm Configuration

The most facility alarm configuration must have the following:

- Log ID or accounting policy ID
- A log source
- A log destination

The following displays an alarm configuration example.

```
A:ALA-12>config>system# alarms
#-----
      no shutdown
      exit
-----
```

9.9 Common Configuration Tasks

The following sections are basic alarm tasks that can be performed.

- [Configuring the Maximum Number of Alarms To Clear](#)

9.9.1 Configuring the Maximum Number of Alarms To Clear

The number of alarms to clear can be configured using the command listed below.

Use the following CLI syntax to configure a log file:

CLI Syntax: config>system
 alarms
 max-cleared max-alarms

The following displays facility alarm configuration example:

```
ALA-12>config>system# alarms
-----
...
max-cleared 100
exit
...
-----
```

9.10 Facility Alarms Configuration Command Reference

9.10.1 Command Hierarchies

- [Facility Alarm Configuration Commands](#)

9.10.1.1 Facility Alarm Configuration Commands

```
config
  — system
    — alarms
      — max-cleared max-alarms
      — [no] shutdown
```

9.10.2 Command Descriptions

9.10.2.1 Generic Commands

alarms

Syntax	alarms
Context	config>system
Description	This command enters the context to configure facility alarm parameters. Alarm support is intended to cover a focused subset of router states that are likely to indicate service impacts (or imminent service impacts) related to the overall state of hardware assemblies (cards, fans, links, and so on).

max-cleared

Syntax	max-cleared <i>max-alarms</i>
Context	config>system>alarms

Description	This command configures the maximum number of cleared alarms that the system will store and display.
Default	500
Parameters	<i>max-alarms</i> — Specifies the maximum number of cleared alarms.

shutdown

Syntax	[no] shutdown
Context	config>system>alarms
Description	This command enables or disables the Facility Alarm functionality. When enabled, the Facility Alarm sub-system tracks active and cleared facility alarms and controls the Alarm LEDs on the CPMs/CFMs. When Facility Alarm functionality is enabled, the alarms are viewed using the show system alarms command(s).
Default	no shutdown

9.11 Facility Alarms Show Command Reference

9.11.1 Command Hierarchies

- [Show Commands](#)

9.11.1.1 Show Commands

```
show
  — system
    — alarms [cleared] [severity severity-level] [count count] [newer-than days]
```

9.11.2 Command Descriptions

9.11.2.1 Show Commands

The command outputs in the following section are examples only; actual displays may differ depending on supported functionality and user configuration.

alarms

Syntax	alarms [cleared] [severity <i>severity-level</i>] [count <i>count</i>] [newer-than <i>days</i>]
Context	show>system
Description	This command displays facility alarms on the system. Alarm support is intended to cover a focused subset of router states that are likely to indicate service impacts (or imminent service impacts) related to the overall state of hardware assemblies (cards, fans, links, and so on).

Output Facility Alarm Output

[Table 98](#) describes the alarms output fields.

Sample Output

Table 98 Show Facility Alarms Output Fields

Label	Description
Index	Alarm index number.
Date/Time	Date and time string for the alarm.
Severity	Severity level of the alarm.
Alarm	Alarm identifier.
Resource	Facility associated with the alarm.
Details	Description of the alarm.

```
A:Dut-A# show system alarms
=====
Alarms [Critical:1 Major:2 Minor:0 Warning:0 Total:3]
=====
Index      Date/Time          Severity    Alarm        Resource
  Details
-----
8          2011/04/01 18:36:43.80 MAJOR      7-2011-1     Power Supply 1
  Power supply 1, power lost

7          2011/04/01 18:35:57.00 MAJOR      7-2005-1     Chassis 1
  Chassis 1: temperature too high

6          2011/04/01 18:35:24.80 CRITICAL   7-2006-1     Fan 1
  Fan 1 failed
=====
```

Cleared alarms table:

```
A:Dut-A# show system alarms cleared
=====
Cleared Alarms [Size:500 Total:5 (not wrapped)]
=====
Index      Date/Time          Severity    Alarm        Resource
  Details
-----
5          2011/04/01 18:11:55.00 MAJOR      7-2005-1     Chassis 1
  Clear Chassis temperature too high alarm

3          2011/04/01 18:11:54.50 CRITICAL   7-2051-1     Power Supply 1
  Clear Power Supply failure

2          2011/04/01 18:11:54.40 CRITICAL   7-2050-1     Power Supply 1
  Clear Power Supply failure

4          2011/04/01 18:11:54.10 MINOR      7-2004-1     Fan 1
  Clear Fan wrong type failure

1          2011/04/01 18:11:54.00 CRITICAL   7-2007-1     Power Supply 1
  Clear Power Supply failure
=====
```

10 Standards and Protocol Support



Note: The information presented is subject to change without notice.

Nokia assumes no responsibility for inaccuracies contained herein.

Access Node Control Protocol (ANCP)

RFC 5851, *Framework and Requirements for an Access Node Control Mechanism in Broadband Multi-Service Networks*

Application Assurance (AA)

3GPP Release 12 (ADC rules over Gx interfaces)

RFC 3507, *Internet Content Adaptation Protocol (ICAP)*

Asynchronous Transfer Mode (ATM)

AF-ILMI-0065.000, *Integrated Local Management Interface (ILMI) Version 4.0*

AF-PHY-0086.001, *Inverse Multiplexing for ATM (IMA) Specification Version 1.1*

AF-TM-0121.000, *Traffic Management Specification Version 4.1*

AF-TM-0150.00, *Addendum to Traffic Management v4.1 optional minimum desired cell rate indication for UBR*

GR-1113-CORE, *Asynchronous Transfer Mode (ATM) and ATM Adaptation Layer (AAL) Protocols Generic Requirements, Issue 1*

GR-1248-CORE, *Generic Requirements for Operations of ATM Network Elements (NEs), Issue 3*

ITU-T I.432.1, *B-ISDN user-network interface - Physical layer specification: General characteristics (02/99)*

ITU-T I.610, *B-ISDN operation and maintenance principles and functions (11/95)*

RFC 1626, *Default IP MTU for use over ATM AAL5*

RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*

Border Gateway Protocol (BGP)

draft-hares-idr-update-attrib-low-bits-fix-01, *Update Attribute Flag Low Bits Clarification*

draft-ietf-idr-add-paths-04, *Advertisement of Multiple Paths in BGP*

draft-ietf-idr-add-paths-guidelines-08, *Best Practices for Advertisement of Multiple Paths in IBGP*

draft-ietf-idr-best-external-03, *Advertisement of the best external route in BGP*
draft-ietf-idr-bgp-flowspec-oid-03, *Revised Validation Procedure for BGP Flow Specifications*
draft-ietf-idr-bgp-gr-notification-01, *Notification Message support for BGP Graceful Restart*
draft-ietf-idr-bgp-optimal-route-reflection-10, *BGP Optimal Route Reflection (BGP-ORR)*
draft-ietf-idr-error-handling-03, *Revised Error Handling for BGP UPDATE Messages*
draft-ietf-idr-link-bandwidth-03, *BGP Link Bandwidth Extended Community*
draft-ietf-sidr-origin-validation-signaling-04, *BGP Prefix Origin Validation State Extended Community*
RFC 1772, *Application of the Border Gateway Protocol in the Internet*
RFC 1997, *BGP Communities Attribute*
RFC 2385, *Protection of BGP Sessions via the TCP MD5 Signature Option*
RFC 2439, *BGP Route Flap Damping*
RFC 2545, *Use of BGP-4 Multiprotocol Extensions for IPv6 Inter-Domain Routing*
RFC 2858, *Multiprotocol Extensions for BGP-4*
RFC 2918, *Route Refresh Capability for BGP-4*
RFC 3107, *Carrying Label Information in BGP-4*
RFC 3392, *Capabilities Advertisement with BGP-4*
RFC 4271, *A Border Gateway Protocol 4 (BGP-4)*
RFC 4360, *BGP Extended Communities Attribute*
RFC 4364, *BGP/MPLS IP Virtual Private Networks (VPNs)*
RFC 4456, *BGP Route Reflection: An Alternative to Full Mesh Internal BGP (IBGP)*
RFC 4486, *Subcodes for BGP Cease Notification Message*
RFC 4659, *BGP/MPLS IP Virtual Private Network (VPN) Extension for IPv6 VPN*
RFC 4684, *Constrained Route Distribution for Border Gateway Protocol/ MultiProtocol Label Switching (BGP/MPLS) Internet Protocol (IP) Virtual Private Networks (VPNs)*
RFC 4724, *Graceful Restart Mechanism for BGP (helper mode)*
RFC 4760, *Multiprotocol Extensions for BGP-4*
RFC 4798, *Connecting IPv6 Islands over IPv4 MPLS Using IPv6 Provider Edge Routers (6PE)*
RFC 4893, *BGP Support for Four-octet AS Number Space*
RFC 5004, *Avoid BGP Best Path Transitions from One External to Another*
RFC 5065, *Autonomous System Confederations for BGP*
RFC 5291, *Outbound Route Filtering Capability for BGP-4*
RFC 5575, *Dissemination of Flow Specification Rules*

RFC 5668, *4-Octet AS Specific BGP Extended Community*
RFC 6810, *The Resource Public Key Infrastructure (RPKI) to Router Protocol*
RFC 6811, *Prefix Origin Validation*
RFC 6996, *Autonomous System (AS) Reservation for Private Use*
RFC 7311, *The Accumulated IGP Metric Attribute for BGP*
RFC 7607, *Codification of AS 0 Processing*
RFC 7674, *Clarification of the Flowspec Redirect Extended Community*
RFC 7752, *North-Bound Distribution of Link-State and Traffic Engineering (TE) Information Using BGP*

Circuit Emulation

RFC 4553, *Structure-Agnostic Time Division Multiplexing (TDM) over Packet (SAToP)*
RFC 5086, *Structure-Aware Time Division Multiplexed (TDM) Circuit Emulation Service over Packet Switched Network (CESoPSN)*
RFC 5287, *Control Protocol Extensions for the Setup of Time-Division Multiplexing (TDM) Pseudowires in MPLS Networks*

Ethernet

IEEE 802.1AB, *Station and Media Access Control Connectivity Discovery*
IEEE 802.1ad, *Provider Bridges*
IEEE 802.1ag, *Connectivity Fault Management*
IEEE 802.1ah, *Provider Backbone Bridges*
IEEE 802.1ak, *Multiple Registration Protocol*
IEEE 802.1aq, *Shortest Path Bridging*
IEEE 802.1ax, *Link Aggregation*
IEEE 802.1D, *MAC Bridges*
IEEE 802.1p, *Traffic Class Expediting*
IEEE 802.1Q, *Virtual LANs*
IEEE 802.1s, *Multiple Spanning Trees*
IEEE 802.1w, *Rapid Reconfiguration of Spanning Tree*
IEEE 802.1X, *Port Based Network Access Control*
IEEE 802.3ab, *1000BASE-T*
IEEE 802.3ac, *VLAN Tag*
IEEE 802.3ad, *Link Aggregation*
IEEE 802.3ae, *10 Gb/s Ethernet*
IEEE 802.3ah, *Ethernet in the First Mile*

IEEE 802.3ba, *40 Gb/s and 100 Gb/s Ethernet*
IEEE 802.3i, *Ethernet*
IEEE 802.3u, *Fast Ethernet*
IEEE 802.3x, *Ethernet Flow Control*
IEEE 802.3z, *Gigabit Ethernet*
ITU-T G.8031/Y.1342, *Ethernet Linear Protection Switching*
ITU-T G.8032/Y.1344, *Ethernet Ring Protection Switching*
ITU-T Y.1731, *OAM functions and mechanisms for Ethernet based networks*

Ethernet VPN (EVPN)

draft-ietf-bess-evpn-overlay-04, *A Network Virtualization Overlay Solution using EVPN*
draft-ietf-bess-evpn-prefix-advertisement-02, *IP Prefix Advertisement in EVPN*
draft-ietf-bess-evpn-proxy-arp-nd-01, *Operational Aspects of Proxy-ARP/ND in EVPN Networks*
draft-ietf-bess-evpn-vpls-seamless-integ-00, *(PBB-)EVPN Seamless Integration with (PBB-)VPLS*
draft-ietf-bess-evpn-vpws-07, *VPWS support in EVPN*
draft-rabadan-bess-evpn-pref-df-02, *Preference-based EVPN DF Election*
draft-snr-bess-pbb-evpn-isid-cmacflush-01, *PBB-EVPN ISID-based CMAC-Flush*
RFC 7432, *BGP MPLS-Based Ethernet VPN*
RFC 7623, *Provider Backbone Bridging Combined with Ethernet VPN (PBB-EVPN)*

Frame Relay

ANSI T1.617 Annex D, *DSS1 - Signalling Specification For Frame Relay Bearer Service*
FRF.1.2, *PVC User-to-Network Interface (UNI) Implementation Agreement*
FRF.12, *Frame Relay Fragmentation Implementation Agreement*
FRF.16.1, *Multilink Frame Relay UNI/NNI Implementation Agreement*
FRF.5, *Frame Relay/ATM PVC Network Interworking Implementation*
FRF2.2, *PVC Network-to-Network Interface (NNI) Implementation Agreement*
ITU-T Q.933 Annex A, *Additional procedures for Permanent Virtual Connection (PVC) status management*

Generalized Multiprotocol Label Switching (GMPLS)

draft-ietf-ccamp-rsvp-te-srlg-collect-04, *RSVP-TE Extensions for Collecting SRLG Information*

- RFC 3471, *Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description*
- RFC 3473, *Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions*
- RFC 4204, *Link Management Protocol (LMP)*
- RFC 4208, *Generalized Multiprotocol Label Switching (GMPLS) User-Network Interface (UNI): Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Support for the Overlay Model*
- RFC 4872, *RSVP-TE Extensions in Support of End-to-End Generalized Multi-Protocol Label Switching (GMPLS) Recovery*
- RFC 5063, *Extensions to GMPLS Resource Reservation Protocol (RSVP) Graceful Restart (helper mode)*
- RFC 5884, *Bidirectional Forwarding Detection (BFD) for MPLS Label Switched Paths (LSPs)*

Intermediate System to Intermediate System (IS-IS)

- draft-ginsberg-isis-mi-bis-01, *IS-IS Multi-Instance* (single topology)
- draft-ietf-isis-mi-02, *IS-IS Multi-Instance*
- draft-kaplan-isis-ext-eth-02, *Extended Ethernet Frame Size Support*
- ISO/IEC 10589:2002, Second Edition, Nov. 2002, *Intermediate system to Intermediate system intra-domain routeing information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode Network Service (ISO 8473)*
- RFC 1195, *Use of OSI IS-IS for Routing in TCP/IP and Dual Environments*
- RFC 2973, *IS-IS Mesh Groups*
- RFC 3359, *Reserved Type, Length and Value (TLV) Codepoints in Intermediate System to Intermediate System*
- RFC 3719, *Recommendations for Interoperable Networks using Intermediate System to Intermediate System (IS-IS)*
- RFC 3787, *Recommendations for Interoperable IP Networks using Intermediate System to Intermediate System (IS-IS)*
- RFC 4971, *Intermediate System to Intermediate System (IS-IS) Extensions for Advertising Router Information*
- RFC 5120, *M-ISIS: Multi Topology (MT) Routing in IS-IS*
- RFC 5130, *A Policy Control Mechanism in IS-IS Using Administrative Tags*
- RFC 5301, *Dynamic Hostname Exchange Mechanism for IS-IS*
- RFC 5302, *Domain-wide Prefix Distribution with Two-Level IS-IS*
- RFC 5303, *Three-Way Handshake for IS-IS Point-to-Point Adjacencies*
- RFC 5304, *IS-IS Cryptographic Authentication*

RFC 5305, *IS-IS Extensions for Traffic Engineering TE*
RFC 5306, *Restart Signaling for IS-IS (helper mode)*
RFC 5307, *IS-IS Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)*
RFC 5308, *Routing IPv6 with IS-IS*
RFC 5309, *Point-to-Point Operation over LAN in Link State Routing Protocols*
RFC 5310, *IS-IS Generic Cryptographic Authentication*
RFC 6213, *IS-IS BFD-Enabled TLV*
RFC 6232, *Purge Originator Identification TLV for IS-IS*
RFC 6233, *IS-IS Registry Extension for Purges*
RFC 6329, *IS-IS Extensions Supporting IEEE 802.1aq Shortest Path Bridging*
RFC 7775, *IS-IS Route Preference for Extended IP and IPv6 Reachability*
RFC 7794, *IS-IS Prefix Attributes for Extended IPv4 and IPv6 Reachability*

Internet Protocol (IP) — Fast Reroute

draft-ietf-rtgwg-lfa-manageability-08, *Operational management of Loop Free Alternates*
RFC 5286, *Basic Specification for IP Fast Reroute: Loop-Free Alternates*
RFC 7431, *Multicast-Only Fast Reroute*
RFC 7490, *Remote Loop-Free Alternate (LFA) Fast Reroute (FRR)*

Internet Protocol (IP) — General

draft-grant-tacacs-02, *The TACACS+ Protocol*
RFC 768, *User Datagram Protocol*
RFC 793, *Transmission Control Protocol*
RFC 854, *Telnet Protocol Specifications*
RFC 1350, *The TFTP Protocol (revision 2)*
RFC 2347, *TFTP Option Extension*
RFC 2348, *TFTP Blocksize Option*
RFC 2349, *TFTP Timeout Interval and Transfer Size Options*
RFC 2428, *FTP Extensions for IPv6 and NATs*
RFC 2784, *Generic Routing Encapsulation (GRE)*
RFC 2865, *Remote Authentication Dial In User Service (RADIUS)*
RFC 2866, *RADIUS Accounting*
RFC 2867, *RADIUS Accounting Modifications for Tunnel Protocol Support*
RFC 2868, *RADIUS Attributes for Tunnel Protocol Support*

- RFC 2869, *RADIUS Extensions*
- RFC 4250, *The Secure Shell (SSH) Protocol Assigned Numbers*
- RFC 4251, *The Secure Shell (SSH) Protocol Architecture*
- RFC 4252, *The Secure Shell (SSH) Authentication Protocol* (publickey, password)
- RFC 4253, *The Secure Shell (SSH) Transport Layer Protocol*
- RFC 4254, *The Secure Shell (SSH) Connection Protocol*
- RFC 4632, *Classless Inter-domain Routing (CIDR): The Internet Address Assignment and Aggregation Plan*
- RFC 5082, *The Generalized TTL Security Mechanism (GTSM)*
- RFC 5176, *Dynamic Authorization Extensions to RADIUS*
- RFC 5656, *Elliptic Curve Algorithm Integration in the Secure Shell Transport Layer (ECDSA)*
- RFC 5880, *Bidirectional Forwarding Detection (BFD)*
- RFC 5881, *Bidirectional Forwarding Detection (BFD) IPv4 and IPv6 (Single Hop)*
- RFC 5883, *Bidirectional Forwarding Detection (BFD) for Multihop Paths*
- RFC 6398, *IP Router Alert Considerations and Usage (MLD)*
- RFC 6528, *Defending against Sequence Number Attacks*
- RFC 6929, *Remote Authentication Dial-In User Service (RADIUS) Protocol Extensions*
- RFC 7130, *Bidirectional Forwarding Detection (BFD) on Link Aggregation Group (LAG) Interfaces*

Internet Protocol (IP) — Multicast

- cisco-ipmulticast/pim-autorp-spec01, *Auto-RP: Automatic discovery of Group-to-RP mappings for IP multicast* (version 1)
- draft-dolganow-bess-mvpn-expl-track-01, *Explicit Tracking with Wild Card Routes in Multicast VPN*
- draft-ietf-idmr-traceroute-ipm-07, *A "traceroute" facility for IP Multicast*
- draft-ietf-l2vpn-vpls-pim-snooping-07, *Protocol Independent Multicast (PIM) over Virtual Private LAN Service (VPLS)*
- RFC 1112, *Host Extensions for IP Multicasting*
- RFC 2236, *Internet Group Management Protocol, Version 2*
- RFC 2375, *IPv6 Multicast Address Assignments*
- RFC 2710, *Multicast Listener Discovery (MLD) for IPv6*
- RFC 3306, *Unicast-Prefix-based IPv6 Multicast Addresses*
- RFC 3376, *Internet Group Management Protocol, Version 3*
- RFC 3446, *Anycast Rendezvous Point (RP) mechanism using Protocol Independent Multicast (PIM) and Multicast Source Discovery Protocol (MSDP)*

-
- RFC 3590, *Source Address Selection for the Multicast Listener Discovery (MLD) Protocol*
- RFC 3618, *Multicast Source Discovery Protocol (MSDP)*
- RFC 3810, *Multicast Listener Discovery Version 2 (MLDv2) for IPv6*
- RFC 3956, *Embedding the Rendezvous Point (RP) Address in an IPv6 Multicast Address*
- RFC 3973, *Protocol Independent Multicast - Dense Mode (PIM-DM): Protocol Specification (Revised) (auto-RP groups)*
- RFC 4541, *Considerations for Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) Snooping Switches*
- RFC 4601, *Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)*
- RFC 4604, *Using Internet Group Management Protocol Version 3 (IGMPv3) and Multicast Listener Discovery Protocol Version 2 (MLDv2) for Source-Specific Multicast*
- RFC 4607, *Source-Specific Multicast for IP*
- RFC 4608, *Source-Specific Protocol Independent Multicast in 232/8*
- RFC 4610, *Anycast-RP Using Protocol Independent Multicast (PIM)*
- RFC 4611, *Multicast Source Discovery Protocol (MSDP) Deployment Scenarios*
- RFC 5059, *Bootstrap Router (BSR) Mechanism for Protocol Independent Multicast (PIM)*
- RFC 5186, *Internet Group Management Protocol Version 3 (IGMPv3) / Multicast Listener Discovery Version 2 (MLDv2) and Multicast Routing Protocol Interaction*
- RFC 5384, *The Protocol Independent Multicast (PIM) Join Attribute Format*
- RFC 5496, *The Reverse Path Forwarding (RPF) Vector TLV*
- RFC 6037, *Cisco Systems' Solution for Multicast in MPLS/BGP IP VPNs*
- RFC 6512, *Using Multipoint LDP When the Backbone Has No Route to the Root*
- RFC 6513, *Multicast in MPLS/BGP IP VPNs*
- RFC 6514, *BGP Encodings and Procedures for Multicast in MPLS/IP VPNs*
- RFC 6515, *IPv4 and IPv6 Infrastructure Addresses in BGP Updates for Multicast VPNs*
- RFC 6516, *IPv6 Multicast VPN (MVPN) Support Using PIM Control Plane and Selective Provider Multicast Service Interface (S-PMSI) Join Messages*
- RFC 6625, *Wildcards in Multicast VPN Auto-Discover Routes*
- RFC 6826, *Multipoint LDP In-Band Signaling for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Path*
- RFC 7246, *Multipoint Label Distribution Protocol In-Band Signaling in a Virtual Routing and Forwarding (VRF) Table Context*

RFC 7385, *IANA Registry for P-Multicast Service Interface (PMSI) Tunnel Type Code Points*

RFC 7716, *Global Table Multicast with BGP Multicast VPN (BGP-MVPN) Procedures*

Internet Protocol (IP) — Version 4

RFC 791, *Internet Protocol*

RFC 792, *Internet Control Message Protocol*

RFC 826, *An Ethernet Address Resolution Protocol*

RFC 951, *Bootstrap Protocol (BOOTP)*

RFC 1034, *Domain Names - Concepts and Facilities*

RFC 1035, *Domain Names - Implementation and Specification*

RFC 1519, *Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy*

RFC 1534, *Interoperation between DHCP and BOOTP*

RFC 1542, *Clarifications and Extensions for the Bootstrap Protocol*

RFC 1812, *Requirements for IPv4 Routers*

RFC 1918, *Address Allocation for Private Internets*

RFC 2003, *IP Encapsulation within IP*

RFC 2131, *Dynamic Host Configuration Protocol*

RFC 2132, *DHCP Options and BOOTP Vendor Extensions*

RFC 2401, *Security Architecture for Internet Protocol*

RFC 3021, *Using 31-Bit Prefixes on IPv4 Point-to-Point Links*

RFC 3046, *DHCP Relay Agent Information Option (Option 82)*

RFC 3768, *Virtual Router Redundancy Protocol (VRRP)*

RFC 4884, *Extended ICMP to Support Multi-Part Messages (ICMPv4 and ICMPv6 Time Exceeded)*

Internet Protocol (IP) — Version 6

RFC 1981, *Path MTU Discovery for IP version 6*

RFC 2460, *Internet Protocol, Version 6 (IPv6) Specification*

RFC 2464, *Transmission of IPv6 Packets over Ethernet Networks*

RFC 2473, *Generic Packet Tunneling in IPv6 Specification*

RFC 2529, *Transmission of IPv6 over IPv4 Domains without Explicit Tunnels*

RFC 3122, *Extensions to IPv6 Neighbor Discovery for Inverse Discovery Specification*

RFC 3315, *Dynamic Host Configuration Protocol for IPv6 (DHCPv6)*

RFC 3587, *IPv6 Global Unicast Address Format*
RFC 3596, *DNS Extensions to Support IP version 6*
RFC 3633, *IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6*
RFC 3646, *DNS Configuration options for Dynamic Host Configuration Protocol for IPv6 (DHCPv6)*
RFC 3736, *Stateless Dynamic Host Configuration Protocol (DHCP) Service for IPv6*
RFC 3971, *SEcure Neighbor Discovery (SEND)*
RFC 3972, *Cryptographically Generated Addresses (CGA)*
RFC 4007, *IPv6 Scoped Address Architecture*
RFC 4193, *Unique Local IPv6 Unicast Addresses*
RFC 4291, *Internet Protocol Version 6 (IPv6) Addressing Architecture*
RFC 4443, *Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification*
RFC 4861, *Neighbor Discovery for IP version 6 (IPv6)*
RFC 4862, *IPv6 Stateless Address Autoconfiguration (router functions)*
RFC 4941, *Privacy Extensions for Stateless Address Autoconfiguration in IPv6*
RFC 5007, *DHCPv6 Leasequery*
RFC 5095, *Deprecation of Type 0 Routing Headers in IPv6*
RFC 5798, *Virtual Router Redundancy Protocol (VRRP) Version 3 for IPv4 and IPv6 (IPv6)*
RFC 5952, *A Recommendation for IPv6 Address Text Representation*
RFC 6106, *IPv6 Router Advertisement Options for DNS Configuration*
RFC 6164, *Using 127-Bit IPv6 Prefixes on Inter-Router Links*

Internet Protocol Security (IPsec)

draft-ietf-ipsec-isakmp-mode-cfg-05, *The ISAKMP Configuration Method*
draft-ietf-ipsec-isakmp-xauth-06, *Extended Authentication within ISAKMP/Oakley (XAUTH)*
RFC 2401, *Security Architecture for the Internet Protocol*
RFC 2403, *The Use of HMAC-MD5-96 within ESP and AH*
RFC 2404, *The Use of HMAC-SHA-1-96 within ESP and AH*
RFC 2405, *The ESP DES-CBC Cipher Algorithm With Explicit IV*
RFC 2406, *IP Encapsulating Security Payload (ESP)*
RFC 2407, *IPsec Domain of Interpretation for ISAKMP (IPsec DoI)*
RFC 2408, *Internet Security Association and Key Management Protocol (ISAKMP)*
RFC 2409, *The Internet Key Exchange (IKE)*

- RFC 2410, *The NULL Encryption Algorithm and Its Use With IPsec*
- RFC 3526, *More Modular Exponential (MODP) Diffie-Hellman group for Internet Key Exchange (IKE)*
- RFC 3566, *The AES-XCBC-MAC-96 Algorithm and Its Use With IPsec*
- RFC 3602, *The AES-CBC Cipher Algorithm and Its Use with IPsec*
- RFC 3706, *A Traffic-Based Method of Detecting Dead Internet Key Exchange (IKE) Peers*
- RFC 3947, *Negotiation of NAT-Traversal in the IKE*
- RFC 3948, *UDP Encapsulation of IPsec ESP Packets*
- RFC 4210, *Internet X.509 Public Key Infrastructure Certificate Management Protocol (CMP)*
- RFC 4211, *Internet X.509 Public Key Infrastructure Certificate Request Message Format (CRMF)*
- RFC 4301, *Security Architecture for the Internet Protocol*
- RFC 4303, *IP Encapsulating Security Payload*
- RFC 4307, *Cryptographic Algorithms for Use in the Internet Key Exchange Version 2 (IKEv2)*
- RFC 4308, *Cryptographic Suites for IPsec*
- RFC 4434, *The AES-XCBC-PRF-128 Algorithm for the Internet Key Exchange Protocol (IKE)*
- RFC 4868, *Using HMAC-SHA-256, HMAC-SHA-384, and HMAC-SHA-512 with IPsec*
- RFC 4945, *The Internet IP Security PKI Profile of IKEv1/ISAKMP, IKEv2 and PKIX*
- RFC 5019, *The Lightweight Online Certificate Status Protocol (OCSP) Profile for High-Volume Environments*
- RFC 5280, *Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile*
- RFC 5998, *An Extension for EAP-Only Authentication in IKEv2*
- RFC 6712, *Internet X.509 Public Key Infrastructure -- HTTP Transfer for the Certificate Management Protocol (CMP)*
- RFC 6960, *X.509 Internet Public Key Infrastructure Online Certificate Status Protocol - OCSP*
- RFC 7296, *Internet Key Exchange Protocol Version 2 (IKEv2)*
- RFC 7321, *Cryptographic Algorithm Implementation Requirements and Usage Guidance for Encapsulating Security Payload (ESP) and Authentication Header (AH)*
- RFC 7383, *Internet Key Exchange Protocol Version 2 (IKEv2) Message Fragmentation*
- RFC 7468, *Textual Encodings of PKIX, PKCS, and CMS Structures*

Label Distribution Protocol (LDP)

- draft-ietf-mpls-ldp-ip-pw-capability-09, Controlling State Advertisements Of Non-negotiated LDP Applications*
- draft-ietf-mpls-ldp-ipv6-15, Updates to LDP for IPv6*
- draft-pdutta-mpls-ldp-adj-capability-00, LDP Adjacency Capabilities*
- draft-pdutta-mpls-ldp-v2-00, LDP Version 2*
- draft-pdutta-mpls-mldp-up-redundancy-00, Upstream LSR Redundancy for Multipoint LDP Tunnels*
- draft-pdutta-mpls-multi-ldp-instance-00, Multiple LDP Instances*
- draft-pdutta-mpls-tldp-hello-reduce-04, Targeted LDP Hello Reduction*
- RFC 3037, LDP Applicability*
- RFC 3478, Graceful Restart Mechanism for Label Distribution Protocol (helper mode)*
- RFC 5036, LDP Specification*
- RFC 5283, LDP Extension for Inter-Area Label Switched Paths (LSPs)*
- RFC 5443, LDP IGP Synchronization*
- RFC 5561, LDP Capabilities*
- RFC 5919, Signaling LDP Label Advertisement Completion*
- RFC 6388, Label Distribution Protocol Extensions for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths*
- RFC 6512, Using Multipoint LDP When the Backbone Has No Route to the Root*
- RFC 6826, Multipoint LDP in-band signaling for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths*
- RFC 7032, LDP Downstream-on-Demand in Seamless MPLS*

Layer Two Tunneling Protocol (L2TP) Network Server (LNS)

- draft-mammoliti-l2tp-accessline-avp-04, Layer 2 Tunneling Protocol (L2TP) Access Line Information Attribute Value Pair (AVP) Extensions*
- RFC 2661, Layer Two Tunneling Protocol "L2TP"*
- RFC 2809, Implementation of L2TP Compulsory Tunneling via RADIUS*
- RFC 3438, Layer Two Tunneling Protocol (L2TP) Internet Assigned Numbers: Internet Assigned Numbers Authority (IANA) Considerations Update*
- RFC 3931, Layer Two Tunneling Protocol - Version 3 (L2TPv3)*
- RFC 4638, Accommodating a Maximum Transit Unit/Maximum Receive Unit (MTU/MRU) Greater Than 1492 in the Point-to-Point Protocol over Ethernet (PPPoE)*
- RFC 4719, Transport of Ethernet Frames over Layer 2 Tunneling Protocol Version 3 (L2TPv3)*

RFC 4951, *Fail Over Extensions for Layer 2 Tunneling Protocol (L2TP) "failover"*

Management

draft-ietf-snmpv3-update-mib-05, *Management Information Base (MIB) for the Simple Network Management Protocol (SNMP)*

draft-ietf-isis-wg-mib-06, *Management Information Base for Intermediate System to Intermediate System (IS-IS)*

draft-ietf-mboned-msdp-mib-01, *Multicast Source Discovery protocol MIB*

draft-ietf-mpls-ldp-mib-07, *Definitions of Managed Objects for the Multiprotocol Label Switching, Label Distribution Protocol (LDP)*

draft-ietf-mpls-lsr-mib-06, *Multiprotocol Label Switching (MPLS) Label Switching Router (LSR) Management Information Base Using SMIv2*

draft-ietf-mpls-te-mib-04, *Multiprotocol Label Switching (MPLS) Traffic Engineering Management Information Base*

draft-ietf-ospf-mib-update-08, *OSPF Version 2 Management Information Base*

draft-ietf-rrp-unified-mib-06, *Definitions of Managed Objects for the VRRP over IPv4 and IPv6 (IPv6)*

ianaaddressfamilynumbers-mib, *IANA-ADDRESS-FAMILY-NUMBERS-MIB*

ianagmplstc-mib, *IANA-GMPLS-TC-MIB*

ianaiftype-mib, *IANAifType-MIB*

ianaiprouteprotocol-mib, *IANA-RTPROTO-MIB*

IEEE8021-CFM-MIB, *IEEE P802.1ag(TM) CFM MIB*

IEEE8021-PAE-MIB, *IEEE 802.1X MIB*

IEEE8023-LAG-MIB, *IEEE 802.3ad MIB*

LLDP-MIB, *IEEE P802.1AB(TM) LLDP MIB*

RFC 1157, *A Simple Network Management Protocol (SNMP)*

RFC 1212, *Concise MIB Definitions*

RFC 1213, *Management Information Base for Network Management of TCP/IP-based Internets: MIB-II*

RFC 1215, *A Convention for Defining Traps for use with the SNMP*

RFC 1724, *RIP Version 2 MIB Extension*

RFC 2021, *Remote Network Monitoring Management Information Base Version 2 using SMIv2*

RFC 2115, *Management Information Base for Frame Relay DTEs Using SMIv2*

RFC 2206, *RSVP Management Information Base using SMIv2*

RFC 2213, *Integrated Services Management Information Base using SMIv2*

RFC 2494, *Definitions of Managed Objects for the DS0 and DS0 Bundle Interface Type*

-
- RFC 2514, *Definitions of Textual Conventions and OBJECT-IDENTITIES for ATM Management*
- RFC 2515, *Definitions of Managed Objects for ATM Management*
- RFC 2570, *SNMP Version 3 Framework*
- RFC 2571, *An Architecture for Describing SNMP Management Frameworks*
- RFC 2572, *Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)*
- RFC 2573, *SNMP Applications*
- RFC 2574, *User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3)*
- RFC 2575, *View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP)*
- RFC 2578, *Structure of Management Information Version 2 (SMIv2)*
- RFC 2579, *Textual Conventions for SMIv2*
- RFC 2580, *Conformance Statements for SMIv2*
- RFC 2787, *Definitions of Managed Objects for the Virtual Router Redundancy Protocol*
- RFC 2819, *Remote Network Monitoring Management Information Base*
- RFC 2856, *Textual Conventions for Additional High Capacity Data Types*
- RFC 2863, *The Interfaces Group MIB*
- RFC 2864, *The Inverted Stack Table Extension to the Interfaces Group MIB*
- RFC 2933, *Internet Group Management Protocol MIB*
- RFC 3014, *Notification Log MIB*
- RFC 3164, *The BSD syslog Protocol*
- RFC 3165, *Definitions of Managed Objects for the Delegation of Management Scripts*
- RFC 3231, *Definitions of Managed Objects for Scheduling Management Operations*
- RFC 3273, *Remote Network Monitoring Management Information Base for High Capacity Networks*
- RFC 3416, *Version 2 of the Protocol Operations for the Simple Network Management Protocol (SNMP)*
- RFC 3417, *Transport Mappings for the Simple Network Management Protocol (SNMP) (SNMP over UDP over IPv4)*
- RFC 3419, *Textual Conventions for Transport Addresses*
- RFC 3498, *Definitions of Managed Objects for Synchronous Optical Network (SONET) Linear Automatic Protection Switching (APS) Architectures*
- RFC 3584, *Coexistence between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework*

- RFC 3592, *Definitions of Managed Objects for the Synchronous Optical Network/ Synchronous Digital Hierarchy (SONET/SDH) Interface Type*
- RFC 3593, *Textual Conventions for MIB Modules Using Performance History Based on 15 Minute Intervals*
- RFC 3635, *Definitions of Managed Objects for the Ethernet-like Interface Types*
- RFC 3637, *Definitions of Managed Objects for the Ethernet WAN Interface Sublayer*
- RFC 3826, *The Advanced Encryption Standard (AES) Cipher Algorithm in the SNMP User-based Security Model*
- RFC 3877, *Alarm Management Information Base (MIB)*
- RFC 3895, *Definitions of Managed Objects for the DS1, E1, DS2, and E2 Interface Types*
- RFC 3896, *Definitions of Managed Objects for the DS3/E3 Interface Type*
- RFC 4001, *Textual Conventions for Internet Network Addresses*
- RFC 4022, *Management Information Base for the Transmission Control Protocol (TCP)*
- RFC 4113, *Management Information Base for the User Datagram Protocol (UDP)*
- RFC 4220, *Traffic Engineering Link Management Information Base*
- RFC 4273, *Definitions of Managed Objects for BGP-4*
- RFC 4292, *IP Forwarding Table MIB*
- RFC 4293, *Management Information Base for the Internet Protocol (IP)*
- RFC 4379, *Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures*
- RFC 4511, *Lightweight Directory Access Protocol (LDAP): The Protocol*
- RFC 4513, *Lightweight Directory Access Protocol (LDAP): Authentication Methods and Security Mechanisms (TLS)*
- RFC 4631, *Link Management Protocol (LMP) Management Information Base (MIB)*
- RFC 4878, *Definitions and Managed Objects for Operations, Administration, and Maintenance (OAM) Functions on Ethernet-Like Interfaces*
- RFC 5101, *Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of IP Traffic Flow Information*
- RFC 5102, *Information Model for IP Flow Information Export*
- RFC 5246, *The Transport Layer Security (TLS) Protocol Version 1.2 (TLS client, RSA public key)*
- RFC 5357, *A Two-Way Active Measurement Protocol (TWAMP) (server, unauthenticated mode)*
- RFC 5938, *Individual Session Control Feature for the Two-Way Active Measurement Protocol (TWAMP)*
- RFC 6020, *YANG - A Data Modeling Language for the Network Configuration Protocol (NETCONF)*

RFC 6038, *Two-Way Active Measurement Protocol (TWAMP) Reflect Octets and Symmetrical Size Features*

RFC 6241, *Network Configuration Protocol (NETCONF)*

RFC 6242, *Using the NETCONF Protocol over Secure Shell (SSH)*

RFC 6243, *With-defaults Capability for NETCONF*

RFC 6424, *Mechanism for Performing Label Switched Path Ping (LSP Ping) over MPLS Tunnels*

RFC 6425, *Detecting Data Plane Failures in Point-to-Multipoint Multiprotocol Label Switching (MPLS) - Extensions to LSP Ping*

RFC 7420, *Path Computation Element Communication Protocol (PCEP) Management Information Base (MIB) Module*

SFLOW-MIB, *sFlow MIB Version 1.3 (Draft 5)*

Multiprotocol Label Switching - Transport Profile (MPLS-TP)

RFC 5586, *MPLS Generic Associated Channel*

RFC 5921, *A Framework for MPLS in Transport Networks*

RFC 5960, *MPLS Transport Profile Data Plane Architecture*

RFC 6370, *MPLS Transport Profile (MPLS-TP) Identifiers*

RFC 6378, *MPLS Transport Profile (MPLS-TP) Linear Protection*

RFC 6426, *MPLS On-Demand Connectivity and Route Tracing*

RFC 6427, *MPLS Fault Management Operations, Administration, and Maintenance (OAM)*

RFC 6428, *Proactive Connectivity Verification, Continuity Check and Remote Defect indication for MPLS Transport Profile*

RFC 6478, *Pseudowire Status for Static Pseudowires*

RFC 7213, *MPLS Transport Profile (MPLS-TP) Next-Hop Ethernet Addressing*

Multiprotocol Label Switching (MPLS)

RFC 3031, *Multiprotocol Label Switching Architecture*

RFC 3032, *MPLS Label Stack Encoding*

RFC 3443, *Time To Live (TTL) Processing in Multi-Protocol Label Switching (MPLS) Networks*

RFC 4023, *Encapsulating MPLS in IP or Generic Routing Encapsulation (GRE)*

RFC 4182, *Removing a Restriction on the use of MPLS Explicit NULL*

RFC 5332, *MPLS Multicast Encapsulations*

RFC 6790, *The Use of Entropy Labels in MPLS Forwarding*

Network Address Translation (NAT)

- draft-ietf-behave-address-format-10, *IPv6 Addressing of IPv4/IPv6 Translators*
- draft-ietf-behave-v6v4-xlate-23, *IP/ICMP Translation Algorithm*
- draft-miles-behave-l2nat-00, *Layer2-Aware NAT*
- draft-nishitani-cgn-02, *Common Functions of Large Scale NAT (LSN)*
- RFC 4787, *Network Address Translation (NAT) Behavioral Requirements for Unicast UDP*
- RFC 5382, *NAT Behavioral Requirements for TCP*
- RFC 5508, *NAT Behavioral Requirements for ICMP*
- RFC 6146, *Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers*
- RFC 6333, *Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion*
- RFC 6334, *Dynamic Host Configuration Protocol for IPv6 (DHCPv6) Option for Dual-Stack Lite*
- RFC 6888, *Common Requirements For Carrier-Grade NATs (CGNs)*
- RFC 7383, *Internet Key Exchange Protocol Version 2 (IKEv2) Message Fragmentation*

Open Shortest Path First (OSPF)

- draft-ietf-ospf-ospfv3-lsa-extend-13, *OSPFv3 LSA Extendibility*
- RFC 1586, *Guidelines for Running OSPF Over Frame Relay Networks*
- RFC 1765, *OSPF Database Overflow*
- RFC 2328, *OSPF Version 2*
- RFC 3101, *The OSPF Not-So-Stubby Area (NSSA) Option*
- RFC 3509, *Alternative Implementations of OSPF Area Border Routers*
- RFC 3623, *Graceful OSPF Restart Graceful OSPF Restart (helper mode)*
- RFC 3630, *Traffic Engineering (TE) Extensions to OSPF Version 2*
- RFC 4203, *OSPF Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)*
- RFC 4222, *Prioritized Treatment of Specific OSPF Version 2 Packets and Congestion Avoidance*
- RFC 4552, *Authentication/Confidentiality for OSPFv3*
- RFC 4576, *Using a Link State Advertisement (LSA) Options Bit to Prevent Looping in BGP/MPLS IP Virtual Private Networks (VPNs)*
- RFC 4577, *OSPF as the Provider/Customer Edge Protocol for BGP/MPLS IP Virtual Private Networks (VPNs)*
- RFC 5185, *OSPF Multi-Area Adjacency*

RFC 5187, *OSPFv3 Graceful Restart (helper mode)*
RFC 5243, *OSPF Database Exchange Summary List Optimization*
RFC 5250, *The OSPF Opaque LSA Option*
RFC 5309, *Point-to-Point Operation over LAN in Link State Routing Protocols*
RFC 5340, *OSPF for IPv6*
RFC 5709, *OSPFv2 HMAC-SHA Cryptographic Authentication*
RFC 5838, *Support of Address Families in OSPFv3*
RFC 6987, *OSPF Stub Router Advertisement*
RFC 7684, *OSPFv2 Prefix/Link Attribute Advertisement*
RFC 7770, *Extensions to OSPF for Advertising Optional Router Capabilities*

OpenFlow

ONF *OpenFlow Switch Specification Version 1.3.1* (OpenFlow-hybrid switches)

Path Computation Element Protocol (PCEP)

draft-alvarez-pce-path-profiles-04, *PCE Path Profiles*
draft-ietf-pce-segment-routing-08, *PCEP Extensions for Segment Routing*
draft-ietf-pce-stateful-pce-14, *PCEP Extensions for Stateful PCE*
RFC 5440, *Path Computation Element (PCE) Communication Protocol (PCEP)*

Point-to-Point Protocol (PPP)

RFC 1332, *The PPP Internet Protocol Control Protocol (IPCP)*
RFC 1377, *The PPP OSI Network Layer Control Protocol (OSINLCP)*
RFC 1661, *The Point-to-Point Protocol (PPP)*
RFC 1662, *PPP in HDLC-like Framing*
RFC 1877, *PPP Internet Protocol Control Protocol Extensions for Name Server Addresses*
RFC 1989, *PPP Link Quality Monitoring*
RFC 1990, *The PPP Multilink Protocol (MP)*
RFC 1994, *PPP Challenge Handshake Authentication Protocol (CHAP)*
RFC 2153, *PPP Vendor Extensions*
RFC 2516, *A Method for Transmitting PPP Over Ethernet (PPPoE)*
RFC 2615, *PPP over SONET/SDH*
RFC 2686, *The Multi-Class Extension to Multi-Link PPP*
RFC 2878, *PPP Bridging Control Protocol (BCP)*

RFC 4638, *Accommodating a Maximum Transit Unit/Maximum Receive Unit (MTU/MRU) Greater Than 1492 in the Point-to-Point Protocol over Ethernet (PPPoE)*

RFC 5072, *IP Version 6 over PPP*

Policy Management and Credit Control

3GPP TS 29.212 Release 11, *Policy and Charging Control (PCC); Reference points (Gx support as it applies to wireline environment (BNG))*

RFC 3588, *Diameter Base Protocol*

RFC 4006, *Diameter Credit-Control Application*

Pseudowire

draft-ietf-l2vpn-vpws-iw-oam-04, *OAM Procedures for VPWS Interworking*

MFA Forum 12.0.0, *Multiservice Interworking - Ethernet over MPLS*

MFA Forum 13.0.0, *Fault Management for Multiservice Interworking v1.0*

MFA Forum 16.0.0, *Multiservice Interworking - IP over MPLS*

MFA Forum 9.0.0, *The Use of Virtual trunks for ATM/MPLS Control Plane Interworking*

RFC 3916, *Requirements for Pseudo- Wire Emulation Edge-to-Edge (PWE3)*

RFC 3985, *Pseudo Wire Emulation Edge-to-Edge (PWE3)*

RFC 4385, *Pseudo Wire Emulation Edge-to-Edge (PWE3) Control Word for Use over an MPLS PSN*

RFC 4446, *IANA Allocations for Pseudowire Edge to Edge Emulation (PWE3)*

RFC 4447, *Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP)*

RFC 4448, *Encapsulation Methods for Transport of Ethernet over MPLS Networks*

RFC 4619, *Encapsulation Methods for Transport of Frame Relay over Multiprotocol Label Switching (MPLS) Networks*

RFC 4717, *Encapsulation Methods for Transport Asynchronous Transfer Mode (ATM) over MPLS Networks*

RFC 4816, *Pseudowire Emulation Edge-to-Edge (PWE3) Asynchronous Transfer Mode (ATM) Transparent Cell Transport Service*

RFC 5085, *Pseudowire Virtual Circuit Connectivity Verification (VCCV): A Control Channel for Pseudowires*

RFC 5659, *An Architecture for Multi-Segment Pseudowire Emulation Edge-to-Edge*

RFC 5885, *Bidirectional Forwarding Detection (BFD) for the Pseudowire Virtual Circuit Connectivity Verification (VCCV)*

RFC 6073, *Segmented Pseudowire*

-
- RFC 6310, *Pseudowire (PW) Operations, Administration, and Maintenance (OAM) Message Mapping*
 - RFC 6391, *Flow-Aware Transport of Pseudowires over an MPLS Packet Switched Network*
 - RFC 6575, *Address Resolution Protocol (ARP) Mediation for IP Interworking of Layer 2 VPNs*
 - RFC 6718, *Pseudowire Redundancy*
 - RFC 6829, *Label Switched Path (LSP) Ping for Pseudowire Forwarding Equivalence Classes (FECs) Advertised over IPv6*
 - RFC 6870, *Pseudowire Preferential Forwarding Status bit*
 - RFC 7023, *MPLS and Ethernet Operations, Administration, and Maintenance (OAM) Interworking*
 - RFC 7267, *Dynamic Placement of Multi-Segment Pseudowires*

Quality of Service (QoS)

- RFC 2430, *A Provider Architecture for Differentiated Services and Traffic Engineering (PASTE)*
- RFC 2474, *Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers*
- RFC 2598, *An Expedited Forwarding PHB*
- RFC 3140, *Per Hop Behavior Identification Codes*
- RFC 3260, *New Terminology and Clarifications for Diffserv*

Resource Reservation Protocol - Traffic Engineering (RSVP-TE)

- draft-newton-mpls-te-dynamic-overbooking-00, *A Diffserv-TE Implementation Model to dynamically change booking factors during failure events*
- RFC 2702, *Requirements for Traffic Engineering over MPLS*
- RFC 2747, *RSVP Cryptographic Authentication*
- RFC 2961, *RSVP Refresh Overhead Reduction Extensions*
- RFC 3097, *RSVP Cryptographic Authentication -- Updated Message Type Value*
- RFC 3209, *RSVP-TE: Extensions to RSVP for LSP Tunnels*
- RFC 3473, *Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource Reservation Protocol-Traffic Engineering (RSVP-TE) Extensions (IF_ID RSVP_HOP object with unnumbered interfaces and RSVP-TE graceful restart helper procedures)*
- RFC 3477, *Signalling Unnumbered Links in Resource ReSerVation Protocol - Traffic Engineering (RSVP-TE)*

- RFC 3564, Requirements for Support of Differentiated Services-aware MPLS Traffic Engineering*
- RFC 3906, Calculating Interior Gateway Protocol (IGP) Routes Over Traffic Engineering Tunnels*
- RFC 4090, Fast Reroute Extensions to RSVP-TE for LSP Tunnels*
- RFC 4124, Protocol Extensions for Support of Diffserv-aware MPLS Traffic Engineering*
- RFC 4125, Maximum Allocation Bandwidth Constraints Model for Diffserv-aware MPLS Traffic Engineering*
- RFC 4127, Russian Dolls Bandwidth Constraints Model for Diffserv-aware MPLS Traffic Engineering*
- RFC 4561, Definition of a Record Route Object (RRO) Node-Id Sub-Object*
- RFC 4875, Extensions to Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label Switched Paths (LSPs)*
- RFC 4950, ICMP Extensions for Multiprotocol Label Switching*
- RFC 5151, Inter-Domain MPLS and GMPLS Traffic Engineering -- Resource Reservation Protocol-Traffic Engineering (RSVP-TE) Extensions*
- RFC 5712, MPLS Traffic Engineering Soft Preemption*
- RFC 5817, Graceful Shutdown in MPLS and Generalized MPLS Traffic Engineering Networks*

Routing Information Protocol (RIP)

- RFC 1058, Routing Information Protocol*
- RFC 2080, RIPng for IPv6*
- RFC 2082, RIP-2 MD5 Authentication*
- RFC 2453, RIP Version 2*

Segment Routing (SR)

- draft-francois-rtgwg-segment-routing-ti-lfa-04, Topology Independent Fast Reroute using Segment Routing*
- draft-gredler-idr-bgp-ls-segment-routing-ext-03, BGP Link-State extensions for Segment Routing*
- draft-ietf-isis-segment-routing-extensions-04, IS-IS Extensions for Segment Routing*
- draft-ietf-mpls-spring-lsp-ping-02, Label Switched Path (LSP) Ping/Trace for Segment Routing Networks Using MPLS Dataplane*
- draft-ietf-ospf-segment-routing-extensions-04, OSPF Extensions for Segment Routing*

Synchronous Optical Networking (SONET)/Synchronous Digital Hierarchy (SDH)

ITU-G.841, *Types and Characteristics of SDH Networks Protection Architecture, issued in October 1998 and as augmented by Corrigendum 1, issued in July 2002*

Timing

GR-1244-CORE, *Clocks for the Synchronized Network: Common Generic Criteria, Issue 3, May 2005*

GR-253-CORE, *SONET Transport Systems: Common Generic Criteria. Issue 3, September 2000*

IEEE 1588-2008, *IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems*

ITU-T G.781, *Synchronization layer functions, issued 09/2008*

ITU-T G.813, *Timing characteristics of SDH equipment slave clocks (SEC), issued 03/2003*

ITU-T G.8261, *Timing and synchronization aspects in packet networks, issued 04/2008*

ITU-T G.8262, *Timing characteristics of synchronous Ethernet equipment slave clock (EEC), issued 08/2007*

ITU-T G.8264, *Distribution of timing information through packet networks, issued 10/2008*

ITU-T G.8265.1, *Precision time protocol telecom profile for frequency synchronization, issued 10/2010*

ITU-T G.8275.1, *Precision time protocol telecom profile for phase/time synchronization with full timing support from the network, issued 07/2014*

RFC 5905, *Network Time Protocol Version 4: Protocol and Algorithms Specification*

Virtual Private LAN Service (VPLS)

RFC 4761, *Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling*

RFC 4762, *Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling*

RFC 5501, *Requirements for Multicast Support in Virtual Private LAN Services*

RFC 6074, *Provisioning, Auto-Discovery, and Signaling in Layer 2 Virtual Private Networks (L2VPNs)*

RFC 7041, *Extensions to the Virtual Private LAN Service (VPLS) Provider Edge (PE) Model for Provider Backbone Bridging*

RFC 7117, *Multicast in Virtual Private LAN Service (VPLS)*

Voice and Video

DVB BlueBook A86, *Transport of MPEG-2 TS Based DVB Services over IP Based Networks*

ETSI TS 101 329-5 Annex E, *QoS Measurement for VoIP - Method for determining an Equipment Impairment Factor using Passive Monitoring*

ITU-T G.1020 Appendix I, *Performance Parameter Definitions for Quality of Speech and other Voiceband Applications Utilizing IP Networks - Mean Absolute Packet Delay Variation & Markov Models*

ITU-T G.107, *The E Model - A computational model for use in planning*

ITU-T P.564, *Conformance testing for voice over IP transmission quality assessment models*

RFC 3550 Appendix A.8, *RTP: A Transport Protocol for Real-Time Applications (estimating the interarrival jitter)*

RFC 4585, *Extended RTP Profile for Real-time Transport Control Protocol (RTCP)-Based Feedback (RTP/AVPF)*

RFC 4588, *RTP Retransmission Payload Format*

Wireless Local Area Network (WLAN) Gateway

3GPP TS 23.402, *Architecture enhancements for non-3GPP accesses (S2a roaming based on GPRS)*

Customer Document and Product Support



Customer Documentation

[Customer Documentation Welcome Page](#)



Technical Support

[Product Support Portal](#)



Documentation Feedback

[Customer Documentation Feedback](#)

