



7210 SERVICE ACCESS SWITCH

7210 SAS OAM and Diagnostics Guide

7210 SAS-M

7210 SAS-T

7210 SAS-X

7210 SAS-R6

7210 SAS-R12

7210 SAS-Mxp

7210 SAS-S

7210 SAS-Sx

Release 9.0.R8

3HE11493AAHTQZZA

Issue: 01

September 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.

© 2013, 2016, 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

Preface	11
About This Guide	11
Audience	12
List of Technical Publications	13
Getting Started	
In This Chapter	15
Alcatel-Lucent 7210 SAS-Series Services Configuration Process	15
Mirror Services	
In This Chapter	17
Service Mirroring	18
Mirror Implementation	19
Mirror Source and Destinations	20
Mirroring Performance	23
Mirroring Configuration	25
Configuration Process Overview	27
Configuration Notes	28
Configuring Service Mirroring with CLI	31
Mirror Configuration Overview	32
Defining Mirrored Traffic	32
Basic Mirroring Configuration	34
Mirror Classification Rules	35
Common Configuration Tasks	37
Configuring a Local Mirror Service	38
Configuring a Remote Mirror Service	40
Service Management Tasks	44
Modifying a Local Mirrored Service	45
Deleting a Local Mirrored Service	46
Modifying a Remote Mirrored Service	47
Deleting a Remote Mirrored Service	49
Mirror Service Command Reference	51
Configuration Commands	55
OAM and SAA	
In This Chapter	79
OAM Overview	81
LSP Diagnostics: LSP Ping and Trace	82
LSP Ping/Trace for an LSP Using a BGP IPv4 Label Route	83
ECMP Considerations	84
Lsp-ping and lsp-trace over Unnumbered IP Interface	86
Downstream Detailed Mapping (DDMAP) TLV	86
Using DDMAP TLV in LSP Stitching and LSP Hierarchy	89
LSP Ping for RSVP P2MP LSP	92
LSP Trace for RSVP P2MP LSP	93
LSP Trace Behavior When S2L Path Traverses a Re-Merge Node	95

Table of Contents

SDP Diagnostics	98
SDP Ping	98
SDP MTU Path Discovery	98
Service Diagnostics	100
VPLS MAC Diagnostics	101
MAC Ping	101
MAC Trace	102
CPE Ping	103
MAC Populate	104
MAC Purge	104
VLL Diagnostics	106
VCCV Ping	106
Automated VCCV-Trace Capability for MS-Pseudowire	109
MPLS-TP On-Demand OAM Commands (Supported only on 7210 SAS-R6, 7210 SAS-R12 and 7210 SAS-T)	113
MPLS-TP Pseudowires: VCCV-Ping/VCCV-Trace	113
MPLS-TP LSPs: LSP-Ping/LSP Trace	114
MPLS-TP Show Commands	116
Static MPLS Labels	116
MPLS-TP Tunnel Configuration	116
MPLS-TP Path configuration	117
MPLS-TP Protection	121
BFD	122
MPLS TP Node Configuration	123
MPLS-TP Interfaces	125
Services using MPLS-TP PWs	126
MPLS-TP DEBUG COMMANDS	128
IP Performance Monitoring (IP PM)	131
Two-Way Active Measurement Protocol (TWAMP)	131
Configuration Notes	131
Two-Way Active Measurement Protocol Light (TWAMP Light)	132
Ethernet Connectivity Fault Management (ETH-CFM)	136
ETH-CFM Building Blocks	138
Loopback	154
Linktrace	155
Continuity Check (CC)	157
Alarm Indication Signal (ETH-AIS Y.1731)	159
Test (ETH-TST Y.1731)	159
Y.1731 Time Stamp Capability	159
Synthetic Loss Measurement (ETH-SL)	161
Configuration Example	163
ETH-CFM QoS Considerations	167
Configuration guidelines	167
OAM Mapping	168
CFM Connectivity Fault Conditions	168
CFM Fault Propagation Methods	169
Epipe Services	171
Fault Propagation to access dot1q/QinQ ports in Access-uplink mode	173
Configuring Fault propagation:	174

Service Assurance Agent Overview176

 Traceroute Implementation176

 NTP176

 Writing SAA Results to Accounting Files177

 Configuring SAA Test Parameters177

Y.1564 Testhead OAM tool179

 Pre-requisites for using the Testhead Tool183

 Generic pre-requisites for use of testhead tool (applicable for all 7210 SAS platforms):183

 Configuration Guidelines186

 Configuring testhead tool parameters190

OAM Performance Monitoring (OAM-PM)192

 Session194

 Standard PM Packets194

 Measurement Intervals195

 Data Structures and Storage201

 Bin Groups204

 Relating the Components205

 Monitoring206

 Accounting Policy Configuration206

 ETH-CFM Configuration207

 Service Configuration207

 OAM-PM Configuration207

 Show and Monitor Commands209

Diagnostics Command Reference217

 OAM Commands217

 OAM Performance Monitoring and Binning Commands229

 SAA Commands231

Tools Command Reference443

 Command Hierarchies443

Common CLI Command Descriptions

 In This Chapter501

 Common Service Commands502

Standards and Protocol Support

Table of Contents

List of Tables

Preface	11
Getting Started	
Table 1: Configuration Process	15
Mirror Services	
Table 2: Combinations of SAPs, spoke-sdp, and remote sources allowed in a mirror service	21
Table 3: Mirroring support for 7210 SAS-M and 7210 SAS-T Access-Uplink mode	23
Table 4: Mirroring support for 7210 SAS-M, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-X, and 7210 SAS-R6 and 7210 SAS-R12.	24
Table 5: Mirror Source Port Requirements	35
OAM and SAA	
Table 6: ETH-CFM Support Matrix for 7210 SAS-M Network Mode	140
Table 7: ETH-CFM Support Matrix for 7210 SAS-M Access-Uplink Mode	141
Table 8: ETH-CFM Support Matrix for 7210 SAS-T Access-Uplink Mode	142
Table 9: ETH-CFM Support Matrix for 7210 SAS-T Network Mode	142
Table 10: ETH-CFM Support Matrix for 7210 SAS-X devices	143
Table 11: ETH-CFM Support Matrix for 7210 SAS-R6 and 7210 SAS-R12 devices	145
Table 12: ETH-CFM Support Matrix for 7210 SAS-Mxp devices	147
Table 13: ETH-CFM Support Matrix for 7210 SAS-Sx/S 1/10GE devices	148
Table 14: ETH-CFM Support Matrix for 7210 SAS-Sx 10/100GE devices	149
Table 15: SAP Encapsulations supported for testhead	189
Table 16: Measurement Interval Start Times	196
Table 17: OAM-PM XML Keywords and MIB Reference	197
Table 18: Request Packet and Behavior	242
Table 19: Request Packet and Behavior	249
Table 20: Request Packet and Behavior	362
Table 21: Request Packet and Behavior	365
Table 22: Output fieldstools dump system-resource sap-ingress-qos	456

List of Figures

Mirror Services

Figure 1:	Service Mirroring	.18
Figure 2:	Local Mirroring Example	.25
Figure 3:	Remote Mirroring Example	.26
Figure 4:	Mirror Configuration and Implementation Flow	.27
Figure 5:	Local Mirrored Service Tasks	.37
Figure 6:	Remote Mirrored Service Tasks	.41

OAM and SAA

Figure 7:	Target FEC Stack TLV for a BGP Labeled IPv4 Prefix	.83
Figure 8:	DDMAP TLV	.87
Figure 9:	FEC Stack Change Sub-TLV	.87
Figure 10:	OAM Control Word Format	.106
Figure 11:	VCCV TLV	.107
Figure 12:	VCCV-Ping Application	.108
Figure 13:	MEP and MIP	.152
Figure 14:	MEP, MIP and MD Levels	.153
Figure 15:	CFM Loopback	.154
Figure 16:	CFM Linktrace	.155
Figure 17:	CFM Continuity Check	.157
Figure 18:	CFM CC Failure Scenario	.157
Figure 19:	SLM Example	.163
Figure 20:	Local Fault Propagation	.174
Figure 21:	7210 acting as traffic generator and traffic analyzer	.180
Figure 22:	OAM-PM Architecture Hierarchy	.192
Figure 23:	Evaluating and Computing Loss and Availability	.204
Figure 24:	Relating OAM-PM Components	.206

About This Guide

All the variants of 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, and 7210 SAS-Sx can be configured in two modes, that is in network mode and in access-uplink mode. In network mode configuration 7210 SAS-M, 7210 SAS-T, and 7210 SAS-Mxp uses IP/MPLS to provide service transport. In access-uplink mode configuration 7210 SAS-M and 7210 SAS-T uses Ethernet QinQ technology to provide service transport. The mode can be selected by configuring the BOF appropriately.

The 7210 SAS-Mxp and 7210 SAS-Sx platforms is configured only in Network Mode. The BOF does not allow the user to configure bof-uplink mode. By default, the mode is set to Network mode. In network mode configuration 7210 SAS-Mxp and 7210 SAS-Sx platforms uses IP/MPLS to provide service transport.

The 7210 SAS-Sx can be operated in standalone mode and satellite mode. The user guide provides features and CLI commands supported in 7210 SAS-Sx standalone mode of operations. Only the Basics System Configuration User Guide contains information on how to boot the 7210 SAS-Sx in satellite mode of operation.

This guide also presents examples to configure and implement various tests.

Notes:

- This user guide is applicable to all 7210 SAS-M, 7210 SAS-T, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-X, 7210 SAS-Mxp, and 7210 SAS-Sx platforms, unless specified otherwise.
- In either mode, it is expected that the user will only configure the required CLI parameters appropriate for the mode that the user intends to use. Unless otherwise noted, most of the configuration is similar in both the network mode and Access uplink mode.
- Only 7210 SAS-M and 7210 SAS-T supports access-uplink mode. 7210 SAS-Mxp and 7210 SAS-Sx platforms support only Network mode. 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12 does not support access-uplink mode. 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12 supports only MPLS uplinks and implicitly operates in network mode.
- On 7210 SAS devices, not all the CLI commands are supported on all the platforms and in all the modes. In many cases, the CLI commands are mentioned explicitly in this

document. In other cases, it is implied and easy to know the CLIs that are not supported on a particular platform.

This document 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.

Audience

This manual is intended for network administrators who are responsible for configuring the 7210 SAS-Series routers. It is assumed that the network administrators have an understanding of networking principles and configurations, routing processes, and protocols and standards, including:

- CLI concepts
- MDA, and port configuration
- QoS policies
- Services

List of Technical Publications

The 7210 SAS-M, T, X, R6, R12, Mxp, S, Sx OS documentation set is composed of the following books:

- 7210 SAS-M, T, X, R6, R12, Mxp, S, Sx OS Basic System Configuration Guide
This guide describes basic system configurations and operations.
- 7210 SAS-M, T, X, R6, R12, Mxp, S, Sx OS System Management Guide
This guide describes system security and access configurations as well as event logging and accounting logs.
- 7210 SAS-M, T, X, R6, R12, Mxp, S, Sx OS Interface Configuration Guide
This guide describes card, Media Dependent Adapter (MDA), link aggregation group (LAG), and port provisioning.
- 7210 SAS-M, T, X, R6, R12, Mxp, S, Sx OS Router Configuration Guide
This guide describes logical IP routing interfaces and associated attributes such as an IP address, port, as well as IP and MAC-based filtering.
- 7210 SAS-M, T, Mxp, S, Sx OS Services Guide
This guide describes how to configure service parameters such as customer information, and user services.
- 7210 SAS-M, T, X, R6, R12, Mxp, S, Sx OS OAM and Diagnostic Guide
This guide describes how to configure features such as service mirroring and Operations, Administration and Management (OAM) tools.
- 7210 SAS-M, T, Mxp, S, Sx Quality of Service Guide
This guide describes how to configure Quality of Service (QoS) policy management.
- 7210-SAS-M, T, X, R6, R12, Mxp, S, Sx OS MPLS Guide
This guide describes how to configure Multiprotocol Label Switching (MPLS) and Label Distribution Protocol (LDP).
- 7210-SAS-M, T, X, R6, R12, Mxp, S, Sx OS Routing Protocols Guide
This guide provides an overview of routing concepts and provides configuration examples for OSPF, IS-IS, and route policies.

Getting Started

In This Chapter

This book provides process flow information to configure service mirroring and Operations, Administration and Management (OAM) tools.

Alcatel-Lucent 7210 SAS-Series Services Configuration Process

[Table 1](#) lists the tasks necessary to configure mirroring, and perform tools monitoring functions.

This guide 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 1: Configuration Process

Area	Task	Chapter
Diagnostics/ Service verification	Mirroring	Mirror Services on page 17
	OAM	OAM and SAA on page 79
Reference	List of IEEE, IETF, and other proprietary entities.	Standards and Protocols section.

In This Chapter

Mirror Services

In This Chapter

This chapter provides information to configure mirroring.

Topics in this chapter include:

- [Service Mirroring on page 18](#)
- [Mirror Implementation on page 19](#)
 - [Mirror Source and Destinations on page 20](#)
 - [Local and Remote Mirroring on page 22](#)
 - [Mirroring Performance on page 23](#)
 - [Mirroring Configuration on page 25](#)
- [Configuration Process Overview on page 27](#)
- [Configuration Notes on page 28](#)
- [Configuring Service Mirroring with CLI on page 31](#)
- [Basic Mirroring Configuration on page 34](#)
- [Common Configuration Tasks on page 37](#)
- [Service Management Tasks on page 44](#)
- [Mirror Service Command Reference on page 51](#)
- [Configuration Commands on page 55](#)

Service Mirroring

When troubleshooting complex operational problems, customer packets can be examined as they traverse the network. Alcatel-Lucent's service mirroring provides the capability to mirror customer packets to allow for trouble shooting and offline analysis.

This capability also extends beyond troubleshooting services. Telephone companies have the ability to obtain itemized calling records and wire-taps where legally required by investigating authorities. The process can be very complex and costly to carry out on data networks. Service Mirroring greatly simplifies these tasks, as well as reduces costs through centralization of analysis tools and skilled technicians.

Original packets are forwarded while a copy is sent out the mirrored port to the mirroring (destination) port. Service mirroring allows an operator to see the actual traffic on a customer's service with a sniffer sitting in a central location. In many cases, this reduces the need for a separate, costly overlay sniffer network.

7210 SAS devices configured in access-uplink mode support only local mirroring.

When using local mirroring user has an option to use NULL SAP or a dot1q SAP or a Q1.* SAP as mirror destination. Use of Dot1q SAP or a Q1.* SAP as the mirror destination allows the mirrored traffic to share the same uplink as the service traffic (when the uplinks are L2 based).

On some 7210 SAS platforms, when using Dot1q SAP or a Q1.* SAP or MPLS SDP as the mirror destination user needs to dedicate the resources of a port for use with mirror application (see below for more details).




Figure 1: Service Mirroring

Mirror Implementation

Mirroring can be configured on ingress or egress of certain service entities (For example, SAPs, ports, filter entries) and they are referred to as mirror sources. For more information, see the [“Mirror Source and Destinations on page 20”](#).

Alcatel-Lucent’s implementation of packet mirroring is based on the following assumptions:

- Ingress and egress packets are mirrored as they appear on the wire. This is important for troubleshooting encapsulation and protocol issues.
When mirroring at ingress, an exact copy of the original ingress packet is sent to the mirror destination while normal forwarding proceeds on the original packet.
- When mirroring is at egress, the system performs normal packet handling on the egress packet, encapsulating it for the destination interface. A copy of the forwarded packet (as seen on the wire) is forwarded to the mirror destination.
 - On M, X, Mxp, Sx/S 1/10GE, R6 with IMMV2, R12, the mirror copy of the packet is a copy of the forwarded copy.
 - On 7210 SAS- T, 7210 SAS-Sx 10/100GE, and 7210 SAS-R6 IMMv1, the mirror copy of the packet is not a exact copy of the forwarded copy in case of port egress mirroring.
 - In 7210 SAS, mirroring at egress takes place before the packet is processed by egress QoS. Hence, there exists a possibility that a packet is dropped by egress QoS mechanisms (because of RED mechanisms and so on) and thus not forwarded, but it is still mirrored.
 - Remote destinations are reached by encapsulating the ingress or egress packet within an SDP, like the traffic for distributed VPN connectivity services. At the remote destination, the tunnel encapsulation is removed and the packet is forwarded out a local SAP.

Mirror Source and Destinations

Mirror sources and destinations have the following characteristics for 7210 SAS devices operating in network mode.

- Mirror source and mirror destination can be on the same node (local mirroring) or on different nodes (remote mirroring).
- Each mirror destination should terminate on a distinct port carrying only null encapsulation or a Dot1q SAP or a Q1.* SAP or a MPLS SDP in case of remote mirroring.
- Packets ingressing a port can have a mirror destination separate from packets egressing another or the same port (the ports must be on the same node).
- Multiple mirror destinations are supported (local only) on a single chassis.

Listed below are the mirror sources and destination characteristics for 7210 SAS devices configured in **access-uplink** mode:

- Mirroring source and destination needs to be on the same node (i.e. only local mirroring is supported).
- A mirror destination can terminate on only one port (NULL SAP or dot1q SAP or a Q1.* SAP).
- Packets ingressing a port can have a mirror destination separate from packets egressing another or the same port.

The following table lists the allowed combinations of SAPs, spoke-sdp and remote sources allowed in a mirror service using different mirror-source-type on 7210 SAS devices configured in network mode:

:

Table 2: Combinations of SAPs, spoke-sdp, and remote sources allowed in a mirror service

Mirror-source-type	Mirror sources Allowed	Mirror Destination Allowed
Local	Port Ingress Port Egress SAP ingress ACL ingress	NULL SAP Dot1q SAP QinQ SAP Spoke-SDP
Remote	remote-source	NULL SAP Dot1q SAP QinQ SAP
Both	Port Ingress Port Egress SAP ingress ACL ingress remote-source	NULL SAP Dot1q SAP QinQ SAP

Local and Remote Mirroring

NOTE: This sections describes the local and remote mirroring that are applicable for the platforms on which the mirroring feature is supported. The 7210 SAS devices configured in access-uplink mode supports only local mirroring and the 7210 SAS devices configured in network mode supports both local and remote mirroring.

The 7210 SAS devices allows multiple concurrent mirroring sessions so traffic from more than one ingress mirror source can be mirrored to the same or different mirror destinations. For more information, see the [“Configuration Notes on page 28”](#).

Remote mirroring uses a service distribution path (SDP) which acts as a logical way of directing traffic from one router to another through a uni-directional (one-way) service tunnel. The SDP terminates at the far-end router which directs packets to the correct destination on that device.

The SDP configuration from the mirrored device to a far-end router requires a return path SDP from the far-end router back to the mirrored router. Each device must have an SDP defined for every remote router to which it provides mirroring services. SDPs must be created first, before services can be configured.

Mirroring Performance

Replication of mirrored packets can, typically, affect performance and should be used carefully.

The following tables mentions about the mirroring that can be performed based on the following criteria (that is, mirror sources):

Table 3: Mirroring support for 7210 SAS-M and 7210 SAS-T Access-Uplink mode

Mirroring	7210 SAS-M	7210 SAS-T
Port (ingress and egress)	Supported	Supported
SAP (ingress only)	Supported	Supported
MAC filter (ingress only)	Supported	Supported
IP filter (ingress only)	Supported	Supported

Table 4: Mirroring support for 7210 SAS-M, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-X, and 7210 SAS-R6 and 7210 SAS-R12.

Platforms	Port (ingress and egress)	SAP (ingress only)	MAC filter (ingress only)	IP filter (ingress only)
7210 SAS-M	Supported	Supported	Supported	Supported
7210 SAS-T	Supported	Supported	Supported	Supported
7210 SAS-Mxp	Supported	Supported	Supported	Supported
7210 SAS-Sx/S 1/10GE	Supported	Supported	Supported	Supported
7210 SAS-Sx 10/100GE	Supported	Supported	Supported	Supported
7210 SAS-X	Supported	Supported	Supported	Supported
7210 SAS-R6	Supported	Supported	Supported	Supported

Mirroring Configuration

Configuring mirroring is similar to creating a uni-direction service. Mirroring requires the configuration of:

- Mirror source — The traffic on a specific point(s) to mirror.
- Mirror destination — The location to send the mirrored traffic, where the sniffer will be located.

Figure 2 depicts a local mirror service configured on ALA-A.

- Port 1/1/2 is specified as the source. Mirrored traffic ingress and egressing this port will be sent to port 1/1/3.
- SAP 1/1/3 is specified as the destination. The sniffer is physically connected to this port. Mirrored traffic ingress and egressing port 1/1/2 is sent here. SAP, encapsulation requirements, and mirror classification parameters are configured.




Figure 2: Local Mirroring Example

Figure 3 depicts a remote mirror service configured as ALA B as the mirror source and ALA A as the mirror destination. Mirrored traffic ingress and egressing port 5/2/1 (the source) on ALA B is handled the following ways:

- Port 5/2/1 is specified as the mirror source port. Parameters are defined to select specific traffic ingress and egressing this port.

Destination parameters are defined to specify where the mirrored traffic is sent. In this case, mirrored traffic sent to a SAP configured as part of the mirror service on port 3/1/3 on ALA A (the mirror destination).

Mirroring Configuration

ALA A decodes the service ID and sends the traffic out of port 3/1/3. The sniffer is physically connected to this port (3/1/3). SAP, encapsulation requirements, packet slicing, and mirror classification parameters are configured in the destination parameters.

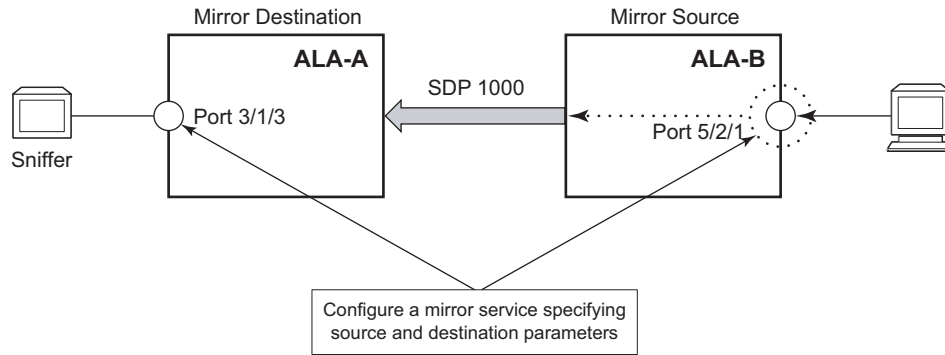


Figure 3: Remote Mirroring Example

Configuration Process Overview

Figure 4 displays the process to provision basic mirroring parameters.

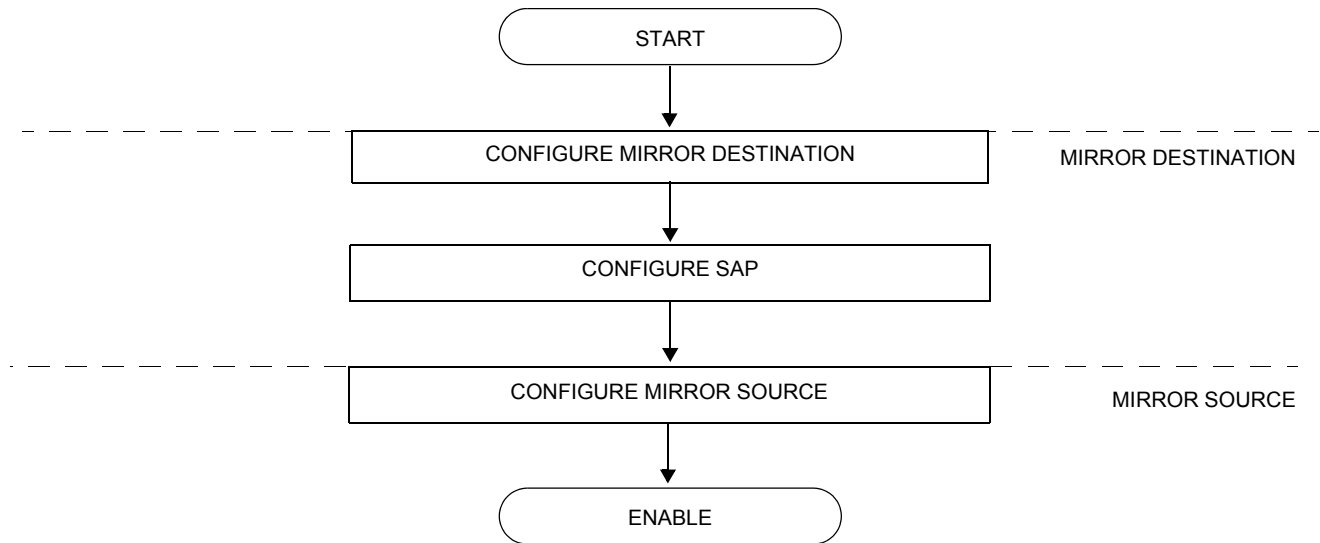


Figure 4: Mirror Configuration and Implementation Flow

Configuration Notes

This section describes mirroring configuration caveats.

- Multiple mirroring service IDs (mirror destinations) may be created within a single system.
- A mirrored source can only have one destination.
- On 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-M, 7210 SAS-T, 7210 SAS-X, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE before using a Dot1q SAP or Q1.* SAP as a mirror destination, the user must configure a port for use with this feature using the command `config> system> loopback-no-svc-port mirror`. The user has an option to use either one of the available virtual internal port resources or a front panel port. The virtual internal port resources available can be determined using the command `show system internal-loopback-ports detail`. No services can be configured on this port. More details of both the commands can be found in the 7210 SAS Interfaces Configuration Guide.
- On 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-M, 7210 SAS-X, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE before using a MPLS SDP as a mirror destination, the user must configure a port for use with this feature using the command `config> system> loopback-no-svc-port mirror`. No services can be configured on this port. The user has an option to use either one of the available virtual internal port resources or a front panel port. The virtual internal port resources available can be determined using the command `show system internal-loopback-ports detail`. More details of both the commands can be found in the 7210 SAS Interfaces Configuration Guide.
- Spoke SDP is supported only on local mirror service type. Please refer to the [Combinations of SAPs, spoke-sdp, and remote sources allowed in a mirror service on page 21](#) above for more information.
- Remote source mirror type service accepts only MPLS labeled traffic from remote sources.
- The destination mirroring service IDs and service parameters are persistent between router (re)boots and are included in the configuration saves.

Mirror source criteria configuration (defined in `debug>mirror>mirror-source`) is not preserved in a configuration save (admin save). Debug mirror source configuration can be saved using `admin>debug-save`.

- Physical layer problems such as collisions, jabbers, etc., are not mirrored. Typically, only complete packets are mirrored.
- Starting and shutting down mirroring:

Mirror destinations:

- The default state for a mirror destination service ID is shutdown. You must issue a **no shutdown** command to enable the feature.
- When a mirror destination service ID is shutdown, mirrored packets associated with the service ID are not accepted from its mirror source. The associated mirror source is put into an operationally down mode. Mirrored packets are not transmitted out the SAP. Each mirrored packet is silently discarded.
- Issuing the `shutdown` command causes the mirror destination service or its mirror source to be put into an administratively down state. Mirror destination service IDs must be shut down first in order to delete a service ID, or SAP association from the system.

Mirror sources:

- The default state for a mirror source for a given mirror-dest service ID is `no shutdown`. Enter a `shutdown` command to deactivate (disable) mirroring from that mirror-source.
- Mirror sources do not need to be shutdown to remove them from the system. When a mirror source is shutdown, mirroring is terminated for all sources defined locally for the mirror destination service ID.

Configuring Service Mirroring with CLI

This section provides information about service mirroring

Topics in this section include:

- [Mirror Configuration Overview on page 32](#)
- [Basic Mirroring Configuration on page 34](#)
 - [Mirror Classification Rules on page 35](#)
- [Common Configuration Tasks on page 37](#)
 - [Configuring a Local Mirror Service on page 38](#)
 - [Configuring a Remote Mirror Service on page 40](#)
- [Service Management Tasks on page 44](#)
 - [Modifying a Local Mirrored Service on page 45](#)
 - [Deleting a Local Mirrored Service on page 46](#)
 - [Modifying a Remote Mirrored Service on page 47](#)
 - [Deleting a Remote Mirrored Service on page 49](#)

Mirror Configuration Overview

7210 SAS node mirroring can be organized in the following logical entities:

- The mirror source is defined as the location from where the traffic should be mirrored. A mirror source could be ingress of service entity or egress of a service entity. The list of mirror sources supported on a given platform is listed above [Mirror Source Port Requirements on page 35](#)
 - A SAP is defined in local mirror services as the mirror destination to where the mirrored packets are sent.
-

Defining Mirrored Traffic

In some scenarios, or when multiple services are configured on the same port, specifying the port does not provide sufficient resolution to separate traffic. In Alcatel-Lucent's implementation of mirroring, multiple source mirroring parameters can be specified to further identify traffic.

Mirroring of packets matching specific filter entries in an IP or MAC filter can be applied to refine what traffic is mirrored to flows of traffic within a service. The IP criteria can be combinations of:

- Source IP address/mask
- Destination IP address/mask
- IP Protocol value
- Source port value (for example, UDP or TCP port)
- Destination port value (for example, UDP or TCP port)
- DiffServ Code Point (DSCP) value
- ICMP code
- ICMP type
- IP fragments
- TCP ACK set/reset
- TCP SYN set/reset

The MAC criteria can be combinations of:

- IEEE 802.1p value/mask
- Source MAC address/mask
- Destination MAC address/mask

- Ethernet Type II Ethernet type value

NOTE: The list of packet fields that are available to match packets in IP and MAC ACLs for different platforms is different. For more information on the lists of packet fields available on different platforms, see the “7210 SAS Router Configuration User Guide”.

Basic Mirroring Configuration

Destination mirroring parameters must include at least:

- A mirror destination ID (same as the mirror source service ID).
- A mirror destination SAP.

Mirror source parameters must include at least:

- A mirror service ID (same as the mirror destination service ID).
- At least one source type (port, SAP, IP filter or MAC filter) specified.

The following example displays a sample configuration of a local mirrored service (ALA-A).

```
*A:ALA-A>config>mirror# info
-----
mirror-dest 103 create
  sap 1/1/1 create
  exit
  no shutdown
  exit
-----
*A:ALA-A>config>mirror#
```

The following displays the mirror source configuration:

```
*A:ALA-A>debug>mirror-source# show debug mirror
debug
  mirror-source 103
    port 1/1/24 egress ingress
    no shutdown
  exit
exit
*A:ALA-A>debug>mirror-source# exit
```

Mirror Classification Rules

Alcatel-Lucent's implementation of mirroring can be performed by configuring parameters to select network traffic according to any of the following entities:

- [Port](#)
- [SAP](#)
- [MAC filter](#)
- [IP filter](#)

Port

The `port` command associates a port to a mirror source. The port is identified by the port ID. The defined port can be Ethernet or a Link Aggregation Group (LAG) ID. When a LAG ID is given as the port ID, mirroring is enabled on all ports making up the LAG.

Mirror sources can be ports in either access or network mode. Port mirroring is supported in the following combinations:

Table 5: Mirror Source Port Requirements

Port Type	Port Mode	Port Encap Type
faste/gige/10gige	access	null, dot1q and QinQ
faste/gige/10gige	access uplink	qinq
faste/gige/10gige	network	null/dot1q
faste/gige/10gige	hybrid	null/dot1q/qinq

CLI Syntax: `debug>mirror-source# port {port-id|lag lag-id} {[egress] [ingress]}`

Example: `*A:ALA-A>debug>mirror-source# port 1/1/2 ingress egress`

SAP

More than one SAP can be associated within a single mirror-source. Each SAP has its own ingress parameter keyword to define which packets are mirrored to the mirror-dest service ID. A SAP that is defined within a mirror destination cannot be used in a mirror source.

CLI Syntax: `debug>mirror-source# sap sap-id {[ingress]}`

Mirror Classification Rules

Example: *A:ALA-A>debug>mirror-source# sap 1/1/4:100 ingress

MAC filter

MAC filters are configured in the **config>filter>mac-filter** context. The **mac-filter** command causes all the packets matching the explicitly defined list of entry IDs to be mirrored to the mirror destination specified by the service-id of the mirror source.

CLI Syntax: debug>mirror-source# mac-filter *mac-filter-id* entry *entry-id* [*entry-id* ...]

Example: *A:ALA-2>debug>mirror-source# mac-filter 12 entry 15 20 25

IP filter

IP filters are configured in the **config>filter>ip-filter** context. The **ip-filter** command causes all the packets matching the explicitly defined list of entry IDs to be mirrored to the mirror destination specified by the service-id of the mirror source.

Ingress mirrored packets are mirrored to the mirror destination prior to any ingress packet modifications.

CLI Syntax: debug>mirror-source# ip-filter *ip-filter-id* entry *entry-id* [*entry-id* ...]

Example: *A:ALA-A>debug>mirror-source# ip-filter 1 entry 20

NOTE: An IP filter cannot be applied to a mirror destination SAP.

Common Configuration Tasks

This section provides a brief overview of the tasks that must be performed to configure local mirror services and provides CLI command syntax. Note that the local mirror source and mirror destination components must be configured under the same service ID context.

Each local mirrored service (Figure 5) (within the same router) requires the following configurations:

1. Specify mirror destination (SAP).
2. Specify mirror source (port, SAP, IP filter, MAC filter).




Figure 5: Local Mirrored Service Tasks

Configuring a Local Mirror Service

To configure a local mirror service, the source and destinations must be located on the same router. Note that local mirror source and mirror destination components must be configured under the same service ID context.

The **mirror-source** commands are used as traffic selection criteria to identify traffic to be mirrored at the source. Each of these criteria are independent. For example, use the **debug>mirror-source>port** *{port-id | lag lag-id}* *{[egress] [ingress]}* command and **debug>mirror-source ip-filter** *ip-filter-id entry entry-id [entry-id...]* command to capture (mirror) traffic that matches a specific IP filter entry and traffic ingressing and egressing a specific port. A filter must be applied to the SAP or interface if only specific packets are to be mirrored.

Use the CLI syntax to configure one or more mirror source parameters:

The **mirror-dest** commands are used to specify where the mirrored traffic is to be sent. Use the following CLI syntax to configure mirror destination parameters:

CLI Syntax: `config>mirror mirror-dest service-id [type {ether}] [create]
description string
sap sap-id [create]
no shutdown`

CLI Syntax: `debug# mirror-source service-id
ip-filter ip-filter-id entry entry-id [entry-id ...]
ipv6-filter ip-filter-id entry entry-id [entry-id ...]
mac-filter mac-filter-id entry entry-id [entry-id ...]
port {port-id|lag lag-id} {[egress] [ingress]}
sap sap-id {[ingress]}
no shutdown`

The following output displays an example of a local mirrored service using a NULL SAP. On ALA-A, mirror service 103 is mirroring traffic matching IP filter 2, entry 1 as well as egress and ingress traffic on port 1/1/23 and sending the mirrored packets to SAP 1/1/24

```
*A:ALA-A>config>mirror# info
-----
mirror-dest 103 create
  sap 1/1/24 create
  exit
  no shutdown
  exit
-----
*A:ALA-A>config>mirror#
```

The following output displays an example of local mirrored service using a dot1q SAP. User needs to configure a front-panel port for use with the mirroring application when the mirror destination is a Dot1q SAP or a Q1.* SAP, as shown below.

```
*A:ALA-A>config>system>
-----
      loopback-no-svc-port mirror 1/1/14
-----

*A:ALA-A>config>mirror# info
-----
      mirror-dest 103 create
          sap 1/1/10:100 create
          exit
          no shutdown
      exit
-----

*A:ALA-A>config>mirror#
```

The following displays the debug mirroring information:

```
*A:ALA-A>debug>mirror-source# show debug mirror
debug
      mirror-source 103
          no shutdown
          port 1/1/23 ingress
          ip-filter 2 entry 1
      exit
exit
*A:ALA-A>debug>mirror-source# exit
```

Configuring a Remote Mirror Service

The source and destination are configured on different routers for remote mirroring. Note that *mirror source* and *mirror destination* parameters must be configured under the same service ID context.

NOTE: Remote Mirroring using MPLS SDP is supported only on 7210 SAS devices configured in Network Mode. It is not supported on 7210 SAS devices configured in access-uplink mode.

The **mirror-source** commands are used as traffic selection criteria to identify traffic to be mirrored at the source. For example, use the **port** *port-id* [*lag-id*] {[**egress**] [**ingress**]}, and **mac-filter** *mac-filter-id* **entry** *entry-id* [*entry-id* ...] commands.

Use the CLI syntax to configure one or more mirror source parameters:

```
CLI Syntax: debug> mirror-source service-id
                ip-filter ip-filter-id entry entry-id [entry-id ...]
                ipv6-filter ip-filter-id entry entry-id [entry-id ...]
                mac-filter mac-filter-id entry entry-id [entry-id ...]
                port {port-id|lag lag-id} {[egress] [ingress]}
                sap sap-id {[ingress]}
                no shutdown
```

The **mirror-dest** commands are used to specify where the mirrored traffic is to be sent, the forwarding class, and the size of the packet. Use the following CLI syntax to configure mirror destination parameters:

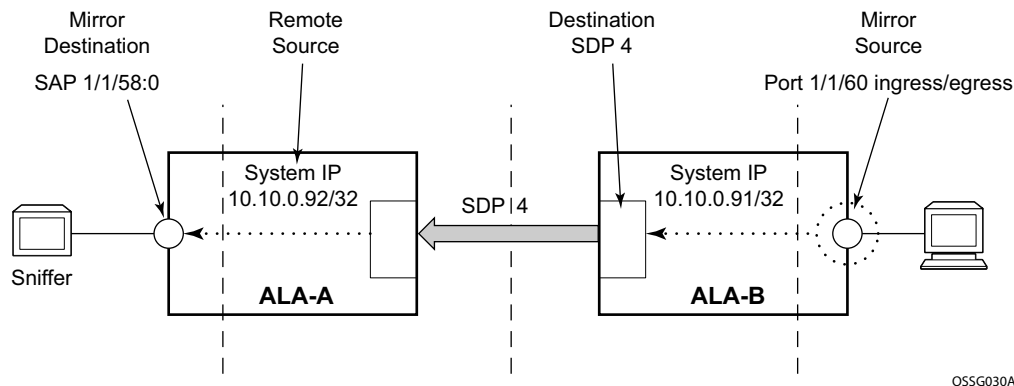
```
CLI Syntax: config>mirror#
mirror-dest service-id

                [create] [type <mirror-type>] [mirror-source-type <mirror-
source-type>]
                description string
                fc fc-name [profile <profile>]
                remote-source
                    far-end <ip-address> [vc-id <vc-id>] [ing-svc-label <i-
ngress-vc-label>|tldp]
                sap sap-id create
                no shutdown
```

The following [Table 6, Remote Mirrored Service Tasks, on page 41](#) displays the mirror destination, which is on ALA-A, configuration for mirror service 1216. This configuration specifies that the mirrored traffic coming from the mirror source (10.10.0.91) is to be directed to SAP /1/58 and states that the service only accepts traffic from far end 10.10.0.92 (ALA-B) with an ingress service

label of 5678. When a forwarding class is specified, then all mirrored packets transmitted to the destination SAP or SDP override the default (be) forwarding class.

Figure 6: Remote Mirrored Service Tasks



OSSG030A

The following example displays the CLI output showing the configuration of remote mirrored service 1216. The traffic ingressing and egressing port 1/1/60 on 10.10.0.92 (ALA-B) will be mirrored to the destination SAP 1/1/58:0 on ALA-A.

The following example displayed is the remote mirror destination configuring the front panel port with mirroring application:

```
*A:7210SAS>config>mirror# info
-----
mirror-dest 23 mirror-source-type remote create
  description "Added by createMirrorDestination 23"
  fc be
  remote-source
    far-end 2.2.2.2 ing-svc-label 14000
  exit
  sap 1/1/4 create
  exit
  no shutdown
exit
mirror-dest 1000 create
  fc be
  spoke-sdp 200:1000 create
    egress
    vc-label 15000
  exit
  no shutdown
exit
  no shutdown
exit
-----
*A:7210SAS>config>mirror# /show system internal-loopback-ports
=====
```

Configuring a Remote Mirror Service

```
Internal Loopback Port Status
=====
```

Port Id	Loopback Type	Application	Service Enabled
1/1/9	Physical	Dot1q-Mirror	No

```
=====
```

The following example displayed is the mirror destination configuration for mirror service 1216 on ALA-A.

```
*A:ALA-A>config>mirror# info
-----
mirror-dest 1000 type ether mirror-source-type remote create
description "Receiving mirror traffic from .91"
remote-source
  far-end 2.2.2.2 tldp
exit
sap 1/1/21:21 create
  egress
  qos 1
  exit
exit
no shutdown
exit
-----
*A:ALA-A>config>mirror#
```

The following example displays the remote mirror destination configured on ALA-B:

```
*A:ALA-B>config>mirror# info
-----
mirror-dest 2000 type ether mirror-source-type local create
no description
no service-name
fc be
no remote-source
spoke-sdp 200:2000 create
  egress
  no vc-label
  exit
  no shutdown
exit
no shutdown
exit
-----
*A:ALA-B>config>mirror#
```

The following example displays the mirror source configuration for ALA-B:

```
*A:ALA-B# show debug mirror
```

```
debug
  mirror-source 1000
    no shutdown
  exit
  mirror-source 2000
    no shutdown
  exit
exit
*A:ALA-B#
```

The following example displays the SDP configuration from ALA-A to ALA-B (SDP 2) and the SDP configuration from ALA-B to ALA-A (SDP 4).

```
*A:ALA-A>config>service>sdp# info
-----
      description "MPLS-10.10.0.91"
      far-end 10.10.0.01
      signalling tldp
      no shutdown
-----
*A:ALA-A>config>service>sdp#

*A:ALA-B>config>service>sdp# info
-----
      description "MPLS-10.10.20.92"
      far-end 10.10.10.103
      signalling tldp
      no shutdown
-----
*A:ALA-B>config>service>sdp#
```

Service Management Tasks

This section discusses the following service management tasks:

- [Modifying a Local Mirrored Service on page 45](#)
- [Deleting a Local Mirrored Service on page 46](#)
- [Modifying a Remote Mirrored Service on page 47](#)
- [Deleting a Remote Mirrored Service on page 49](#)

Use the following command syntax to modify an existing mirrored service:

CLI Syntax: config>mirror#
mirror-dest *service-id* [type {ether}]
description *description-string*
no description
sap *sap-id*
no sap
[no] shutdown

CLI Syntax: debug
[no] mirror-source *service-id*
ip-filter *ip-filter-id* entry *entry-id* [*entry-id...*]
no ip-filter *ip-filter-id*
no ip-filter entry *entry-id* [*entry-id...*]
ipv6-filter *ip-filter-id* entry *entry-id* [*entry-id...*]
no ipv6-filter *ip-filter-id*
no ipv6-filter entry *entry-id* [*entry-id...*]
mac-filter *mac-filter-id* entry *entry-id* [*entry-id...*]
no mac-filter *mac-filter-id*
no mac-filter *mac-filter-id* entry *entry-id* [*entry-id...*]
[no] port {*port-id*|lag *lag-id*} {[egress][ingress]}
[no] sap *sap-id* {[ingress]}
[no] shutdown

Modifying a Local Mirrored Service

Existing mirroring parameters can be modified in the CLI. The changes are applied immediately. The service must be shut down if changes to the SAP are made.

The following example displays commands to modify parameters for a basic local mirroring service.

```
Example: config>mirror# mirror-dest 103
config>mirror>mirror-dest# shutdown
config>mirror>mirror-dest# no sap
config>mirror>mirror-dest# sap 1/1/5 create
config>mirror>mirror-dest>sap$ exit
config>mirror>mirror-dest# no shutdown
debug# mirror-source 103
debug>mirror-source# no port 1/1/23
debug>mirror-source# port 1/1/7 ingress egress
```

The following displays the local mirrored service modifications:

```
*A:ALA-A>config>mirror# info
-----
mirror-dest 103 create
      no shutdown
      sap 1/1/5 create
      exit

*A:ALA-A>debug>mirror-source# show debug mirror
debug
      mirror-source 103
      no shutdown
      port 1/1/7 egress ingress
      exit
*A:ALA-A>debug>mirror-source#
```

Deleting a Local Mirrored Service

Existing mirroring parameters can be deleted in the CLI. A shutdown must be issued on a service level in order to delete the service. It is not necessary to shut down or remove SAP or port references to delete a local mirrored service.

The following example displays commands to delete a local mirrored service.

```
Example:ALA-A>config>mirror# mirror-dest 103
config>mirror>mirror-dest# shutdown
config>mirror>mirror-dest# exit
config>mirror# no mirror-dest 103
config>mirror# exit
```

Modifying a Remote Mirrored Service

Existing mirroring parameters can be modified in the CLI. The changes are applied immediately. The service must be shut down if changes to the SAP are made.

In the following example, the mirror destination is changed from 10.10.10.2 (ALA-B) to 10.10.10.3 (SR3). Note that the mirror-dest service ID on ALA-B must be shut down first before it can be deleted.

The following example displays commands to modify parameters for a remote mirrored service.

```
Example:*A:ALA-A>config>mirror# mirror-dest 104
config>mirror>mirror-dest# remote-source
config>mirror>mirror-dest>remote-source# no far-end 10.10.10.2
remote-source# far-end 10.10.10.3 ing-svc-label 3500
```

```
*A:ALA-B>config>mirror# mirror-dest 104
config>mirror>mirror-dest# shutdown
config>mirror>mirror-dest# exit
config>mirror# no mirror-dest 104
```

```
SR3>config>mirror# mirror-dest 104 create
config>mirror>mirror-dest# sdp 4 egr-svc-label 3500
config>mirror>mirror-dest# no shutdown
config>mirror>mirror-dest# exit all
```

```
SR3># debug
debug# mirror-source 104
debug>mirror-source# port 551/1/2 ingress egress
debug>mirror-source# no shutdown
```

```
*A:ALA-A>config>mirror# info
```

```
-----
mirror-dest 104 create
  remote-source
    far-end 2.2.2.2 tldp
  exit
  sap 1/1/21:21 create
    egress
      qos 1
    exit
  exit
  no shutdown
exit
```

```
A:SR3>config>mirror# info
```

```
-----
mirror-dest 104 create
spoke-sdp 200:2000 create
```

Modifying a Remote Mirrored Service

```
        no shutdown
        exit
-----
A:SR3>config>mirror#

A:SR3# show debug mirror
debug
    mirror-source 104
    no shutdown
```


Deleting a Remote Mirrored Service

Existing mirroring parameters can be deleted in the CLI. A shut down must be issued on a service level to delete the service. It is not necessary to shut down or remove SAP, or far-end references to delete a remote mirrored service.

To delete a mirror service, the spoke-SDP service has to be deleted from the service. Mirror destinations must be shut down first before they are deleted.

```
Example:*A:ALA-A>config>mirror# mirror-dest 105
config>mirror>mirror-dest# shutdown
config>mirror>mirror-dest# exit
config>mirror# no mirror-dest 105
config>mirror# exit
```

```
*A:ALA-B>config>mirror# mirror-dest 105
config>mirror>mirror-dest# shutdown
config>mirror>mirror-dest# exit
config>mirror# no mirror-dest 105
config>mirror# exit
```

The mirror-destination service ID 105 was removed from the configuration on ALA-A and ALA-B, thus, does not appear in the `info` command output.

```
*A:ALA-A>config>mirror# info
-----
-----
*A:ALA-A>config>mirror# exit
```

```
*A:ALA-B>config>mirror# info
-----
-----
*A:ALA-B>config>mirror# exit
```

Since the mirror destination was removed from the configuration on ALA-B, the port information was automatically removed from the `debug mirror-source` configuration.

```
*A:ALA-B# show debug mirror
debug
exit
*A:ALA-B#
```

Deleting a Remote Mirrored Service

Mirror Service Command Reference

- [Mirror Configuration Commands for 7210 SAS-M, 7210 SAS-T, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE in Network mode on page 51](#)
- [Mirror Configuration Commands for 7210 SAS devices configured in Access-uplink mode on page 51](#)
- [Mirror Configuration Commands for 7210 SAS-X, 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12 on page 52](#)
- [Show Commands on page 53](#)
- [Debug Commands on page 53](#)

Mirror Configuration Commands for 7210 SAS-M, 7210 SAS-T, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE in Network mode

```

config
  — mirror
    — mirror-dest service-id [type mirror-type] [mirror-source-type mirror-source-type] [create]
    — no mirror-dest service-id
      — description description-string
      — no description
      — [no] mirror-dest [fc-name] profile { profile }
        — no remote-source
        — remote-source
        — far-end ip-address [vc-id vc-id] [ing-svc-label ingress-vc-label] tldp
        — no far-end ip-address
      — sap sap-id [create]
      — no sap
      — service-name service-name
      — [no] service-name
      — [no] shutdown
      — no spoke-sdp sdp-id:vc-id
      — spoke-sdp sdp-id:vc-id [create]
        — egress
          — no vc-label [egress-vc-label]
          — vc-label egress-vc-label
          — no shutdown
          — shutdown

```

Mirror Configuration Commands for 7210 SAS devices configured in Access-uplink mode

```

config
  — mirror
    — mirror-dest service-id [type mirror-type] [create]

```

- **no mirror-dest** *service-id*
 - **description** *description-string*
 - **no description**
 - **[no] mirror-dest** [*fc-name*] **profile** { *profile* }
 - **sap** *sap-id* [**create**]
 - **no sap**
 - **service-name** *service-name*
 - **[no]service-name**
 - **[no] shutdown**

Mirror Configuration Commands for 7210 SAS-X, 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

- ```
config
 — mirror
 — mirror-dest service-id [type encap-type][mirror-source-type mirror-source-type] [create]
```
- **no mirror-dest** *service-id*
    - **description** *description-string*
    - **no description**
    - **[no] mirror-dest** [*fc-name*]
    - **no remote-source**
    - **remote-source**
      - **far-end** *ip-address* [*vc-id* *vc-id*] [*ing-svc-label* *ingress-vc-label*] **tldp**
      - **no far-end** *ip-address*
    - **sap** *sap-id* [**create**]
    - **no sap**
      - **[no] egress**
        - **[no] qos** *policy-id*
    - **service-name** *service-name*
    - **[no]service-name**
    - **[no] shutdown**
    - **no spoke-sdp** *sdp-id:vc-id*
    - **spoke-sdp** *sdp-id:vc-id* [**create**]
      - **egress**
        - **no vc-label** [*egress-vc-label*]
        - **vc-label** *egress-vc-label*
        - **no shutdown**
        - **shutdown**

## Show Commands

```
show
— debug [application]
— mirror mirror-dest [service-id]
— service
 — service-using mirror
```

## Debug Commands

```
debug
— [no] mirror-source service-id
 — [no] ip-filter ip-filter-id [entry entry-id]
 — [no] ipv6-filter ipv6-filter-id [entry entry-id]
 — [no] mac-filter mac-filter-id [entry entry-id...]
 — [no] port {port-id | lag lag-id} [egress] [ingress]
 — [no] sap sap-id {[ingress] [egress]}
 — [no] shutdown
```



---

## Configuration Commands

---

### Generic Commands

#### description

**Syntax** **description** *description-string*  
**no description**

**Context** config>mirror>mirror-dest

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

.This command creates a text description stored in the configuration file for a configuration context to help the administrator identify the content of the file.

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

**Default** There is no default description associated with the configuration context.

**Parameters** *description-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>mirror>mirror-dest  
 config>mirror>mirror-dest>spoke-sdp>egress  
 debug>mirror-source

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

The **shutdown** command administratively disables an 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 **shutdown** command administratively disables an entity. 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.

Unlike other commands and parameters where the default state is not indicated in the configuration file, **shutdown** and **no shutdown** are always indicated in system generated configuration files.

## Generic Commands

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

### Default

See Special Cases below.

### Special Cases

**Mirror Destination** — When a mirror destination service ID is shutdown, mirrored packets associated with the service ID are not accepted from the mirror source device. The associated mirror source is put into an operationally down mode. Mirrored packets are not transmitted out of the SAP. Each mirrored packet is silently discarded. If the mirror destination is a SAP, the SAP's discard counters are increased.

The **shutdown** command places the mirror destination service or mirror source into an administratively down state. The **mirror-dest** service ID must be shut down in order to delete the service ID, SAP association from the system.

The default state for a mirror destination service ID is **shutdown**. A **no shutdown** command is required to enable the service.

**Mirror Source** — Mirror sources do not need to be shutdown in order to remove them from the system.

When a mirror source is **shutdown**, mirroring is terminated for all sources defined locally for the **mirror-dest** service ID.

The default state for a mirror source for a given **mirror-dest** service ID is **no shutdown**. A **shutdown** command is required to disable mirroring from that mirror-source.



---

## Mirror Destination Configuration Commands

### mirror-dest

**Syntax** **mirror-dest** *service-id* [**type** *encap-type*] [**mirror-source-type** *mirror-source-type*][**create**]  
**no mirror-dest**

**Context** config>mirror

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command creates a context to set up a service that is intended for packet mirroring. It is configured as a service to allow mirrored packets to be directed locally (within the same device), over the core of the network and have a far end device decode the mirror encapsulation.

The **mirror-dest** service is comprised of destination parameters that define where the mirrored packets are to be sent. It also specifies whether the defined *service-id* will receive mirrored packets from far end devices over the network core.

The **mirror-dest** service IDs are persistent between boots of the router and are included in the configuration backups. The local sources of mirrored packets for the service ID are defined within the **debug mirror mirror-source** command that references the same *service-id*.

The **mirror-dest** command is used to create or edit a service ID for mirroring purposes. If the *service-id* does not exist within the context of all defined services, the **mirror-dest** service is created and the context of the CLI is changed to that service ID. If the *service-id* exists within the context of defined **mirror-dest** services, the CLI context is changed for editing parameters on that service ID. If the *service-id* exists within the context of another service type, an error message is returned and CLI context is not changed from the current context.

The **no** form of the command removes a mirror destination from the system. The **mirror-source** associations with the **mirror-dest** *service-id* do not need to be removed or shutdown first. The **mirror-dest** *service-id* must be shutdown before the service ID can be removed. When the service ID is removed, all **mirror-source** commands that have the service ID defined will also be removed from the system.

**Default** No packet mirroring services are defined.

**Parameters** *service-id* — The service id identifies the service in the service domain. This ID is unique to this service and cannot be used by any other service, regardless of service type. The same service ID must be configured on every device that this particular service is defined on.

If a particular service ID already exists for a service, then the same value cannot be used to create a mirror destination service ID with the same value.

For example:

If an Epipe service-ID **11** exists, then a mirror destination service-ID **11** cannot be created. If a VPLS service-ID **12** exists, then a mirror destination service-ID **12** cannot be created.

## Mirror Destination Configuration Commands

If an IES service-ID **13** exists, then a mirror destination service-ID **13** cannot be created.

**Values** *service-id:* 1 — 2147483647

**type** *encap-type* — The type describes the encapsulation supported by the mirror service.

**Values** ether

*mirror-source-type* — This allows scaling of mirror services that can be used only with remote mirror sources, while limiting the mirror services that can be used by local mirror sources or by both local and remote mirror sources. For more information, see Table 2, Combinations of SAPs, spoke-sdp, and remote sources allowed in a mirror service, on page 21.

**Values** local | remote | both

**local** — indicates that the mirror service can only be used by local mirror sources.

**remote** — indicates that the mirror service can only be used by remote mirror sources.

**both** — indicates that the mirror service can be used by both local and remote mirror sources.

**Default** local

**Values**

### fc

**Syntax** **fc fc-name**  
**no fc**

**Context** config>mirror>mirror-dest

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command specifies a forwarding class for all mirrored copy of the packets transmitted to the destination SAP overriding the default (be) forwarding class. All packets are sent with the same class of service to minimize out of sequence issues. The mirrored copy of the packet does not inherit the forwarding class of the original packet.

When the destination is on a SAP, it pulls buffers from the queue associated with the fc-name and the shaping and scheduling treatment given to the packet is as per the user configuration for that queue.

FC can be assigned only when the mirror source is local. When the mirror source is remote, the network QoS ingress policies that are applied to all the traffic received on the network port and network IP interface are also applied to mirror traffic.

#### NOTES:

- On 7210 SAS-M and 7210 SAS-T, all SAPs configured on a port use the port-based egress queues. If the mirror destination SAP (that is, dot1q SAP or a Q1.\* SAP) is configured to share an uplink with service traffic, mirrored copy of the traffic sent out of the Dot1q or Q1.\* SAP will share the port-based egress

queues with the other service traffic. User is provided an option to assign the profile mirrored copy to the packet, so that during congestion mirrored copy of the packets marked as out-of-profile is dropped before in-profile service traffic (and possibly in-profile mirrored traffic, if user has configured mirrored traffic to be in-profile). The profile is used to determine the slope policy to use for the packet and determines the packet's drop precedence. Additionally, if marking is enabled, it determines the marking value to be used in the packet header.

- On 7210 SAS-X, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE, SAP based egress QoS policy can be used on the mirrored destination SAP, allowing users to control the bandwidth allocated for mirrored traffic.
- On 7210 SAS-Mxp, 7210 SAS-R6 and 7210 SAS-R12, SAP based egress queue QoS policy is used when the "port-based egress queuing is disabled" on the mirrored destination SAP, allowing users to control the amount of bandwidth allocated for mirrored traffic. If port-based queuing is enabled, all SAPs configured on a port use the port-based egress queues.

The no form of the command returns the mirror-dest service ID forwarding class to the default forwarding class.

**Default** The best effort (be) forwarding class is associated with the mirror-dest service ID and profile is out.

**Parameters** *fc-name* — The name of the forwarding class with which to associate mirrored service traffic. The forwarding class name must already be defined within the system. If the *fc-name* does not exist, an error will be returned and the *fc* command will have no effect. If the *fc-name* does exist, the forwarding class associated with *fc-name* will override the default forwarding class.

**Values** be, l2, af, l1, h2, ef, h1, nc

*profile* — The profile to assign to mirrored copy of the service traffic. The profile is used to determine the slope policy to use for the packet and determines the packet's drop precedence. Additionally, if marking is enabled, it determines the marking value to be used in the packet header. A value of in marks the traffic as in-profile traffic and results in use of high slope parameters. A value of out marks the traffic as out-of-profile and results in use of low slope parameters.

**Values** in, out

**Default** out

## far-end

**Syntax** **far-end** *ip-address* [**vc-id** *vc-id*] [**ing-svc-label** *ing-vc-label* | **tldp**]  
**no far-end** *ip-addr*

**Context** config>mirror>mirror-dest>remote-source

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command defines the remote device and configures parameters for mirror destination services on other devices allowed to mirror to the mirror destination service ID.

## Mirror Destination Configuration Commands

The **far-end** command is used within the context of the **remote-source** node. It allows the definition of accepted remote sources for mirrored packets to this *mirror-dest-service-id*. If a far end router has not been specified, packets sent to the router are discarded.

The **far-end** command is used to define a remote source that may send mirrored packets to this 7210 SAS for handling by this **mirror-dest** *service-id*.

The **ing-svc-label** keyword must be given to manually define the expected ingress service label. This ingress label must also be manually defined on the far end address through the **mirror-dest** SDP binding keyword **egr-svc-label**.

The **no** form of the command deletes a far end address from the allowed remote senders to this **mirror-dest** service. All **far-end** addresses are removed when **no remote-source** is executed. All signaled ingress service labels are withdrawn from the far end address affected. All manually defined *ing-svc-label* are removed.

**Default** No far end service ingress addresses are defined.

**Parameters** *ip-address* — The service IP address (system IP address) of the remote device sending mirrored traffic to this mirror destination service. If 0.0.0.0 is specified, any remote is allowed to send to this service.

**Values** 1.0.0.1 — 223.255.255.254

The ingress service label must be manually defined using the **ing-svc-label** keyword. On the far end 7210 SAS, the associated SDP **egr-svc-label** must be manually set and equal to the label defined in **ing-svc-label**.

**vc-id** *vc-id* — Specifies the virtual circuit identifier.

**Values** 1 — 4294967295

**ing-svc-label** *ing-vc-label* — Specifies the ingress service label for mirrored service traffic on the **far end** device for manually configured mirror service labels.

The defined *ing-svc-label* is entered into the ingress service label table which causes ingress packet with that service label to be handled by this **mirror-dest** service.

The specified *ing-svc-label* must not have been used for any other service ID and must match the far end expected specific *egr-svc-label* for this 7210 SAS. It must be within the range specified for manually configured service labels defined on this 7210 SAS. It may be reused for other far end addresses on this *mirror-dest-service-id*.

**Values** 2048 — 18431

**tl dp** — Specifies that the label is obtained through signaling via the LDP.

### remote-source

**Syntax** [no] **remote-source**

**Context** config>mirror>mirror-dest

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command configures remote devices to mirror traffic to this device for mirror service egress. Optionally, deletes all previously defined remote mirror ingress devices.

The remote-source context allows the creation of a ‘sniffer farm’ to consolidate expensive packet capture and diagnostic tools to a central location. Remote areas of the access network can be monitored via normal service provisioning techniques.

Specific far-end routers can be specified with the **far-end** command allowing them to use this router as the destination for the same *mirror-dest-service-id*.

The **remote-source** node allows the source of mirrored packets to be on remote 7210 SAS devices. The local 7210 SAS will configure its network ports to forward packets associated with the *service-id* to the destination SAP. When **remote-source far-end** addresses are configured, an SDP is not allowed as a destination.

By default, the **remote-source** context contains no **far-end** addresses. When no **far-end** addresses have been specified, network remote devices will not be allowed to mirror packets to the local 7210 SAS as a mirror destination. Packets received from unspecified **far-end** addresses will be discarded at network ingress.

The **no** form of the command restores the *service-id* to the default condition to not allow a remote 7210 SAS access to the mirror destination. The **far-end** addresses are removed without warning.

**Default** No remote source devices defined

## sap

**Syntax** **sap sap-id [create]**  
**no sap**

**Context** config>mirror>mirror-dest

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command creates a service access point (SAP) within a mirror destination service. The SAP is owned by the mirror destination service ID.

The SAP is defined with port and encapsulation parameters to uniquely identify the (mirror) SAP on the interface and within the box. The specified SAP must define an Ethernet port with a null SAP or a Dot1q SAP or a Q1.\* SAP.

Only one SAP can be created within a **mirror-dest** service ID. If the defined SAP has not been created on any service within the system, the SAP is created and the context of the CLI will change to the newly created SAP. In addition, the port cannot be a member of a multi-link bundle, LAG, APS group or IMA bundle.

If the defined SAP exists in the context of another service ID, **mirror-dest** or any other type, an error is generated.

Mirror destination SAPs can be created on Ethernet interfaces that have been defined as an access port or access-uplink port. If the interface is defined as network, the SAP creation returns an error.

**NOTE:** When using a dot1q SAP or a Q1.\* SAP as a mirror destination, users must allocated resources of another port for use by this features. Please refer the mirror configuration notes above [Configuration Notes on page 28](#).

## Mirror Destination Configuration Commands

When the **no** form of this command is used on a SAP created by a mirror destination service ID, the SAP with the specified port and encapsulation parameters is deleted.

**Default** No default SAP for the mirror destination service defined.

**Parameters** *sap-id* — Specifies the physical port identifier portion of the SAP definition. See [Common CLI Command Descriptions on page 501](#) for command syntax.

### service-name

**Syntax** **service-name** *service-name*  
**no service-name**

**Context** config>mirror>mirror-dest

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command configures an optional service name, up to 64 characters in length, which adds a name identifier to a given service to then use that service name in configuration references as well as display and use service names in show commands throughout the system. This helps the service provider/administrator to identify and manage services within the 7210 SAS platforms.

All services are required to assign a service ID to initially create a service. However, either the service ID or the service name can be used to identify and reference a given service once it is initially created.

**Parameters** *service-name* — Specifies a unique service name to identify the service. Service names may not begin with an integer (0-9).

### spoke-sdp

**Syntax** **spoke-sdp** *sdp-id:vc-id* [**create**] [**no-endpoint**]  
**spoke-sdp** *sdp-id:vc-id* [**create**] **endpoint** *name*  
**no sdp** *sdp-id:vc-id*

**Context** config>mirror>mirror-dest

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command binds an existing (mirror) service distribution path (SDP) to the mirror destination service ID.

The operational state of the SDP dictates the operational state of the SDP binding to the mirror destination. If the SDP is shutdown or operationally down, then SDP binding is down. Once the binding is defined and the service and SDP are operational, the far-end router defined in the **config service sdp** *sdp-id* **far-end** parameter is considered part of the service ID.

Only one SDP can be associated with a mirror destination service ID. If a second **sdp** command is executed after a successful SDP binding, an error occurs and the command has no effect on the existing configuration. A **no sdp** command must be issued before a new SDP binding can be attempted.

An SDP is a logical mechanism that ties a far end router to a specific service without having to define the far-end SAP. Each SDP represents a method to reach a router.

The other method is Multi-Protocol Label Switching (MPLS) encapsulation. router routers support both signaled and non-signaled LSPs (Label Switched Path) though the network. Non-signaled paths are defined at each hop through the network. Signaled paths are protocol communicated from end to end using RSVP. Paths may be manually defined or a constraint based routing protocol (OSPF-TE or CSPF) can be used to determine the best path with specific constraints.

SDPs are created and then bound to services. Many services can be bound to a single SDP. The operational and administrative state of the SDP controls the state of the SDP binding to the service.

An egress service label (Martini VC-Label), used by the SDP to differentiate each service bound to the SDP to the far-end router, must be obtained manually or through signaling with the far end. If manually configured, it must match the **ing-svc-label** defined for the local router.

**NOTE:** When using remote mirroring with spoke-SDP configured as a mirror destination, users must allocated resources of another port for use by this features. Please refer the mirror configuration notes above [Configuration Notes on page 28](#).

The **no** form of the command removes the SDP binding from the mirror destination service. Once removed, no packets are forwarded to the far-end (destination) router from that mirror destination service ID.

**Default** No default SDP ID is bound to a mirror destination service ID. If no SDP is bound to the service, the mirror destination will be local and cannot be to another router over the core network.

**Parameters** *sdp-id[:vc-id]* — A locally unique SDP identification (ID) number. The SDP ID must exist. If the SDP ID does not exist, an error will occur and the command will not execute.

For mirror services, the *vc-id* defaults to the *service-id*. However, there are scenarios where the *vc-id* is being used by another service. In this case, the SDP binding cannot be created. So, to avoid this, the mirror service SDP bindings now accepts *vc-ids*.

**Values** 1 — 17407

*vc-id* — Specifies the virtual circuit identifier.

**Values** 1 — 4294967295

**endpoint name** — specifies the name of the endpoint associated with the SAP.

**no endpoint** — Removes the association of a SAP or a SDP with an explicit endpoint name.

## egress

**Syntax** **egress**

**Context** config>mirror>mirror-dest>spoke-sdp

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command enters the context to configure spoke SDP egress parameters.

## Mirror Destination Configuration Commands

### vc-label

|                    |                                                                                                                                                                                                                    |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>vc-label</b> <i>egress-vc-label</i><br><b>no vc-label</b> [ <i>egress-vc-label</i> ]                                                                                                                            |
| <b>Context</b>     | config>mirror>mirror-dest>spoke-sdp>egress                                                                                                                                                                         |
| <b>Description</b> | <b>Platforms Supported:</b> 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.<br>This command configures the spoke-SDP egress VC label. |
| <b>Parameters</b>  | <i>egress-vc-label</i> — A VC egress value that indicates a specific connection.<br><b>Values</b> 16 — 1048575                                                                                                     |

### egress

|                    |                                                                                                                                                                                                                                              |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>egress</b>                                                                                                                                                                                                                                |
| <b>Context</b>     | config> mirror> sap                                                                                                                                                                                                                          |
| <b>Description</b> | <b>Platforms Supported:</b> 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.<br>This command provides the context to configure QoS egress policies for this SAP. |

### qos

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no]</b> qos <i>policy-id</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Context</b>     | config> mirror> sap> egress                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b> | <b>Platforms Supported:</b> 7210 SAS-Mxp, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.<br>This command allows user to configure the QoS policy for the mirror destination SAP egress. The SAP egress QoS policy to use is specified using the <i>policy-id</i> > and must have been configured before associating this policy with the SAP. The SAP egress policy can be configured using the commands under the context config> qos> sap-egress.<br>When a SAP egress policy is associated with the SAP configured as a mirror destination, the queue associated with FC specified with the CLI command config> mirror> mirror-dest> fc, is used for traffic sent out of the mirror destination SAP. The policy allows the user to specify the amount of buffers, the WRED policy, the shaping rate and the marking values to use for the mirrored copy.<br>The no form of the command associates the default SAP egress QoS policy with the SAP. |
| <b>Default</b>     | no qos                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |



**Parameters**    *policy-id* — Specifies the QoS policy to associated with SAP egress. The QoS policy referred to by the *policy-id* is configured using the commands under `config> qos> sap-egress`.

---

## Mirror Source Configuration Commands

### mirror-source

**Syntax** [no] **mirror-source** *service-id*

**Context** debug

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command configures mirror source parameters for a mirrored service.

The **mirror-source** command is used to enable mirroring of packets specified by the association of the **mirror-source** to sources of packets defined within the context of the *mirror-dest-service-id*. The mirror destination service must already exist within the system.

A mirrored packet cannot be mirrored to multiple destinations. If a mirrored packet is properly referenced by multiple mirror sources (for example, a SAP on one **mirror-source** and a port on another **mirror-source**), then the packet is mirrored to a single *mirror-dest-service-id* based on the following hierarchy:

1. Filter entry
2. Service access port (SAP)
3. Physical port

The hierarchy is structured so the most specific match criteria has precedence over a less specific match. For example, if a **mirror-source** defines a port and a SAP on that port, then the SAP mirror-source is accepted and the mirror-source for the port is ignored because of the hierarchical order of precedence.

The **mirror-source** configuration is not saved when a configuration is saved. A **mirror-source** manually configured within an ASCII configuration file will not be preserved if that file is overwritten by a **save** command. Define the **mirror-source** within a file associated with a **config exec** command to make a **mirror-source** persistent between system reboots.

By default, all **mirror-dest** service IDs have a **mirror-source** associated with them. The **mirror-source** is not technically created with this command. Instead the service ID provides a contextual node for storing the current mirroring sources for the associated **mirror-dest** service ID. The **mirror-source** is created for the mirror service when the operator enters the **debug>mirror-source svcId** for the first time. The **mirror-source** is also automatically removed when the **mirror-dest** service ID is deleted from the system.

The **no** form of the command deletes all related source commands within the context of the **mirror-source** *service-id*. The command does not remove the service ID from the system.

**Default** No mirror source match criteria is defined for the mirror destination service.

**Parameters** *service-id* — The mirror destination service ID for which match criteria will be defined. The *service-id* must already exist within the system.

**Values** *service-id:* 1 — 2147483647

## ip-filter

**Syntax** **ip-filter** *ip-filter-id* **entry** *entry-id* [*entry-id* ...]  
**no ip-filter** *ip-filter-id* **entry** *entry-id*

**Context** debug>mirror-source

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command enables mirroring of packets that match specific entries in an existing IP filter.

The **ip-filter** command directs packets which match the defined list of entry IDs to be mirrored to the mirror destination referenced by the *mirror-dest-service-id* of the **mirror-source**.

The IP filter must already exist in order for the command to execute. Filters are configured in the **config>filter** context. If the IP filter does not exist, an error will occur. If the filter exists but has not been associated with a SAP or IP interface, an error is not generated but mirroring will not be enabled (there are no packets to mirror). Once the IP filter is defined to a SAP or IP interface, mirroring is enabled.

If the IP filter is defined as ingress, only ingress packets are mirrored. Ingress mirrored packets are mirrored to the mirror destination prior to any ingress packet modifications.

If the IP filter is defined as egress, only egress packets are mirrored. Egress mirrored packets are mirrored to the mirror destination after all egress packet modifications.

An *entry-id* within an IP filter can only be mirrored to a single mirror destination. If the same *entry-id* is defined multiple times, an error occurs and only the first **mirror-source** definition is in effect.

By default, no packets matching any IP filters are mirrored. Mirroring of IP filter entries must be explicitly defined.

The **no ip-filter** command, without the **entry** keyword, removes mirroring on all *entry-id*'s within the *ip-filter-id*.

When the **no** command is executed with the **entry** keyword and one or more *entry-id*'s, mirroring of that list of *entry-id*'s is terminated within the *ip-filter-id*. If an *entry-id* is listed that does not exist, an error will occur and the command will not execute. If an *entry-id* is listed that is not currently being mirrored, no error will occur for that *entry-id* and the command will execute normally.

**Default** IP filter mirroring is not defined.

**Parameters** *ip-filter-id* — The IP filter ID whose entries are mirrored. If the *ip-filter-id* does not exist, an error will occur and the command will not execute. Mirroring of packets will commence once the *ip-filter-id* is defined on a SAP or IP interface.

**Values** 1 - 65535

**entry** *entry-id* [*entry-id* ...] — The IP filter entries to use as match criteria for packet mirroring. The **entry** keyword begins a list of *entry-id*'s for mirroring. Multiple *entry-id* entries may be specified with a single command. Each *entry-id* must be separated by a space.

If an *entry-id* does not exist within the IP filter, an error occurs and the command will not execute.

If the filter's *entry-id* is renumbered within the IP filter definition, the old *entry-id* is removed but the new *entry-id* must be manually added to the configuration to include the new (renumbered) entry's cri-

## Mirror Source Configuration Commands

teria.

**Values** 1 - 65535

### ipv6-filter

**Syntax** **ipv6-filter** *ip-filter-id* **entry** *entry-id* [*entry-id* ...]  
**no ipv6-filter** *ip-filter-id* **entry** *entry-id*

**Context** debug>mirror-source

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command enables mirroring of packets that match specific entries in an existing IPv6 filter.

The **ipv6-filter** command directs packets which match the defined list of entry IDs to be mirrored to the mirror destination referenced by the *mirror-dest-service-id* of the **mirror-source**.

The IPv6 filter must already exist in order for the command to execute. Filters are configured in the **config>filter** context. If the IPv6 filter does not exist, an error will occur. If the filter exists but has not been associated with a SAP or IP interface, an error is not generated but mirroring will not be enabled (there are no packets to mirror). Once the IPv6 filter is defined to a SAP or IP interface, mirroring is enabled.

If the IPv6 filter is defined as ingress, only ingress packets are mirrored. Ingress mirrored packets are mirrored to the mirror destination prior to any ingress packet modifications.

If the IPv6 filter is defined as egress, only egress packets are mirrored. Egress mirrored packets are mirrored to the mirror destination after all egress packet modifications.

An *entry-id* within an IPv6 filter can only be mirrored to a single mirror destination. If the same *entry-id* is defined multiple times, an error occurs and only the first **mirror-source** definition is in effect.

By default, no packets matching any IPv6 filters are mirrored. Mirroring of IPv6 filter entries must be explicitly defined.

The **no ipv6-filter** command, without the **entry** keyword, removes mirroring on all *entry-id*'s within the *ipv6-filter-id*.

When the **no** command is executed with the **entry** keyword and one or more *entry-id*'s, mirroring of that list of *entry-id*'s is terminated within the *ipv6-filter-id*. If an *entry-id* is listed that does not exist, an error will occur and the command will not execute. If an *entry-id* is listed that is not currently being mirrored, no error will occur for that *entry-id* and the command will execute normally.

**Default** IPv6 filter mirroring is not defined.

**Parameters** *ipv6-filter-id* — The IPv6 filter ID whose entries are mirrored. If the *ipv6-filter-id* does not exist, an error will occur and the command will not execute. Mirroring of packets will commence once the *ipv6-filter-id* is defined on a SAP or IP interface.

**Values** 1 - 65535

**entry** *entry-id* [*entry-id* ...] — The IPv6 filter entries to use as match criteria for packet mirroring. The **entry** keyword begins a list of *entry-id*'s for mirroring. Multiple *entry-id* entries may be specified with a single

command. Each *entry-id* must be separated by a space.

If an *entry-id* does not exist within the IPv6 filter, an error occurs and the command will not execute.

If the filter's *entry-id* is renumbered within the IPv6 filter definition, the old *entry-id* is removed but the new *entry-id* must be manually added to the configuration to include the new (renumbered) entry's criteria.

**Values** 1 - 65535

## mac-filter

**Syntax** **mac-filter** *mac-filter-id* **entry** *entry-id* [*entry-id* ...]  
**no mac-filter** *mac-filter-id*  
**no mac-filter** *mac-filter-id* **entry** *entry-id* [*entry-id* ...]

**Context** debug>mirror-source

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command enables mirroring of packets that match specific entries in an existing MAC filter.

The **mac-filter** command directs packets which match the defined list of entry IDs to be mirrored to the mirror destination referenced by the *mirror-dest-service-id* of the **mirror-source**.

The MAC filter must already exist in order for the command to execute. Filters are configured in the config>filter context. If the MAC filter does not exist, an error will occur. If the filter exists but has not been associated with a SAP or IP interface, an error is not generated but mirroring will not be enabled (there are no packets to mirror). Once the filter is defined to a SAP or MAC interface, mirroring is enabled.

If the MAC filter is defined as ingress, only ingress packets are mirrored. Ingress mirrored packets are mirrored to the mirror destination prior to any ingress packet modifications.

The **no mac-filter** command, without the **entry** keyword, removes mirroring on all *entry-id*'s within the *mac-filter-id*.

When the **no** command is executed with the **entry** keyword and one or more *entry-id*'s, mirroring of that list of *entry-id*'s is terminated within the *mac-filter-id*. If an *entry-id* is listed that does not exist, an error will occur and the command will not execute. If an *entry-id* is listed that is not currently being mirrored, no error will occur for that *entry-id* and the command will execute normally.

**Default** No MAC filter mirroring defined.

**Parameters** *mac-filter-id* — The MAC filter ID whose entries are mirrored. If the *mac-filter-id* does not exist, an error will occur and the command will not execute. Mirroring of packets will commence once the *mac-filter-id* is defined on a SAP.

**Values** 1 - 65535

**entry** *entry-id* [*entry-id* ...] — The MAC filter entries to use as match criteria for packet mirroring. The **entry** keyword begins a list of *entry-id*'s for mirroring. Multiple *entry-id* entries may be specified with a single command. Each *entry-id* must be separated by a space. Up to 8 entry IDs may be specified in a

## Mirror Source Configuration Commands

single command.

Each *entry-id* must exist within the *mac-filter-id*. If the *entry-id* is renumbered within the MAC filter definition, the old *entry-id* is removed from the list and the new *entry-id* will need to be manually added to the list if mirroring is still desired.

If no *entry-id* entries are specified in the command, mirroring will not occur for that MAC filter ID. The command will have no effect.

**Values** 1 - 65535

### port

**Syntax** **port** {*port-id* | **lag** *lag-id*} {[**egress**] [**ingress**]}  
**no port** {*port-id* | **lag** *lag-id*} [**egress**] [**ingress**]

**Context** debug>mirror-source

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command enables mirroring of traffic ingressing or egressing a port (Ethernet port, or Link Aggregation Group (LAG)).

The **port** command associates a port or LAG to a mirror source. The port is identified by the *port-id*. The defined port may be Ethernet, access or access uplink. access. A port may be a single port or a Link Aggregation Group (LAG) ID. When a LAG ID is given as the *port-id*, mirroring is enabled on all ports making up the LAG. Either a LAG port member *or* the LAG port can be mirrored.

The port is only referenced in the mirror source for mirroring purposes. If the port is removed from the system, the mirroring association will be removed from the mirror source.

The same port may not be associated with multiple mirror source definitions with the **ingress** parameter defined. The same port may not be associated with multiple mirror source definitions with the **egress** parameter defined.

If a SAP is mirrored on an access port, the SAP mirroring will have precedence over the access port mirroring when a packet matches the SAP mirroring criteria. Filter and label mirroring destinations will also precedence over a port-mirroring destination.

If the port is not associated with a **mirror-source**, packets on that port will not be mirrored. Mirroring may still be defined for a SAP or filter entry, which will mirror based on a more specific criteria.

The **no port** command disables port mirroring for the specified port. Mirroring of packets on the port may continue due to more specific mirror criteria. If the **egress** or **ingress** parameter keywords are specified in the **no** command, only the ingress or egress mirroring condition will be removed.

**Default** No ports are defined.

**Parameters** *port-id* — Specifies the port ID.

**Values** SAS-M,X,T-[1..26]  
SAS-Mxp - [1..28]

SAS-R6 and SAS-R12- Depends on the type of IMM card used. For more information about the number of ports supported on different IMM cards see, Systems Basic User Guide.  
 SAS-Sx/S 1/10GE- For 24 port variant - 1/1/25 and 1/1/26  
 For 48 port variant - 1/1/49 and 1/1/50.  
 For SAS-Sx 10/100GE port variant - 1/1/68

*lag-id* — The LAG identifier, expressed as a decimal integer.

**Values** SAS-M,X,T, Mxp, Sx -[1..25]  
 SAS-R6 and SAS-R12- [1..63]  
 SAS-Sx 1/10GE and SAS-Sx 10/100GE- [1..56]

**egress** — Specifies that packets egressing the port should be mirrored. Egress packets are mirrored to the mirror destination after egress packet modification.

**ingress** — Specifies that packets ingressing the port should be mirrored. Ingress packets are mirrored to the mirror destination prior to ingress packet modification.

## sap

**Syntax** **no sap** *sap-id* {[ingress] }  
**no sap** *sap-id* {[ingress] }

**Context** debug>mirror-source

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command enables mirroring of traffic ingressing or egressing a service access port (SAP). A SAP that is defined within a mirror destination cannot be used in a mirror source. The mirror source SAP referenced by the *sap-id* is owned by the service ID of the service in which it was created. The SAP is only referenced in the mirror source name for mirroring purposes. The mirror source association does not need to be removed before deleting the SAP from its service ID. If the SAP is deleted from its service ID, the mirror association is removed from the mirror source.

More than one SAP can be associated within a single **mirror-source**. Each SAP has its own **ingress** parameter keywords to define which packets are mirrored to the mirror destination.

The SAP must be valid and properly configured. If the associated SAP does not exist, an error occurs and the command will not execute.

The same SAP cannot be associated with multiple mirror source definitions for ingress packets.

If a particular SAP is not associated with a mirror source name, then that SAP will not have mirroring enabled for that mirror source.

The **no** form of the command disables mirroring for the specified SAP. All mirroring for that SAP on ingress and egress is terminated. Mirroring of packets on the SAP can continue if more specific mirror criteria is configured. If the **egress** or **ingress** parameter keywords are specified in the **no** command, only the ingress or egress mirroring condition is removed.

**Default** No SAPs are defined by default.

## Mirror Source Configuration Commands

- Parameters**
- sap-id* — Specifies the physical port identifier portion of the SAP definition. See [Common CLI Command Descriptions on page 501](#) for command syntax.
  - ingress** — Specifies that packets ingressing the SAP should be mirrored. Ingress packets are mirrored to the mirror destination prior to ingress packet modification.



---

## Show Commands

### debug

**Syntax** `debug [application]`

**Context** `show`

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command displays set debug points.

**Parameters** *application* — Display which debug points have been set.

**Values:** service, ip, ospf, ospf3, mtrace, isis, mpls, rsvp, ldp, mirror, system, filter, subscriber-mgmt, radius, lag, oam

**Output**

```
*A:alul# show debug
debug
 mirror-source 101
 port 1/1/1 ingress
 no shutdown
 exit
 mirror-source 102
 port 1/1/3 egress
 no shutdown
 exit
exit
*A:alul#
```

### service-using

**Syntax** `service-using [mirror]`

**Context** `show>service`

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

Displays mirror services.

If no optional parameters are specified, all services defined on the system are displayed.

**Parameters** **mirror** — Displays mirror services.

**Output** **Show Service-Using Mirror** — The following table describes service-using mirror output fields:

| Label            | Description                                                                       |
|------------------|-----------------------------------------------------------------------------------|
| Service Id       | The service identifier.                                                           |
| Type             | Specifies the service type configured for the service ID.                         |
| Adm              | The desired state of the service.                                                 |
| Opr              | The operating state of the service.                                               |
| CustomerID       | The ID of the customer who owns this service.                                     |
| Last Mgmt Change | The date and time of the most recent management-initiated change to this service. |

### Sample Output

```

A:ALA-48# show service service-using mirror
=====
Services [mirror]
=====
ServiceId Type Adm Opr CustomerId Last Mgmt Change

218 Mirror Up Down 1 04/08/2007 13:49:57
318 Mirror Down Down 1 04/08/2007 13:49:57
319 Mirror Up Down 1 04/08/2007 13:49:57
320 Mirror Up Down 1 04/08/2007 13:49:57
1000 Mirror Down Down 1 04/08/2007 13:49:57
1216 Mirror Up Down 1 04/08/2007 13:49:57
1412412 Mirror Down Down 1 04/08/2007 13:49:57

Matching Services : 7
=====
A:ALA-48#

```

## mirror

**Syntax** `mirror mirror-dest service-id`

**Context** show

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command displays mirror configuration and operation information.

**Parameters** *service-id* — Specify the mirror service ID.

**Values** [1..2147483648] | svc-name:64 char max

**Output** **Mirroring Output** — The following table describes the mirroring output fields:

| Label                  | Description                                                                                                                   |
|------------------------|-------------------------------------------------------------------------------------------------------------------------------|
| Service Id             | The service ID associated with this mirror destination.                                                                       |
| Type                   | Entries in this table have an implied storage type of “volatile”. The configured mirror source information is not persistent. |
| Admin State            | Up — The mirror destination is administratively enabled.<br>Down — The mirror destination is administratively disabled.       |
| Oper State             | Up — The mirror destination is operationally enabled.<br>Down — The mirror destination is operationally disabled.             |
| Forwarding Class       | The forwarding class for all packets transmitted to the mirror destination.                                                   |
| Remote Sources         | Yes — A remote source is configured.<br>No — A remote source is not configured.                                               |
| Destination SAP        | The ID of the access port where the Service Access Point (SAP) associated with this mirror destination service is defined.    |
| Egr QoS Policy         | This value indicates the egress QoS policy ID. A value of 0 indicates that no QoS policy is specified.                        |
| mirror sources allowed | This value tells the user the type of mirror sources allowed to be configured.                                                |

### Sample Output

```
*A:7210SAS>config>mirror>mirror-dest$ show mirror mirror-dest
```

```

=====
Mirror Services
=====
Id Type Adm Opr Destination SDP Lbl/
 SAP QoS
 Src
 Allowed

1 Ether Down Down None n/a 0 Local
1000 Ether Up Down SDP 400 (1.1.1.1) Pending 0 Local
2000 Ether Up Up SAP 1/1/17:17 1 0 Remote
=====
*A:7210SAS>config>mirror>mirror-dest$

```

### Sample Output for Network mode

```

*A:7210SAS>config>mirror>mirror-dest$ show mirror mirror-dest 1

=====
Mirror Service
=====
Service Id : 1 Type : Ether
Description : (Not Specified)
Admin State : Down Oper State : Down
Mirror Sources Allowed : Local
Forwarding Class : be Remote Sources: No
Profile : out

=====
Mirror Services SDP
=====
SdpId IP Addr CfgLabel Signal EgrLabel

No Matching Entries
=====

Local Sources

Admin State : Up

No Mirror Sources configured
=====
*A:7210SAS>config>mirror>mirror-dest$

```

### Sample Output for Access-Uplink mode

```

*A:7210SAS# show mirror mirror-dest 1000

=====
Mirror Service
=====
Service Id : 1000 Type : Ether
Description : (Not Specified)
Admin State : Up Oper State : Down

```

Mirror Sources Allowed : Local  
Profile : out  
Destination SAP : 1/1/1

-----  
Local Sources  
-----

Admin State : Up

-Port 1/1/1 Ing

=====  
\*A:7210SAS#



# OAM and SAA

---

## In This Chapter

This chapter provides information about the Operations, Administration and Management (OAM) and Service Assurance Agent (SAA) commands available in the CLI for troubleshooting services.

Topics in this chapter include:

- [OAM Overview on page 81](#)
  - [LSP Diagnostics: LSP Ping and Trace on page 82](#)
  - [SDP Diagnostics on page 98](#)
  - [Service Diagnostics on page 100](#)
  - [VPLS MAC Diagnostics on page 101](#)
  - [VLL Diagnostics on page 106](#)
- [IP Performance Monitoring \(IP PM\) on page 131](#)
  - [Two-Way Active Measurement Protocol \(TWAMP\) on page 131](#)
  - [Two-Way Active Measurement Protocol Light \(TWAMP Light\) on page 132](#)
- [Ethernet Connectivity Fault Management \(ETH-CFM\) on page 136](#)
- [Synthetic Loss Measurement \(ETH-SL\) on page 161](#)

## In This Chapter

- [Service Assurance Agent Overview on page 83](#)
  - [SAA Application on page 209](#)
- [Y.1564 Testhead OAM tool on page 179](#)
- [OAM Performance Monitoring \(OAM-PM\) on page 192](#)
  - [Session on page 194](#)
  - [Standard PM Packets on page 194](#)
  - [Measurement Intervals on page 195](#)
  - [Data Structures and Storage on page 201](#)
  - [Bin Groups on page 204](#)
  - [Relating the Components on page 205](#)
  - [Monitoring on page 206](#)



## OAM Overview

Delivery of services requires a number of operations occur properly and at different levels in the service delivery model. For example, operations such as the association of packets to a service, must be performed properly in the forwarding plane for the service to function properly. In order to verify that a service is operational, a set of in-band, packet-based Operation, Administration, and Maintenance (OAM) tools is required, with the ability to test each of the individual packet operations.

For in-band testing, the OAM packets closely resemble customer packets to effectively test the customer's forwarding path, but they are distinguishable from customer packets so they are kept within the service provider's network and not forwarded to the customer.

The suite of OAM diagnostics supplement the basic IP ping and traceroute operations with diagnostics specialized for the different levels in the service delivery model. There are diagnostics for services.

**NOTE:** The following OAM features are not available on 7210 SAS-M and 7210 SAS-T devices configured in Access uplink mode:

- LSP Diagnostics
  - SDP Diagnostics
  - Service Diagnostics
  - VPLS MAC Diagnostics
  - VLL Diagnostics
-

## LSP Diagnostics: LSP Ping and Trace

In the following paragraphs the discussion is generalized. Users need to consider the following restrictions while reading through it:

- On 7210, LDP LER ECMP is not supported. The description below on LDP ECMP as an LER node does not apply to 7210 nodes.
- LDP LSR ECMP is supported only on 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-R6 and 7210 SAS-R12. The description on LDP ECMP as LSR node applies only to these nodes.
- 7210 nodes support LDP and BGP 3107 labeled routes. The 7210 does not support LDP FEC stitching to BGP 3107 labeled route to LDP FEC stitching and vice-versa. The description below about stitching of BGP 3107 labeled routes to LDP FEC is provided to explain the behavior in an end-to-end network solution deployed using 7210 and 7x50 nodes, with 7210 nodes acting as the LER node.
- 7210 nodes do not support ECMP for BGP 3107 labeled routes. The description on use of ECMP routes for BGP 3107 labeled routes does not apply to 7210. It is provided below for completeness and better understanding of the behavior in end-to-end network solution deployed using 7210 and 7750 nodes.

The router LSP diagnostics are implementations of LSP ping and LSP trace based on RFC 4379, Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures. LSP ping provides a mechanism to detect data plane failures in MPLS LSPs. LSP ping and LSP trace are modeled after the ICMP echo request/reply used by ping and trace to detect and localize faults in IP networks.

For a given LDP FEC, RSVP P2P LSP, or BGP IPv4 Label Router, LSP ping verifies whether the packet reaches the egress label edge router (LER), while in LSP trace mode, the packet is sent to the control plane of each transit label switched router (LSR) which performs various checks to see if it is actually a transit LSR for the path.

The downstream mapping TLV is used in `lsp-ping` and `lsp-trace` to provide a mechanism for the sender and responder nodes to exchange and validate interface and label stack information for each downstream of an LDP FEC or an RSVP LSP and at each hop in the path of the LDP FEC or RSVP LSP.


Two downstream mapping TLVs are supported. The original Downstream Mapping (DSMAP) TLV defined in RFC 4379 and the new Downstream Detailed Mapping (DDMAP) TLV defined in RFC 6424.

When the responder node has multiple equal cost next-hops for an LDP FEC prefix, the downstream mapping TLV can further be used to exercise a specific path of the ECMP set using the path-destination option. The behavior in this case is described in the ECMP sub-section below.

## LSP Ping/Trace for an LSP Using a BGP IPv4 Label Route

This feature adds support of the target FEC stack TLV of type BGP Labeled IPv4 /32 Prefix as defined in RFC 4379.

The new TLV is structured as follows:



**Figure 7: Target FEC Stack TLV for a BGP Labeled IPv4 Prefix**

The user issues a LSP ping using the existing CLI command and specifying a new type of prefix:

```
oam lsp-ping bgp-label prefix ip-prefix/mask [src-ip-address ip-address] [fc fc-name]
[size octets] [ttl label-ttl] [send-count send-count] [timeout timeout] [interval interval]
[path-destination ip-address [interface if-name | next-hop ip-address]] [detail]
```

The path-destination option is used for exercising specific ECMP paths in the network when the LSR performs hashing on the MPLS packet.

Similarly, the user issues a LSP trace using the following command:

```
oam lsp-trace bgp-label prefix ip-prefix/mask [src-ip-address ip-address] [fc fc-name]
[max-fail no-response-count] [probe-count probes-per-hop] [size octets] [min-ttl min-label-ttl]
[max-ttl max-label-ttl] [timeout timeout] [interval interval] [path-destination ip-address
[interface if-name | next-hop ip-address]] [detail]
```

The following are the procedures for sending and responding to an LSP ping or LSP trace packet. These procedures are valid when the downstream mapping is set to the DSMAP TLV. The detailed procedures with the DDMAP TLV are presented in [Using DDMAP TLV in LSP Stitching and LSP Hierarchy](#).

1. The next-hop of a BGP label route for a core IPv4 /32 prefix is always resolved to an LDP FEC or an RSVP LSP. Thus the sender node encapsulates the packet of the echo request message with a label stack which consists of the LDP/RSVP outer label and the BGP inner label.

If the packet expires on an RSVP or LDP LSR node which does not have context for the BGP label IPv4 /32 prefix, it validates the outer label in the stack and if the validation is successful it replies the same way as it does today when it receives an echo request message for an LDP FEC which is stitched to a BGP IPv4 label route. In other words it replies with return code 8 Label switched at stack-depth <RSC>.

2. An LSR node which is the next-hop for the BGP label IPv4 /32 prefix as well as the LER node which originated the BGP label IPv4 prefix have full context for the BGP IPv4 target FEC stack and can thus perform full validation of it.
3. If the BGP IPv4 label route is stitched to an LDP FEC, the egress LER for the resulting LDP FEC will not have context for the BGP IPv4 target FEC stack in the echo request message and replies with return code 4 Replying router has no mapping for the FEC at stack- depth <RSC>. This is the same behavior as that of an LDP FEC which is stitched to a BGP IPv4 label route when the echo request message reaches the egress LER for the BGP prefix.

Note that only BGP label IPv4 /32 prefixes are supported since these are usable as tunnels on nodes. BGP label IPv6 /128 prefixes are not currently usable as tunnels on a node and as such are not supported in LSP ping/trace.

## ECMP Considerations

When the responder node has multiple equal cost next-hops for an LDP FEC or a BGP label IPv4 prefix, it replies in the Downstream Mapping TLV with the downstream information of the outgoing interface which is part of the ECMP next-hop set for the prefix.

Note however that when BGP label route is resolved to an LDP FEC (of the BGP next-hop of the BGP label route), ECMP can exist at both the BGP and LDP levels. The following selection of next-hop is performed in this case:

1. For each BGP ECMP next-hop of the label route, a single LDP next-hop is selected even if multiple LDP ECMP next-hops exist. Thus, the number of ECMP next-hops for the BGP IPv4 label route will be equal to the number of BGP next-hops.
2. ECMP for a BGP IPv4 label route is only supported at PE router (BGP label push operation) and not at ABR/ASBR (BGP label swap operation). Thus at an LSR, a BGP IPv4 label route will be resolved to a single BGP next-hop which itself is resolved to a single LDP next-hop.
3. LSP trace will return one downstream mapping TLV for each next-hop of the BGP IPv4 label route. Furthermore, it will return exactly the LDP next-hop the data path programmed for each BGP next-hop.

The following description of the behavior of LSP ping and LSP trace makes a reference to a FEC in a generic way and which can represent an LDP FEC or a BGP IPv4 label route. In addition the reference to a downstream mapping TLV means either the DSMAP TLV or the DDMAP TLV.

1. If the user initiates an `lsp-trace` of the FEC without the **path-destination** option specified, then the sender node will not include multi-path information in the Downstream Mapping TLV in the echo request message (multipath type=0). In this case, the responder node will reply with a Downstream Mapping TLV for each outgoing interface which is part of the ECMP next-hop set for the FEC. Note however the sender node will select the first Downstream Mapping TLV only for the subsequent echo request message with incrementing TTL.
2. If the user initiates an `lsp-ping` of the FEC with the **path-destination** option specified, then the sender node will not include the Downstream Mapping TLV. However, the user can use the **interface** option, part of the same **path-destination** option, to direct the echo request message at the sender node to be sent out a specific outgoing interface which is part of an ECMP path set for the FEC.
3. If the user initiates an `lsp-trace` of the FEC with the **path-destination** option specified but configured not to include a downstream mapping TLV in the MPLS echo request message using the CLI command **downstream-map-tlv {none}**, then the sender node will not include the Downstream Mapping TLV. However, the user can use the **interface** option, part of the same **path-destination** option, to direct the echo request message at the sender node to be sent out a specific outgoing interface which is part of an ECMP path set for the FEC.
4. If the user initiates an `lsp-trace` of the FEC with the **path-destination** option specified, then the sender node will include the multipath information in the Downstream Mapping TLV in the echo request message (multipath type=8). The **path-destination** option allows the user to exercise a specific path of a FEC in the presence of ECMP. This is performed by having the user enter a specific address from the 127/8 range which is then inserted in the multipath type 8 information field of the Downstream Mapping TLV. The CPM code at each LSR in the path of the target FEC runs the same hash routine as the data path and replies in the Downstream Mapping TLV with the specific outgoing interface the packet would have been forwarded to if it did not expire at this node and if DEST IP field in the packet's header was set to the 127/8 address value inserted in the multipath type 8 information.
5. The **ldp-treetrace** tool always uses the multipath type=8 and inserts a range of 127/8 addresses instead of a single address in order to exercise multiple ECMP paths of an LDP FEC. As such, it behaves the same way as the **lsp-trace** with the **path-destination** option enabled described above.
6. Note that the **path-destination** option can also be used to exercise a specific ECMP path of an LDP FEC, which is tunneled over a RSVP LSP or of an LDP FEC stitched to a BGP FEC in the presence of BGP ECMP paths. The user must however enable the use of the new DDMAP

TLV either globally (**config>test-oam>mpls-echo-request-downstream-map ddmmap**) or within the specific **ldp-treetrace** or **lsp-trace** test (**downstream-map-tlv ddmmap** option).

---

## Lsp-ping and Lsp-trace over Unnumbered IP Interface

Lsp-ping and p2mp-lsp-ping operate over a network using unnumbered links without any changes. Lsp-trace, p2mp-lsp-trace and ldp-treetrace are modified such that the unnumbered interface is properly encoded in the downstream mapping (DSMAP/DDMAP) TLV.

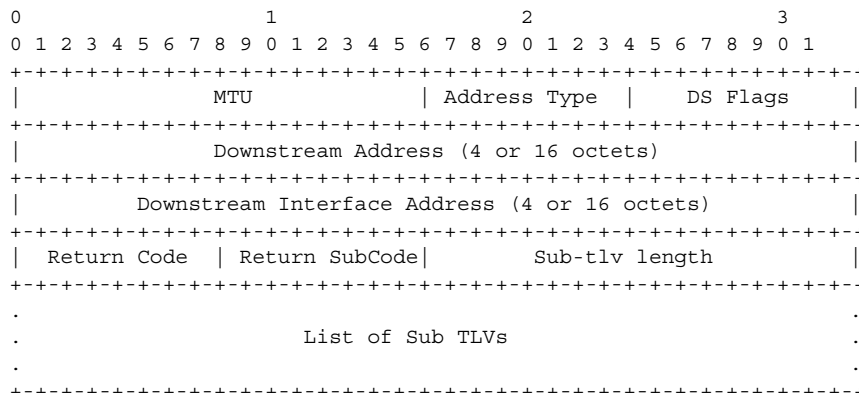
In a RSVP P2P or P2MP LSP, the upstream LSR encodes the downstream router-id in the “Downstream IP Address” field and the local unnumbered interface index value in the “Downstream Interface Address” field of the DSMAP/DDMAP TLV as per RFC 4379. Both values are taken from the TE database.

In a LDP unicast FEC or mLDP P2MP FEC, the interface index assigned by the peer LSR is not readily available to the LDP control plane. In this case, the alternative method described in RFC 4379 is used. The upstream LSR sets the Address Type to IPv4 Unnumbered, the Downstream IP Address to a value of 127.0.0.1, and the interface index is set to 0. If an LSR receives an echo-request packet with this encoding in the DSMAP/DDMAP TLV, it will bypass interface verification but continue with label validation.

---

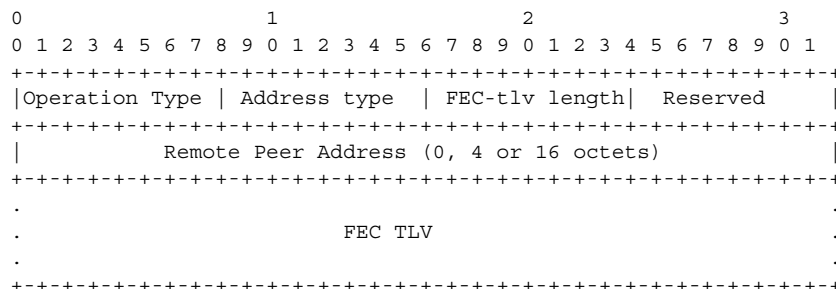
## Downstream Detailed Mapping (DDMAP) TLV

The DDMAP TLV provides with exactly the same features as the existing DSMAP TLV, plus the enhancements to trace the details of LSP stitching and LSP hierarchy. The latter is achieved using a new sub-TLV of the DDMAP TLV called the FEC stack change sub-TLV. The following are the structures of these two objects as defined in RFC 6424.



**Figure 8: DDMAP TLV**

The DDMAP TLV format is derived from the DSMAP TLV format. The key change is that variable length and optional fields have been converted into sub-TLVs. The fields have the same use and meaning as in RFC 4379.



**Figure 9: FEC Stack Change Sub-TLV**

The operation type specifies the action associated with the FEC stack change. The following operation types are defined.

| Type # | Operation |
|--------|-----------|
| 1      | Push      |
| 2      | Pop       |

More details on the processing of the fields of the FEC stack change sub-TLV are provided later in this section.

The user can configure which downstream mapping TLV to use globally on a system by using the following command:

```
configure test-oam mpls-echo-request-downstream-map {dsmap | ddmmap}
```

This command specifies which format of the downstream mapping TLV to use in all LSP trace packets and LDP tree trace packets originated on this node. The Downstream Mapping (DSMAP) TLV is the original format in RFC 4379 and is the default value. The Downstream Detailed Mapping (DDMAP) TLV is the new enhanced format specified in RFC 6424.

This command applies to LSP trace of an RSVP P2P LSP, a MPLS-TP LSP, a BGP IPv4 Label Route, or LDP unicast FEC, and to LDP tree trace of a unicast LDP FEC. It does not apply to LSP trace of an RSVP P2MP LSP which always uses the DDMAP TLV.

The global DSMAP/DDMAP setting impacts the behavior of both OAM LSP trace packets and SAA test packets of type **lsp-trace** and is used by the sender node when one of the following events occurs:

1. An SAA test of type **lsp-trace** is created (not modified) and no value is specified for the per-test **downstream-map-tlv {dsmap|ddmap|none}** option. In this case the SAA test **downstream-map-tlv** value defaults to the global **mpls-echo-request-downstream-map** value.
2. An OAM test of type **lsp-trace** test is executed and no value is specified for the per-test **downstream-map-tlv {dsmap|ddmap|none}** option. In this case, the OAM test **downstream-map-tlv** value defaults to the global **mpls-echo-request-downstream-map** value.

A consequence of the rules above is that a change to the value of **mpls-echo-request-downstream-map** option does not affect the value inserted in the downstream mapping TLV of existing tests.

The following are the details of the processing of the new DDMAP TLV:

1. When either the DSMAP TLV or the DDMAP TLV is received in an echo request message, the responder node will include the same type of TLV in the echo reply message with the proper downstream interface information and label stack information.
2. If an echo request message without a Downstream Mapping TLV (DSMAP or DDMAP) expires at a node which is not the egress for the target FEC stack, the responder node always includes the DSMAP TLV in the echo reply message. This can occur in the following cases:
  - a. The user issues a LSP trace from a sender node with a **min-ttl** value higher than 1 and a **max-ttl** value lower than the number of hops to reach the egress of the target FEC



stack. This is the sender node behavior when the global configuration or the per-test setting of the DSMAP/DDMAP is set to DSMAP.

- b. The user issues a LSP ping from a sender node with a **ttl** value lower than the number of hops to reach the egress of the target FEC stack. This is the sender node behavior when the global configuration of the DSMAP/DDMAP is set to DSMAP.
  - c. The behavior in (a) is changed when the global configuration or the per-test setting of the Downstream Mapping TLV is set to DDMAP. The sender node will include in this case the DDMAP TLV with the Downstream IP address field set to the all-routers multicast address as per Section 3.3 of RFC 4379. The responder node then bypasses the interface and label stack validation and replies with a DDMAP TLV with the correct downstream information for the target FEC stack.
3. A sender node never includes the DSMAP or DDMAP TLV in an lsp-ping message.

## Using DDMAP TLV in LSP Stitching and LSP Hierarchy

In addition to performing the same features as the DSMAP TLV, the new DDMAP TLV addresses the following scenarios:

1. Full validation of an LDP FEC stitched to a BGP IPv4 label route. In this case, the LSP trace message is inserted from the LDP LSP segment or from the stitching point.
2. Full validation of a BGP IPv4 label route stitched to an LDP FEC. The LSP trace message is inserted from the BGP LSP segment or from the stitching point.
3. Full validation of an LDP FEC which is stitched to a BGP LSP and stitched back into an LDP FEC. In this case, the LSP trace message is inserted from the LDP segments or from the stitching points.
4. Full validation of an LDP FEC tunneled over an RSVP LSP using LSP trace.
5. Full validation of a BGP IPv4 label route tunneled over an RSVP LSP or an LDP FEC.

In order to properly check a target FEC which is stitched to another FEC (stitching FEC) of the same or a different type, or which is tunneled over another FEC (tunneling FEC), it is necessary for the responding nodes to provide details about the FEC manipulation back to the sender node. This is achieved via the use of the new FEC stack change sub-TLV in the Downstream Detailed Mapping TLV (DDMAP) defined in RFC 6424.

When the user configures the use of the DDMAP TLV on a trace for an LSP that does not undergo stitching or tunneling operation in the network, the procedures at the sender and responder nodes are the same as in the case of the existing DSMAP TLV.

This feature however introduces changes to the target FEC stack validation procedures at the sender and responder nodes in the case of LSP stitching and LSP hierarchy. These changes pertain to the processing of the new FEC stack change sub-TLV in the new DDMAP TLV and the new return code 15 Label switched with FEC change. The following is a description of the main changes which are a superset of the rules described in Section 4 of RFC 6424 to allow greater scope of interoperability with other vendor implementations.

---

### Responder Node Procedures

1. As a responder node, the node will always insert a global return code of either 3 Replying router is an egress for the FEC at stack-depth <RSC> or 14 See DDMAP TLV for Return Code and Return Subcode.
2. When the responder node inserts a global return code of 3, it will not include a DDMAP TLV.
3. When the responder node includes the DDMAP TLV, it inserts a global return code 14 See DDMAP TLV for Return Code and Return Subcode and:
  - a. On a success response, include a return code of 15 in the DDMAP TLV for each downstream which has a FEC stack change TLV.
  - b. On a success response, include a return code 8 Label switched at stack-depth <RSC> in the DDMAP TLV for each downstream if no FEC stack change sub-TLV is present.
  - c. On a failure response, include an appropriate error return code in the DDMAP TLV for each downstream.
4. A tunneling node indicates that it is pushing a FEC (the tunneling FEC) on top of the target FEC stack TLV by including a FEC stack change sub-TLV in the DDMAP TLV with a FEC operation type value of PUSH. It also includes a return code 15 Label switched with FEC change. The downstream interface address and downstream IP address fields of the DDMAP TLV are populated for the pushed FEC. The remote peer address field in the FEC stack change sub-TLV is populated with the address of the control plane peer for the pushed FEC. The Label stack sub-TLV provides the full label stack over the downstream interface.
5. A node that is stitching a FEC indicates that it is performing a POP operation for the stitched FEC followed by a PUSH operation for the stitching FEC and potentially one PUSH operation for the transport tunnel FEC. It will thus include two or more FEC stack change

sub-TLVs in the DDMAP TLV in the echo reply message. It also includes and a return code 15 Label switched with FEC change. The downstream interface address and downstream address fields of the DDMAP TLV are populated for the stitching FEC. The remote peer address field in the FEC stack change sub-TLV of type POP is populated with a null value (0.0.0.0). The remote peer address field in the FEC stack change sub-TLV of type PUSH is populated with the address of the control plane peer for the tunneling FEC. The Label stack sub-TLV provides the full label stack over the downstream interface.

6. If the responder node is the egress for one or more FECs in the target FEC Stack, then it must reply with no DDMAP TLV and with a return code 3 Replying router is an egress for the FEC at stack-depth <RSC>. RSC must be set to the depth of the topmost FEC. This operation is iterative in a sense that at the receipt of the echo reply message the sender node will pop the topmost FEC from the target stack FEC TLV and resend the echo request message with the same TTL value as explained in (5) below. The responder node will thus perform exactly the same operation as described in this step until all FECs are popped or until the topmost FEC in the target FEC stack TLV matches the tunneled or stitched FEC. In the latter case, processing of the target FEC stack TLV follows again steps (1) or (2).

## Sender Node Procedures

1. If the echo reply message contains the return code 14 See DDMAP TLV for Return Code and Return Subcode and the DDMAP TLV has a return code 15 Label switched with FEC change, the sender node adjusts the target FEC Stack TLV in the echo request message for the next value of the TTL to reflect the operation on the current target FEC stack as indicated in the FEC stack change sub-TLV received in the DDMAP TLV of the last echo reply message. In other words, one FEC is popped at most and one or more FECs are pushed as indicated.
2. If the echo reply message contains the return code 3 Replying router is an egress for the FEC at stack-depth <RSC>, then:
  - a. If the value for the label stack depth specified in the Return Sub-Code (RSC) field is the same as the depth of current target FEC Stack TLV, then the sender node considers the trace operation complete and terminates it. A responder node will cause this case to occur as per step (6) of the responder node procedures.
  - b. If the value for the label stack depth specified in the Return Sub-Code (RSC) field is different from the depth of the current target FEC Stack TLV, the sender node must continue the LSP trace with the same TTL value after adjusting the target FEC stack TLV by removing the top FEC. Note this step will continue iteratively until the value for the label stack depth specified in the Return Sub-Code (RSC) field is the same as the depth of current target FEC Stack TLV and in which case step (a) is performed. A

responder node will cause this case to occur as per step (6) of the responder node procedures.

- c. If a DDMAP TLV with or without a FEC stack change sub-TLV is included, then the sender node must ignore it and processing is performed as per steps (a) or (b) above. A responder node will not cause this case to occur but a third party implementation may do.
3. As a sender node, the can accept an echo-reply message with the global return code of either 14 (with DDMAP TLV return code of 15 or 8), or 15 and process properly the FEC stack change TLV as per step (1) of the sender node procedures.
4. If an LSP ping is performed directly to the egress LER of the stitched FEC, there is no DDMAP TLV included in the echo request message and thus the responder node, which is the egress node, will still reply with return code 4 `Replying router has no mapping for the FEC at stack- depth <RSC>`. This case cannot be resolved with this feature.

Note the following limitation when a BGP IPv4 label route is resolved to an LDP FEC which itself is resolved to an RSVP LSP all on the same node. This 2-level LSP hierarchy is not supported as a feature on the SROS but user is not prevented from configuring it. In that case, user and OAM packets are forwarded by the sender node using two labels (T-LDP and BGP). The LSP trace will fail on the downstream node with return code 1 `Malformed echo request received` since there is no label entry for the RSVP label.

---

## LSP Ping for RSVP P2MP LSP

**NOTE:** RSVP-based signalling for P2MP LSPs is supported only on 7210 SAS-T, SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-R6 and 7210 SAS-R12 devices.

An LSP ping can be generated by entering the following OAM command:

```
oam p2mp-lsp-ping lsp-name [p2mp-instance instance-name [s2l-dest-addr ip-address [...up to 5 max]]] [fc fc-name] [size octets] [ttl label-ttl] [timeout timeout] [detail]
```

The echo request message is sent on the active P2MP instance and is replicated in the data path over all branches of the P2MP LSP instance. By default, all egress LER nodes which are leaves of the P2MP LSP instance will reply to the echo request message.

The user can reduce the scope of the echo reply messages by explicitly entering a list of addresses for the egress LER nodes that are required to reply. A maximum of five addresses can be specified in a single execution of the **p2mp-lsp-ping** command. If all five egress LER nodes are router nodes, they will be able to parse the list of egress LER addresses and will reply. Note that RFC 6425 specifies that only the top address in the P2MP egress identifier TLV must be inspected by an

egress LER. When interoperating with other implementations, the router egress LER will respond if its address is anywhere in the list. Furthermore, if another vendor implementation is the egress LER, only the egress LER matching the top address in the TLV may respond.

If the user enters the same egress LER address more than once in a single **p2mp-lsp-ping** command, the head-end node displays a response to a single one and displays a single error warning message for the duplicate ones. When queried over SNMP, the head-end node issues a single response trap and issues no trap for the duplicates.

The **timeout** parameter should be set to the time it would take to get a response from all probed leaves under no failure conditions. For that purpose, its range extends to 120 seconds for a **p2mp-lsp-ping** from a 10 second **lsp-ping** for P2P LSP. The default value is 10 seconds.

The router head-end node displays a “Send\_Fail” error when a specific S2L path is down only if the user explicitly listed the address of the egress LER for this S2L in the ping command.

Similarly, the router head-end node displays the timeout error when no response is received for an S2L after the expiry of the timeout timer only if the user explicitly listed the address of the egress LER for this S2L in the **ping** command.

The user can configure a specific value of the **tll** parameter to force the echo request message to expire on a router branch node or a bud LSR node. The latter replies with a downstream mapping TLV for each branch of the P2MP LSP in the echo reply message. Note that a maximum of 16 downstream mapping TLVs can be included in a single echo reply message. It also sets the multipath type to zero in each downstream mapping TLV and will thus not include any egress address information for the reachable egress LER nodes for this P2MP LSP.

If the router ingress LER node receives the new multipath type field with the list of egress LER addresses in an echo reply message from another vendor implementation, it will ignore but will not cause an error in processing the downstream mapping TLV.

If the ping expires at an LSR node which is performing a re-merge or cross-over operation in the data path between two or more ILMs of the same P2MP LSP, there will be an echo reply message for each copy of the echo request message received by this node.

The output of the command without the **detail** parameter specified provides a high-level summary of error codes and/or success codes received.

The output of the command with the **detail** parameter specified shows a line for each replying node as in the output of the LSP ping for a P2P LSP.

The display is delayed until all responses are received or the timer configured in the timeout parameter expired. No other CLI commands can be entered while waiting for the display. A **control-C** (^C) command will abort the ping operation.

## LSP Trace for RSVP P2MP LSP

**NOTE:** RSVP based signalling for P2MP LSPs is supported only on 7210 SAS-T, SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-R6 and 7210 SAS-R12 devices.

An LSP trace can be generated by entering the following OAM command:

```
oam p2mp-lsp-trace lsp-name p2mp-instance instance-name s2l-dest-addr ip-address
[fc fc-name] [size octets] [max-fail no-response-count] [probe-count probes-per-hop]
[min-ttl min-label-ttl] [max-ttl max-label-ttl] [timeout timeout] [interval interval] [detail]
```


The LSP trace capability allows the user to trace the path of a single S2L path of a P2MP LSP. Its operation is similar to that of the **p2mp-lsp-ping** command but the sender of the echo reply request message includes the downstream mapping TLV to request the downstream branch information from a branch LSR or bud LSR. The branch LSR or bud LSR will then also include the downstream mapping TLV to report the information about the downstream branches of the P2MP LSP. An egress LER does not include this TLV in the echo response message.

The **probe-count** parameter operates in the same way as in LSP trace on a P2P LSP. It represents the maximum number of probes sent per TTL value before giving up on receiving the echo reply message. If a response is received from the traced node before reaching maximum number of probes, then no more probes are sent for the same TTL. The sender of the echo request then increments the TTL and uses the information it received in the downstream mapping TLV to start sending probes to the node downstream of the last node which replied. This continues until the egress LER for the traced S2L path replied.

Since the command traces a single S2L path, the **timeout** and **interval** parameters keep the same value range as in LSP trace for a P2P LSP.

The following supported options in **lsp-trace** for P2P LSP are not applicable: **path**, **prefix**, **path-destination**, and [**interface** | **next-hop**].


The P2MP LSP Trace makes use of the Downstream Detailed Mapping (DDMAP) TLV. The following excerpt from RFC 6424 details the format of the new DDMAP TLV entered in the **path-destination** belongs to one of the possible outgoing interface of the FEC.





## LSP Trace for RSVP P2MP LSP

S2L2 = ABDE (to leaf E)



1. Tracing S2L1 when ILM on interface C-B blocked at node B:
  - For TTL=1, A gets a response from C only as B does not have S2L1 on the ILM on interface A-B.
  - For TTL=2, assume A gets first the response from B which indicates a success. It then builds the next probe with TTL=3. B will only pass the copy of the message arriving on interface A-B and will drop the one arriving on interface C-B (treats it like a data packet since it does not expire at node B). This copy will expire at F. However F will return a "DSMappingMismatched" error because the DDMAP TLV was the one provided by node B in TTL=2 step. The trace will abort at this point in time. However, A knows it got a second response from Node D for TTL=2 with a "DSMappingMismatched" error.
  - If A gets the response from D first with the error code, it waits to see if it gets a response from B or it times out. In either case, it will log this status as **multiple replies received per probe** in the last probe history and aborts the trace.
2. Tracing S2L2 when ILM on interface A-B blocked at node B:
  - For TTL=1, B responds with a success. C does not respond as it does not have an ILM for S2L2.
  - For TTL=2, B drops the copy coming on interface A-B. It receives a copy coming on interface B-C but will drop it as the ILM does not contain S2L2. Node A times out. Next, node A generates a probe with TTL=3 without a DDMAP TLV. This time node D will respond with a success and will include its downstream DDMAP TLV to node E. The rest of the path will be discovered correctly. The traced path for S2L2 will look something like: A-B-(\*)-D-E.

The router ingress LER detects a re-merge condition when it receives two or more replies to the same probe, such as the same TTL value. It displays the following message to the user regardless if the trace operation successfully reached the egress LER or was aborted earlier:

```
Probe returned multiple responses. Result may be inconsistent.
```

This warning message indicates to the user the potential of a re-merge scenario and that a **p2mp-lsp-ping** command for this S2L should be used to verify that the S2L path is not defective.



The router ingress LER behavior is to always proceed to the next ttl probe when it receives an OK response to a probe or when it times out on a probe. If however it receives replies with an error return code, it must wait until it receives an OK response or it times out. If it times out without receiving an OK reply, the LSP trace must be aborted.

The following are possible echo reply messages received and corresponding ingress LER behavior:

- One or more error return codes + OK: display OK return code. Proceed to next **ttl** probe. Display warning message at end of trace.
  - OK + One or more error return codes: display OK return code. Proceed to next **ttl** probe right after receiving the OK reply but keep state that more replies received. Display warning message at end of trace.
  - OK + OK: should not happen for re-merge but would continue trace on 1st OK reply. This is the case when one of the branches of the P2MP LSP is activating the P2P bypass LSP. In this case, the head-end node will get a reply from both a regular P2MP LSR which has the ILM for the traced S2L and from an LSR switching the P2P bypass for other S2Ls. The latter does not have context for the P2MP LSP being tunneled but will respond after doing a label stack validation.
  - One error return code + timeout: abort LSP trace and display error code. Ingress LER cannot tell the error is due to a re-merge condition.
  - More than one error return code + timeout: abort LSP trace and display first error code. Display warning message at end of trace.
  - Timeout on probe without any reply: display "\*" and proceed to next **ttl** probe.
-

## SDP Diagnostics

The 7210 SAS SDP diagnostics are SDP ping and SDP MTU path discovery.

---

### SDP Ping

SDP ping performs in-band uni-directional or round-trip connectivity tests on SDPs. The SDP ping OAM packets are sent in-band, in the tunnel encapsulation, so it will follow the same path as traffic within the service. The SDP ping response can be received out-of-band in the control plane, or in-band using the data plane for a round-trip test.

For a uni-directional test, SDP ping tests:

- Egress SDP ID encapsulation
- Ability to reach the far-end IP address of the SDP ID within the SDP encapsulation
- Path MTU to the far-end IP address over the SDP ID
- Forwarding class mapping between the near-end SDP ID encapsulation and the far-end tunnel termination

For a round-trip test, SDP ping uses a local egress SDP ID and an expected remote SDP ID. Since SDPs are uni-directional tunnels, the remote SDP ID must be specified and must exist as a configured SDP ID on the far-end 7210 SAS. SDP round trip testing is an extension of SDP connectivity testing with the additional ability to test:

- Remote SDP ID encapsulation
  - Potential service round trip time
  - Round trip path MTU
  - Round trip forwarding class mapping
- 

### SDP MTU Path Discovery

In a large network, network devices can support a variety of packet sizes that are transmitted across its interfaces. This capability is referred to as the Maximum Transmission Unit (MTU) of network interfaces. It is important to understand the MTU of the entire path end-to-end when provisioning services, especially for virtual leased line (VLL) services where the service must support the ability to transmit the largest customer packet.

The Path MTU discovery tool provides a powerful tool that enables service provider to get the exact MTU supported by the network's physical links between the service ingress and service termination points (accurate to one byte).

## Service Diagnostics

Alcatel-Lucent's Service ping feature provides end-to-end connectivity testing for an individual service. Service ping operates at a higher level than the SDP diagnostics in that it verifies an individual service and not the collection of services carried within an SDP.

Service ping is initiated from a 7210 SAS router to verify round-trip connectivity and delay to the far-end of the service. Alcatel-Lucent's implementation functions for MPLS tunnels and tests the following from edge-to-edge:

- Tunnel connectivity
- VC label mapping verification
- Service existence
- Service provisioned parameter verification
- Round trip path verification
- Service dynamic configuration verification

## VPLS MAC Diagnostics

While the LSP ping, SDP ping and service ping tools enable transport tunnel testing and verify whether the correct transport tunnel is used, they do not provide the means to test the learning and forwarding functions on a per-VPLS-service basis.

It is conceivable, that while tunnels are operational and correctly bound to a service, an incorrect Forwarding Information Base (FIB) table for a service could cause connectivity issues in the service and not be detected by the ping tools. Alcatel-Lucent has developed VPLS OAM functionality to specifically test all the critical functions on a per-service basis. These tools are based primarily on the IETF document draft-stokes-vkompella-ppvpn-hvpls-oam-xx.txt, *Testing Hierarchical Virtual Private LAN Services*.

The VPLS OAM tools are:

- **MAC Ping** — Provides an end-to-end test to identify the egress customer-facing port where a customer MAC was learned. MAC ping can also be used with a broadcast MAC address to identify all egress points of a service for the specified broadcast MAC.
- **MAC Trace** — Provides the ability to trace a specified MAC address hop-by-hop until the last node in the service domain. An SAA test with MAC trace is considered successful when there is a reply from a far-end node indicating that they have the destination MAC address on an egress SAP or the CPM.
- **CPE Ping** — Provides the ability to check network connectivity to the specified client device within the VPLS. CPE ping will return the MAC address of the client, as well as the SAP and PE at which it was learned.
- **MAC Populate** — Allows specified MAC addresses to be injected in the VPLS service domain. This triggers learning of the injected MAC address by all participating nodes in the service. This tool is generally followed by MAC ping or MAC trace to verify if correct learning occurred.
- **MAC Purge** — Allows MAC addresses to be flushed from all nodes in a service domain.

---

## MAC Ping

For a MAC ping test, the destination MAC address (unicast or multicast) to be tested must be specified. A MAC ping packet can be sent through the control plane or the data plane. When sent by the control plane, the ping packet goes directly to the destination IP in a UDP/IP OAM packet. If it is sent by the data plane, the ping packet goes out with the data plane format.

In the control plane, a MAC ping is forwarded along the flooding domain if no MAC address bindings exist. If MAC address bindings exist, then the packet is forwarded along those paths (if

they are active). Finally, a response is generated only when there is an egress SAP binding to that MAC address. A control plane request is responded to via a control reply only.

In the data plane, a MAC ping is sent with a VC label TTL of 255. This packet traverses each hop using forwarding plane information for next hop, VC label, etc. The VC label is swapped at each service-aware hop, and the VC TTL is decremented. If the VC TTL is decremented to 0, the packet is passed up to the management plane for processing. If the packet reaches an egress node, and would be forwarded out a customer facing port, it is identified by the OAM label below the VC label and passed to the management plane.

MAC pings are flooded when they are unknown at an intermediate node. They are responded to only by the egress nodes that have mappings for that MAC address.

---

## MAC Trace

A MAC trace functions like an LSP trace with some variations. Operations in a MAC trace are triggered when the VC TTL is decremented to 0.

Like a MAC ping, a MAC trace can be sent either by the control plane or the data plane.

For MAC trace requests sent by the control plane, the destination IP address is determined from the control plane mapping for the destination MAC. If the destination MAC is known to be at a specific remote site, then the far-end IP address of that SDP is used. If the destination MAC is not known, then the packet is sent unicast, to all SDPs in the service with the appropriate squelching.

A control plane MAC traceroute request is sent via UDP/IP. The destination UDP port is the LSP ping port. The source UDP port is whatever the system gives (note that this source UDP port is really the demultiplexor that identifies the particular instance that sent the request, when correlating the reply). The source IP address is the system IP of the sender.

When a traceroute request is sent via the data plane, the data plane format is used. The reply can be via the data plane or the control plane.

A data plane MAC traceroute request includes the tunnel encapsulation, the VC label, and the OAM, followed by an Ethernet DLC, a UDP and IP header. If the mapping for the MAC address is known at the sender, then the data plane request is sent down the known SDP with the appropriate tunnel encapsulation and VC label. If it is not known, then it is sent down every SDP (with the appropriate tunnel encapsulation per SDP and appropriate egress VC label per SDP binding).

The tunnel encapsulation TTL is set to 255. The VC label TTL is initially set to the min-ttl (default is 1). The OAM label TTL is set to 2. The destination IP address is the all-routers multicast address. The source IP address is the system IP of the sender.

The destination UDP port is the LSP ping port. The source UDP port is whatever the system gives (note that this source UDP port is really the demultiplexor that identifies the particular instance that sent the request, when correlating the reply).

The Reply Mode is either 3 (i.e., reply via the control plane) or 4 (i.e., reply through the data plane), depending on the reply-control option. By default, the data plane request is sent with Reply Mode 3 (control plane reply) Reply Mode 4 (data plane reply).

The Ethernet DLC header source MAC address is set to either the system MAC address (if no source MAC is specified) or to the specified source MAC. The destination MAC address is set to the specified destination MAC. The EtherType is set to IP.

---

## CPE Ping

The MAC ping OAM tool makes it possible to detect whether a particular MAC address has been learned in a VPLS.

The **cpe-ping** command extends this capability to detecting end-station IP addresses inside a VPLS. A CPE ping for a specific destination IP address within a VPLS will be translated to a MAC-ping towards a broadcast MAC address. Upon receiving such a MAC ping, each peer PE within the VPLS context will trigger an ARP request for the specific IP address. The PE receiving a response to this ARP request will report back to the requesting 7210 SAS. It is encouraged to use the source IP address of 0.0.0.0 to prevent the provider's IP address of being learned by the CE.

## MAC Populate

MAC populate is used to send a message through the flooding domain to learn a MAC address as if a customer packet with that source MAC address had flooded the domain from that ingress point in the service. This allows the provider to craft a learning history and engineer packets in a particular way to test forwarding plane correctness.

The MAC populate request is sent with a VC TTL of 1, which means that it is received at the forwarding plane at the first hop and passed directly up to the management plane. The packet is then responded to by populating the MAC address in the forwarding plane, like a conventional learn although the MAC will be an OAM-type MAC in the FIB to distinguish it from customer MAC addresses.

This packet is then taken by the control plane and flooded out the flooding domain (squelching appropriately, the sender and other paths that would be squelched in a typical flood).

This controlled population of the FIB is very important to manage the expected results of an OAM test. The same functions are available by sending the OAM packet as a UDP/IP OAM packet. It is then forwarded to each hop and the management plane has to do the flooding.

Options for MAC populate are to force the MAC in the table to type OAM (in case it already existed as dynamic or static or an OAM induced learning with some other binding), to prevent new dynamic learning to over-write the existing OAM MAC entry, to allow customer packets with this MAC to either ingress or egress the network, while still using the OAM MAC entry.

Finally, an option to flood the MAC populate request causes each upstream node to learn the MAC, for example, populate the local FIB with an OAM MAC entry, and to flood the request along the data plane using the flooding domain.

An age can be provided to age a particular OAM MAC after a different interval than other MACs in a FIB.

---

## MAC Purge

MAC purge is used to clear the FIBs of any learned information for a particular MAC address. This allows one to do a controlled OAM test without learning induced by customer packets. In addition to clearing the FIB of a particular MAC address, the purge can also indicate to the control plane not to allow further learning from customer packets. This allows the FIB to be clean, and be populated only via a MAC Populate.

MAC purge follows the same flooding mechanism as the MAC populate.



A UDP/IP version of this command is also available that does not follow the forwarding notion of the flooding domain, but the control plane notion of it.

## VLL Diagnostics

---

### VCCV Ping


VCCV ping is used to check connectivity of a VLL in-band. It checks that the destination (target) PE is the egress for the Layer 2 FEC. It provides a cross-check between the data plane and the control plane. It is in-band, meaning that the VCCV ping message is sent using the same encapsulation and along the same path as user packets in that VLL. This is equivalent to the LSP ping for a VLL service. VCCV ping reuses an LSP ping message format and can be used to test a VLL configured over an MPLS SDP.

---

### VCCV-Ping Application

VCCV effectively creates an IP control channel within the pseudowire between PE1 and PE2. PE2 should be able to distinguish on the receive side VCCV control messages from user packets on that VLL. There are three possible methods of encapsulating a VCCV message in a VLL which translates into three types of control channels:

1. Use of a Router Alert Label immediately above the VC label. This method has the drawback that if ECMP is applied to the outer LSP label (for example, transport label), the VCCV message will not follow the same path as the user packets. This effectively means it will not troubleshoot the appropriate path. This method is supported by the 7210 SAS.
2. Use of the OAM control word as illustrated in [Figure 10](#).




**Figure 10: OAM Control Word Format**

The first nibble is set to 0x1. The Format ID and the reserved fields are set to 0 and the channel type is the code point associated with the VCCV IP control channel as specified in the PWE3 IANA registry (RFC 4446). The channel type value of 0x21 indicates that the Associated Channel carries an IPv4 packet.

The use of the OAM control word assumes that the draft-martini control word is also used on the user packets. This means that if the control word is optional for a VLL and is not configured, the 7210 SAS M, X PE node will only advertise the router alert label as the CC capability in the Label Mapping message. This method is supported by the 7210 SAS M.

3. Set the TTL in the VC label to 1 to force PE2 control plane to process the VCCV message. This method is not guaranteed to work under all circumstances. For instance, the draft mentions some implementations of penultimate hop popping overwrite the TTL field. This method is not supported by the 7210 SAS M.

When sending the label mapping message for the VLL, PE1 and PE2 must indicate which of the above OAM packet encapsulation methods (for example, which control channel type) they support. This is accomplished by including an optional VCCV TLV in the pseudowire FEC Interface Parameter field. The format of the VCCV TLV is shown in [Figure 11](#).



**Figure 11: VCCV TLV**

Note that the absence of the optional VCCV TLV in the Interface parameters field of the pseudowire FEC indicates the PE has no VCCV capability.

The Control Channel (CC) Type field is a bitmask used to indicate if the PE supports none, one, or many control channel types.

- 0x00 None of the following VCCV control channel types are supported
- 0x01 PWE3 OAM control word (see [Figure 10](#))
- 0x02 MPLS Router Alert Label
- 0x04 MPLS inner label TTL = 1

If both PE nodes support more than one of the CC types, then a 7210 SAS PE will make use of the one with the lowest type value. For instance, OAM control word will be used in preference to the MPLS router alert label.

The Connectivity Verification (CV) bitmask field is used to indicate the specific type of VCCV packets to be sent over the VCCV control channel. The valid values are:

0x00 None of the below VCCV packet type are supported.

0x01 ICMP ping. Not applicable to a VLL over a MPLS SDP and as such is not supported by the 7210 SAS.

0x02 LSP ping. This is used in VCCV-Ping application and applies to a VLL over an MPLS SDP. This is supported by the 7210 SAS.

A VCCV ping is an LSP echo request message as defined in RFC 4379. It contains an L2 FEC stack TLV which must include within the sub-TLV type 10 “FEC 128 Pseudowire”. It also contains a field which indicates to the destination PE which reply mode to use. There are four reply modes defined in RFC 4379:

Reply mode, meaning:

1. Do not reply. This mode is supported by the 7210 SAS.
2. Reply via an IPv4/IPv6 UDP packet. This mode is supported by the 7210 SAS.
3. Reply with an IPv4/IPv6 UDP packet with a router alert. This mode sets the router alert bit in the IP header and is not be confused with the CC type which makes use of the router alert label. This mode is not supported by the 7210 SAS.
4. Reply via application level control channel. This mode sends the reply message inband over the pseudowire from PE2 to PE1. PE2 will encapsulate the Echo Reply message using the CC type negotiated with PE1. This mode is supported by the 7210 SAS.

The reply is an LSP echo reply message as defined in RFC 4379. The message is sent as per the reply mode requested by PE1. The return codes supported are the same as those supported in the 7210 SAS LSP ping capability.

The VCCV ping feature is in addition to the service ping OAM feature which can be used to test a service between 7210 SAS M nodes. The VCCV ping feature can test connectivity of a VLL with any third party node which is compliant to RFC 5085.




Figure 12: VCCV-Ping Application

## VCCV-Ping in a Multi-Segment Pseudowire

Pseudowire switching is a method for scaling a large network of VLL or VPLS services by removing the need for a full mesh of T-LDP sessions between the PE nodes as the number of these nodes grow over time. Pseudowire switching is also used whenever there is a need to deploy a VLL service across two separate routing domains.

In the network, a Termination PE (T-PE) is where the pseudowire originates and terminates.

VCCV ping is extended to be able to perform the following OAM functions:

1. VCCV ping to a destination PE. A VLL FEC Ping is a message sent by T-PE1 to test the FEC at T-PE2. The operation at T-PE1 and T-PE2 is the same as in the case of a single-segment pseudowire. The pseudowire switching node, S-PE1, pops the outer label, swaps the inner (VC) label, decrements the TTL of the VC label, and pushes a new outer label. The 7210 SAS M PE1 node does not process the VCCV OAM Control Word unless the VC label TTL expires. In that case, the message is sent to the CPM for further validation and processing. This is the method described in draft-hart-pwe3-segmented-pw-vcv.

## Automated VCCV-Trace Capability for MS-Pseudowire

Although tracing of the MS-pseudowire path is possible using the methods explained in previous sections, these require multiple manual iterations and that the FEC of the last pseudowire segment to the target T-PE/S-PE be known a priori at the node originating the echo request message for each iteration. This mode of operation is referred to as a “ping” mode.

The automated VCCV-trace can trace the entire path of a pseudowire with a single command issued at the T-PE or at an S-PE. This is equivalent to LSP-trace and is an iterative process by which the ingress T-PE or T-PE sends successive VCCV-ping messages with incrementing the TTL value, starting from TTL=1.

The method is described in draft-hart-pwe3-segmented-pw-vcv, *VCCV Extensions for Segmented Pseudo-Wire*, and is pending acceptance by the PWE3 working group. In each iteration, the source T-PE or S-PE builds the MPLS echo request message in a way similar to [VCCV Ping on page 106](#). The first message with TTL=1 will have the next-hop S-PE T-LDP session source address in the Remote PE Address field in the pseudowire FEC TLV. Each S-PE which terminates and processes the message will include in the MPLS echo reply message the FEC 128 TLV corresponding the pseudowire segment to its downstream node. The inclusion of the FEC TLV in the echo reply message is allowed in RFC 4379, *Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures*. The source T-PE or S-PE can then build the next echo reply message with TTL=2 to test the next-next hop for the MS-pseudowire. It will copy the FEC TLV it received in the echo reply message into the new echo request message. The process is terminated when the reply is from the

egress T-PE or when a timeout occurs. If specified, the max-ttl parameter in the vccv-trace command will stop on SPE before reaching T-PE.

The results VCCV-trace can be displayed for a fewer number of pseudowire segments of the end-to-end MS-pseudowire path. In this case, the min-ttl and max-ttl parameters are configured accordingly. However, the T-PE/S-PE node will still probe all hops upto min-ttl in order to correctly build the FEC of the desired subset of segments.

Note that this method does not require the use of the downstream mapping TLV in the echo request and echo reply messages.

---

### **VCCV for Static Pseudowire Segments**

MS pseudowire is supported with a mix of static and signaled pseudowire segments. However, VCCV ping and VCCV-trace is allowed until at least one segment of the MS pseudowire is static. Users cannot test a static segment but also, cannot test contiguous signaled segments of the MS-pseudowire. VCCV ping and VCCV trace is not supported in static-to-dynamic configurations.

## Detailed VCCV-Trace Operation

In [Figure on page 69](#) a trace can be performed on the MS-pseudowire originating from T-PE1 by a single operational command. The following process occurs:

1. T-PE1 sends a VCCV echo request with TTL set to 1 and a FEC 128 containing the pseudowire information of the first segment (pseudowire1 between T-PE1 and S-PE) to S-PE for validation.
2. S-PE validates the echo request with the FEC 128. Since it is a switching point between the first and second segment it builds an echo reply with a return code of 8 and includes the FEC 128 of the second segment (pseudowire2 between S-PE and T-PE2) and sends the echo reply back to T-PE1.
3. T-PE1 builds a second VCCV echo request based on the FEC128 in the echo reply from the S-PE. It increments the TTL and sends the next echo request out to T-PE2. Note that the VCCV echo request packet is switched at the S-PE datapath and forwarded to the next downstream segment without any involvement from the control plane.
4. T-PE2 receives and validates the echo request with the FEC 128 of the pseudowire2 from T-PE1. Since T-PE2 is the destination node or the egress node of the MS-pseudowire it replies to T-PE1 with an echo reply with a return code of 3 (egress router) and no FEC 128 is included.
5. T-PE1 receives the echo reply from T-PE2. T-PE1 is made aware that T-PE2 is the destination of the MS pseudowire because the echo reply does not contain the FEC 128 and because its return code is 3. The trace process is completed.

## Control Plane Processing of a VCCV Echo Message in a MS-Pseudowire

---

### Sending a VCCV Echo Request

When in the ping mode of operation, the sender of the echo request message requires the FEC of the last segment to the target S-PE/T-PE node. This information can either be configured manually or be obtained by inspecting the corresponding sub-TLV's of the pseudowire switching point TLV. However, the pseudowire switching point TLV is optional and there is no guarantee that all S-PE nodes will populate it with their system address and the pseudowire-id of the last pseudowire segment traversed by the label mapping message. Thus the 7210 SAS M implementation will always make use of the user configuration for these parameters.

---

### Receiving an VCCV Echo Request

Upon receiving a VCCV echo request the control plane on S-PEs (or the target node of each segment of the MS pseudowire) validates the request and responds to the request with an echo reply consisting of the FEC 128 of the next downstream segment and a return code of 8 (label switched at stack-depth) indicating that it is an S-PE and not the egress router for the MS-pseudowire.

If the node is the T-PE or the egress node of the MS-pseudowire, it responds to the echo request with an echo reply with a return code of 3 (egress router) and no FEC 128 is included.

---

### Receiving an VCCV Echo Reply

The operation to be taken by the node that receives the echo reply in response to its echo request depends on its current mode of operation such as ping or trace.

In ping mode, the node may choose to ignore the target FEC 128 in the echo reply and report only the return code to the operator.



## MPLS-TP On-Demand OAM Commands (Supported only on 7210 SAS-R6, 7210 SAS-R12 and 7210 SAS-T)

**NOTE:** This feature is supported on 7210 SAS-T network mode, 7210 SAS-R6, and 7210 SAS-R12 devices. This feature is not supported on 7210 SAS-M, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE and 7210 SAS-Mxp devices.

Ping and Trace tools for PWs and LSPs are supported with both IP encapsulation and the MPLS-TP on demand CV channel for non-IP encapsulation (0x025).

---

### MPLS-TP Pseudowires: VCCV-Ping/VCCV-Trace

For vccv-ping and vccv-trace commands:

- Sub-type **static** must be specified. This indicates to the system that the rest of the command contains parameters that are applied to a static PW with a static PW FEC.
- Add the ability to specify the non-IP ACH channel type (0x0025). This is known as the **non-ip control-channel**. This is the default for type static. GAL is not supported for PWs.
- If the ip-control-channel is specified as the encapsulation, then the IPv4 channel type is used (0x0021). In this case, a destination IP address in the 127/8 range is used, while the source address in the UDP/IP packet is set to the system IP address, or may be explicitly configured by the user with the src-ip-address option. This option is only valid if the IPv4 control-channel is specified.
- The reply mode are always assumed to be the same application level control channel type for type static.
- Allow an MPLS-TP global-id and node-id specified under the spoke-sdps with a given sdp-id:vc-id, used for MPLS-TP PW MEPs, or node-id (prefix) only for MIPs.
- The following CLI command description shows the options that are only allowed if the type static option is configured. All other options are blocked.
- As in the existing implementation, the downstream mapping and detailed downstream mapping TLVs (DSMAP/DDMAP TLVs) is not supported on PWs.

```
vccv-ping static <sdp-id:vc-id> [dest-global-id <global-id> dest-node-id <node-id>] [control-channel ipv4 | non-ip] [fc <fc-name> [profile {in|out}]] [size <octets>] [count <send-count>] [timeout <timeout>] [interval <interval>] [ttl <vc-label-ttl>] [src-ip-address <ip-address>]
vccv-trace static <sdp-id:vc-id> [size <octets>] [min-ttl <min-vc-label-ttl>] [max-ttl <max-vc-label-ttl>] [max-fail <no-response-count>] [probe-count <probe-count>] [control-channel ipv4 | non-ip] [timeout <timeout-value>] [interval <interval-value>] [fc <fc-name> [profile {in|out}]] [src-ip-address <ip-address>] [detail]
```

If the spoke-sdp referred to by sdp-id:vc-id has an MPLS-TP PW-Path-ID defined, then those parameters are used to populate the static PW TLV in the target FEC stack of the vccv-ping or vccv-trace packet. If a Global-ID and Node-ID are specified in the command, then these values are used to populate the destination node TLV in the vccv-ping or vccv-trace packet.

The global-id/node-id are only used as the target node identifiers if the vccv-ping is not end-to-end (for example, a TTL is specified in the vccv-ping/trace command and it is < 255); otherwise, the value in the PW Path ID is used. For vccv-ping, the dest-node-id may be entered as a 4-octet IP address <a.b.c.d> or 32-bit integer <1..4294967295>. For vccv-trace, the destination node-id and global-id are taken from the spoke-sdp context.

The same command syntax is applicable for SAA tests configured under configure saa test a type.

## MPLS-TP LSPs: LSP-Ping/LSP Trace

For lsp-ping and lsp-trace commands:

- The sub-type **static** must be specified. This indicates to the system that the rest of the command contains parameters specific to a LSP identified by a static LSP FEC.
- The 7210 SAS supports the use of the G-ACh with non-IP encapsulation or labeled encapsulation with IP de-multiplexing for both the echo request and echo reply for LSP-Ping and LSP-Trace on LSPs with a static LSP FEC (such as MPLS-TP LSPs).
- It is possible to specify the target MPLS-TP MEP/MIP identifier information for LSP Ping. If the target global-id and node-id are not included in the lsp-ping command, then these parameters for the target MEP ID are taken from the context of the LSP. The **tunnel-number** <tunnel-num> and **lsp-num** <lsp-num> for the far-end MEP are always taken from the context of the path under test.

```
lsp-ping static <lsp-name>
 [force]
 [path-type [active|working|protect]]
 [fc <fc-name> [profile {in|out}]]
 [size <octets>]
 [ttl <label-ttl>]
 [send-count <send-count>]
 [timeout <timeout>]
 [interval <interval>]
 [src-ip-address <ip-address>]
 [dest-global-id <dest-global-id> dest-node-id dest-node-id]
 [control-channel none | non-ip] [detail]
lsp-trace static <lsp-name>
 [force]
 [path-type [active|working|protect]]
 [fc <fc-name> [profile {in|out}]]
 [max-fail <no-response-count>]
 [probe-count <probes-per-hop>]
 [size <octets>]
```

```

[min-ttl <min-label-ttl>]
[max-ttl <max-label-ttl>]
[timeout <timeout>]
[interval <interval>]
[src-ip-address <ip-address>]
 [control-channel none | non-ip]
[downstream-map-tlv <dsmap|ddmap>]
[detail]

```

The following commands are only valid if the sub-type **static** option is configured, implying that **lsp-name** refers to an MPLS-TP tunnel LSP:

**path-type**. Values: active, working, protect. Default: active.

**dest-global-id** <global-id> **dest-node-id** <node-id>: Default: the **to** global-id:node-id from the LSP ID.

**control-channel**: If this is set to none, then IP encapsulation over an LSP is used with a destination address in the 127/8 range. The source address is set to the system IP address, unless the user specifies a source address using the **src-ip-address** option. If this is set to non-ip, then non-IP encapsulation over a G-ACh with channel type 0x00025 is used. This is the default for sub-type static. Note that the encapsulation used for the echo reply is the same as the encapsulation used for the echo request.

**downstream-map-tlv**: LSP Trace commands with this option can only be executed if the control-channel is set to none. The DSMAP/DDMAP TLV is only included in the echo request message if the egress interface is either a numbered IP interface, or an unnumbered IP interface. The TLV will not be included if the egress interface is of type **unnumbered-mpls-tp**.

For **lsp-ping**, the **dest-node-id** may be entered as a 4-octet IP address <a.b.c.d> or 32-bit integer <1..4294967295>. For **lsp-trace**, the destination **node-id** and **global-id** are taken from the spoke-sdp context.

The send mode and reply mode are always taken to be an application level control channel for MPLS-TP.

The **force** parameter causes an LSP ping echo request to be sent on an LSP that has been brought oper-down by BFD (LSP-Ping echo requests would normally be dropped on oper-down LSPs). This parameter is not applicable to SAA.

The LSP ID used in the LSP Ping packet is derived from a context lookup based on **lsp-name** and **path-type** (active/working/protect).

**Dest-global-id** and **dest-node-id** refer to the target global/node id. They do not need to be entered for end-to-end ping and trace, and the system will use the destination global id and node id from the LSP ID.

The same command syntax is applicable for SAA tests configured under **configure>saa>test**.

## MPLS-TP Show Commands

### Static MPLS Labels

The following new commands show the details of the static MPLS labels.

**show>router>mpls-labels>label <start-label> [<end-label> [in-use|<label-owner>]]**

**show>router>mpls-labels>label-range**

An example output is as follows:

```
*A:mlstp-dutA# show router mpls
mpls mpls-labels
*A:mlstp-dutA# show router mpls label
label label-range
*A:mlstp-dutA# show router mpls label-range

=====
Label Ranges
=====
Label Type Start Label End Label Aging Total Available

Static-lsp 32 16415 - 16364
Static-svc 16416 32799 - 16376
Dynamic 32800 131071 0 98268
=====
```

### MPLS-TP Tunnel Configuration

These should show the configuration of a given tunnel.

**show>router>mpls>tp-lsp**

A sample output is as follows:

```
*A:mlstp-dutA# show router mpls tp-lsp
- tp-lsp [<lsp-name>] [status {up|down}] [from <ip-address>|to <ip-address>]
 [detail]
- tp-lsp [<lsp-name>] path [protect|working] [detail]
- tp-lsp [<lsp-name>] protection

<lsp-name> : [32 chars max] - accepts * as wildcard char
<path> : keyword - Display LSP path information.
<protection> : keyword - Display LSP protection information.
<up|down> : keywords - Specify state of the LSP
```

```

<ip-address> : a.b.c.d
<detail> : keyword - Display detailed information.
*A:mlstp-dutA# show router mpls tp-lsp
path
protection
to <a.b.c.d>
<lsp-name>
 "lsp-32" "lsp-33" "lsp-34" "lsp-35" "lsp-36" "lsp-37" "lsp-38" "lsp-39"
 "lsp-40" "lsp-41"
status {up|down}
from <ip-address>
detail

*A:mlstp-dutA# show router mpls tp-lsp "lsp-
"lsp-32" "lsp-33" "lsp-34" "lsp-35" "lsp-36" "lsp-37" "lsp-38" "lsp-39"
"lsp-40" "lsp-41"
*A:mlstp-dutA# show router mpls tp-lsp "lsp-32"

=====
MPLS MPLS-TP LSPs (Originating)
=====
LSP Name To Tun Protect Adm Opr
 Id Id Path

lsp-32 0.0.3.234 32 No Up Up

LSPs : 1
=====
*A:mlstp-dutA# show router mpls tp-lsp "lsp-32" detail

=====
MPLS MPLS-TP LSPs (Originating) (Detail)
=====
Type : Originating

LSP Name : lsp-32
LSP Type : MplsTp
From Node Id : 0.0.3.233+
Adm State : Up
LSP Up Time : 0d 04:50:47
Transitions : 1
DestGlobalId : 42

LSP Tunnel ID : 32
To Node Id : 0.0.3.234
Oper State : Up
LSP Down Time : 0d 00:00:00
Path Changes : 2
DestTunnelNum : 32

```

## MPLS-TP Path configuration.

This can reuse and augment the output of the current show commands for static LSPs. They should also show if BFD is enabled on a given path. If this referring to a transit path, this should also display (among others) the path-id (7 parameters) for a given transit-path-name, or the transit-path-name for a given the path-id (7 parameters)

```
show>router>mpls>tp-lsp>path
```

## MPLS-TP Show Commands

A sample output is as follows:

```

=====
*A:mlstp-dutA# show router mpls tp-lsp path
=====
MPLS-TP LSP Path Information
=====
LSP Name : lsp-32 To : 0.0.3.234
Admin State : Up Oper State : Up

Path NextHop InLabel OutLabel Out I/F Admin Oper

Working 32 32 32 AtoB_1 Up Down
Protect 2080 2080 2080 AtoC_1 Up Up
=====
LSP Name : lsp-33 To : 0.0.3.234
Admin State : Up Oper State : Up

Path NextHop InLabel OutLabel Out I/F Admin Oper

Working 33 33 33 AtoB_1 Up Down
Protect 2082 2082 2082 AtoC_1 Up Up
=====
LSP Name : lsp-34 To : 0.0.3.234
Admin State : Up Oper State : Up

Path NextHop InLabel OutLabel Out I/F Admin Oper

Working 34 34 34 AtoB_1 Up Down
Protect 2084 2084 2084 AtoC_1 Up Up
=====
LSP Name : lsp-35 To : 0.0.3.234
Admin State : Up Oper State : Up

Path NextHop InLabel OutLabel Out I/F Admin Oper

Working 35 35 35 AtoB_1 Up Down
Protect 2086 2086 2086 AtoC_1 Up Up
=====
LSP Name : lsp-36 To : 0.0.3.234
Admin State : Up Oper State : Up

Path NextHop InLabel OutLabel Out I/F Admin Oper

Working 36 36 36 AtoB_1 Up Down
Protect 2088 2088 2088 AtoC_1 Up Up
=====
LSP Name : lsp-37 To : 0.0.3.234
Admin State : Up Oper State : Up

```

```

Path NextHop InLabel OutLabel Out I/F Admin Oper

Working 37 37 AtoB_1 Up Down
Protect 2090 2090 AtoC_1 Up Up

```

```

=====
LSP Name : lsp-38 To : 0.0.3.234
Admin State : Up Oper State : Up

```

```

Path NextHop InLabel OutLabel Out I/F Admin Oper

Working 38 38 AtoB_1 Up Down
Protect 2092 2092 AtoC_1 Up Up

```

```

=====
LSP Name : lsp-39 To : 0.0.3.234
Admin State : Up Oper State : Up

```

```

Path NextHop InLabel OutLabel Out I/F Admin Oper

Working 39 39 AtoB_1 Up Down
Protect 2094 2094 AtoC_1 Up Up

```

```

=====
LSP Name : lsp-40 To : 0.0.3.234
Admin State : Up Oper State : Up

```

```

Path NextHop InLabel OutLabel Out I/F Admin Oper

Working 40 40 AtoB_1 Up Down
Protect 2096 2096 AtoC_1 Up Up

```

```

=====
LSP Name : lsp-41 To : 0.0.3.234
Admin State : Up Oper State : Up

```

```

Path NextHop InLabel OutLabel Out I/F Admin Oper

Working 41 41 AtoB_1 Up Down
Protect 2098 2098 AtoC_1 Up Up

```

```
*A:mlstp-dutA# show router mpls tp-lsp "lsp-32" path working
```

```

=====
MPLS-TP LSP Working Path Information
LSP: "lsp-32"

```

```

=====
LSP Name : lsp-32 To : 0.0.3.234
Admin State : Up Oper State : Up

```

```

Path NextHop InLabel OutLabel Out I/F Admin Oper

Working 32 32 AtoB_1 Up Down

```

```
*A:mlstp-dutA# show router mpls tp-lsp "lsp-32" path protect
```

## MPLS-TP Show Commands

```

=====
MPLS-TP LSP Protect Path Information
 LSP: "lsp-32"
=====
LSP Name : lsp-32 To : 0.0.3.234
Admin State : Up Oper State : Up

Path NextHop InLabel OutLabel Out I/F Admin Oper

Protect 2080 2080 AtoC_1 Up Up
=====
*A:mlstp-dutA# show router mpls tp-lsp "lsp-32" path protect detail

=====
MPLS-TP LSP Protect Path Information
 LSP: "lsp-32" (Detail)
=====
LSP Name : lsp-32 To : 0.0.3.234
Admin State : Up Oper State : Up

Protect path information

Path Type : Protect LSP Num : 2
Path Admin : Up Path Oper : Up
Out Interface : AtoC_1 Next Hop Addr : n/a
In Label : 2080 Out Label : 2080
Path Up Time : 0d 04:52:17 Path Dn Time : 0d 00:00:00
Active Path : Yes Active Time : 0d 00:52:56

MEP information
MEP State : Up BFD : cc
OAM Templ : privatebed-oam-template CC Status : inService
CV Status : unknown
Protect Templ : privatebed-protection-template WTR Count Down: 0 seconds
RX PDU : SF (1,1) TX PDU : SF (1,1)
Defects :

=====
*A:mlstp-dutA# show router mpls tp-lsp "lsp-32" path working detail

=====
MPLS-TP LSP Working Path Information
 LSP: "lsp-32" (Detail)
=====
LSP Name : lsp-32 To : 0.0.3.234
Admin State : Up Oper State : Up

Working path information

Path Type : Working LSP Num : 1
Path Admin : Up Path Oper : Down
Down Reason : ccFault ifDn
Out Interface : AtoB_1 Next Hop Addr : n/a
In Label : 32 Out Label : 32
Path Up Time : 0d 00:00:00 Path Dn Time : 0d 00:53:01
Active Path : No Active Time : n/a

```



```

MEP information
MEP State : Up
OAM Templ : privatebed-oam-template
BFD : cc
CC Status : outOfService
CV Status : unknown
=====
*A:mlstp-dutA#

```

## MPLS-TP Protection

These should show the protection configuration for a given tunnel, which path in a tunnel is currently working and which is protect, and whether the working or protect is currently active.

**show>router>mpls>tp-lsp>protection**

A sample output is as follows:

```

*A:mlstp-dutA# show router mpls tp-lsp protection
=====
MPLS-TP LSP Protection Information
Legend: W-Working, P-Protect,
=====
LSP Name Admin Oper Path Ingr/Egr Act. Rx PDU
 State State State Label Path Tx PDU

lsp-32 Up Up W Down 32/32 No SF (1,1)
 P Up 2080/2080 Yes SF (1,1)
lsp-33 Up Up W Down 33/33 No SF (1,1)
 P Up 2082/2082 Yes SF (1,1)
lsp-34 Up Up W Down 34/34 No SF (1,1)
 P Up 2084/2084 Yes SF (1,1)
lsp-35 Up Up W Down 35/35 No SF (1,1)
 P Up 2086/2086 Yes SF (1,1)
lsp-36 Up Up W Down 36/36 No SF (1,1)
 P Up 2088/2088 Yes SF (1,1)
lsp-37 Up Up W Down 37/37 No SF (1,1)
 P Up 2090/2090 Yes SF (1,1)
lsp-38 Up Up W Down 38/38 No SF (1,1)
 P Up 2092/2092 Yes SF (1,1)
lsp-39 Up Up W Down 39/39 No SF (1,1)
 P Up 2094/2094 Yes SF (1,1)
lsp-40 Up Up W Down 40/40 No SF (1,1)
 P Up 2096/2096 Yes SF (1,1)
lsp-41 Up Up W Down 41/41 No SF (1,1)
 P Up 2098/2098 Yes SF (1,1)

No. of MPLS-TP LSPs: 10
=====

```

## BFD

The existing show>router>bfd context should be enhanced for MPLS-TP, as follows:

**show>router>bfd>mpls-tp-lsp**

Displays the MPLS –TP paths for which BFD is enabled.

**show>router>bfd>session [src <ip-address> [dest <ip-address> | detail]]|[mpls-tp-path <lsp-id...> [detail]]**

Should be enhanced to show the details of the BFD session on a particular MPLS-TP path, where lsp-id is the fully qualified lsp-id to which the BFD session is in associated.

A sample output is as follows:

```
*A:mlstp-dutA# show router bfd
- bfd

 bfd-template - Display BFD Template information
 interface - Display Interfaces with BFD
 session - Display session information

*A:mlstp-dutA# show router bfd bfd-template "privatebed-bfd-template"

=====
BFD Template privatebed-bfd-template
=====
Template Name : privatebed-* Template Type : cpmNp
Transmit Timer : 10 msec Receive Timer : 10 msec
CV Transmit Interval : 1000 msec
Template Multiplier : 3 Echo Receive Interval : 100 msec

Mpls-tp Association
privatebed-oam-template
=====
* indicates that the corresponding row element may have been truncated.
*A:mlstp-dutA# show router bfd session

=====
BFD Session
=====
Interface/Lsp Name State Tx Intvl Rx Intvl Multipl
 Remote Address/Info Protocols Tx Pkts Rx Pkts Type

wp::lsp-32 Down (1) 1000 1000 3
 0::0.0.0.0 mplsTp N/A N/A cpm-np
wp::lsp-33 Down (1) 1000 1000 3
 0::0.0.0.0 mplsTp N/A N/A cpm-np
wp::lsp-34 Down (1) 1000 1000 3
 0::0.0.0.0 mplsTp N/A N/A cpm-np
wp::lsp-35 Down (1) 1000 1000 3
```

```

0::0.0.0.0 mplsTp N/A N/A cpm-np
wp::lsp-36 Down (1) 1000 1000 3
0::0.0.0.0 mplsTp N/A N/A cpm-np
wp::lsp-37 Down (1) 1000 1000 3
0::0.0.0.0 mplsTp N/A N/A cpm-np
wp::lsp-38 Down (1) 1000 1000 3
0::0.0.0.0 mplsTp N/A N/A cpm-np
wp::lsp-39 Down (1) 1000 1000 3
0::0.0.0.0 mplsTp N/A N/A cpm-np
wp::lsp-40 Down (1) 1000 1000 3
0::0.0.0.0 mplsTp N/A N/A cpm-np
wp::lsp-41 Down (1) 1000 1000 3
0::0.0.0.0 mplsTp N/A N/A cpm-np
pp::lsp-32 Up (3) 1000 1000 3
0::0.0.0.0 mplsTp N/A N/A cpm-np
pp::lsp-33 Up (3) 1000 1000 3
0::0.0.0.0 mplsTp N/A N/A cpm-np
pp::lsp-34 Up (3) 1000 1000 3
0::0.0.0.0 mplsTp N/A N/A cpm-np
pp::lsp-35 Up (3) 1000 1000 3
0::0.0.0.0 mplsTp N/A N/A cpm-np
pp::lsp-36 Up (3) 1000 1000 3
0::0.0.0.0 mplsTp N/A N/A cpm-np
pp::lsp-37 Up (3) 1000 1000 3
0::0.0.0.0 mplsTp N/A N/A cpm-np
pp::lsp-38 Up (3) 1000 1000 3
0::0.0.0.0 mplsTp N/A N/A cpm-np
pp::lsp-39 Up (3) 1000 1000 3
0::0.0.0.0 mplsTp N/A N/A cpm-np
pp::lsp-40 Up (3) 1000 1000 3
0::0.0.0.0 mplsTp N/A N/A cpm-np
pp::lsp-41 Up (3) 1000 1000 3
0::0.0.0.0 mplsTp N/A N/A cpm-np

No. of BFD sessions: 20

wp = Working path pp = Protecting path
=====

```

## MPLS TP Node Configuration

Displays the Global ID, Node ID and other general MPLS-TP configurations for the node.

**show>router>mpls>mpls-tp**

A sample output is as follows:

```
*A:mlstp-dutA# show router mpls mpls-tp
- mpls-tp
```

```
oam-template - Display MPLS-TP OAM Template information
protection-tem* - Display MPLS-TP Protection Template information
```

## MPLS-TP Show Commands

status - Display MPLS-TP system configuration  
transit-path - Display MPLS-TP Tunnel information

```
*A:mlstp-dutA# show router mpls mpls-tp oam-template
```

```
=====
MPLS-TP OAM Templates
=====
```

```
Template Name : privatebed-oam-template Router ID : 1
BFDD Template : privatebed-bfd-template Hold-Down Time: 0 centiseconds
Hold-Up Time : 20 deciseconds
```

```
*A:mlstp-dutA# show router mpls mpls-tp protection-template
```

```
=====
MPLS-TP Protection Templates
=====
```

```
Template Name : privatebed-protection-template Router ID : 1
Protection Mode: one2one Direction : bidirectional
Revertive : revertive Wait-to-Restore: 300sec
Rapid-PSC-Timer: 10ms Slow-PSC-Timer : 5sec
```

```
*A:mlstp-dutA# show router mpls mpls-tp status
```

```
=====
MPLS-TP Status
=====
```

```
Admin Status : Up
Global ID : 42 Node ID : 0.0.3.233
Tunnel Id Min : 1 Tunnel Id Max : 4096
```

```
*A:mlstp-dutA# show router mpls mpls-tp transit-path
```

```
- transit-path [<path-name>] [detail]
```

```
<path-name> : [32 chars max]
<detail> : keyword - Display detailed information.
```

```
A:mlstp-dutC# show router mpls mpls-tp transit-path
```

```
- transit-path [<path-name>] [detail]
```

```
<path-name> : [32 chars max]
<detail> : keyword - Display detailed information.
```

```
A:mlstp-dutC# show router mpls mpls-tp transit-path
```

```
<path-name>
"tp-32" "tp-33" "tp-34" "tp-35" "tp-36" "tp-37" "tp-38" "tp-39"
"tp-40" "tp-41"
detail
```

```
A:mlstp-dutC# show router mpls mpls-tp transit-path "tp-32"
```

```

MPLS-TP Transit tp-32 Path Information
=====
Path Name : tp-32
Admin State : Up Oper State : Up

Path NextHop InLabel OutLabel Out I/F

FP 2080 2081 CtoB_1
RP 2081 2080 CtoA_1
=====
A:mplstp-dutC# show router mpls mpls-tp transit-path "tp-32" detail

=====
MPLS-TP Transit tp-32 Path Information (Detail)
=====
Path Name : tp-32
Admin State : Up Oper State : Up

Path ID configuration
Src Global ID : 42 Dst Global ID : 42
Src Node ID : 0.0.3.234 Dst Node ID : 0.0.3.233
LSP Number : 2 Dst Tunnel Num: 32

Forward Path configuration
In Label : 2080 Out Label : 2081
Out Interface : CtoB_1 Next Hop Addr : n/a

Reverse Path configuration
In Label : 2081 Out Label : 2080
Out Interface : CtoA_1 Next Hop Addr : n/a
=====
A:mplstp-dutC#

```

## MPLS-TP Interfaces

The existing show>router>interface command should be enhanced to display mpls-tp specific information.

The following is a sample output:

```

*A:mlstp-dutA# show router interface "AtoB_1"

=====
Interface Table (Router: Base)
=====
Interface-Name Adm Opr (v4/v6) Mode Port/SapId
IP-Address PfxState

AtoB_1 Down Down/-- Network 1/2/3:1
 Unnumbered If[system] n/a

Interfaces : 1

```

## Services using MPLS-TP PWs

The show>service command should be updated to display MPLS-TP specific information such as the PW Path ID and control channel status signaling parameters.

The following is a sample output:

```
*A:mlstp-dutA# show service id 1 all

=====
Service Detailed Information
=====
Service Id : 1 Vpn Id : 0
Service Type : Epipe
Name : (Not Specified)
Description : (Not Specified)
Customer Id : 1 Creation Origin : manual
Last Status Change: 12/03/2012 15:26:20
Last Mgmt Change : 12/03/2012 15:24:57
Admin State : Up Oper State : Up
MTU : 1514
Vc Switching : False
SAP Count : 1 SDP Bind Count : 1
Per Svc Hashing : Disabled
Force QTag Fwd : Disabled

ETH-CFM service specifics

Tunnel Faults : ignore

Service Destination Points (SDPs)

Sdp Id 32:1 - (0.0.3.234:42)

Description : (Not Specified)
SDP Id : 32:1 Type : Spoke
Spoke Descr : (Not Specified)
VC Type : Ether VC Tag : n/a
Admin Path MTU : 0 Oper Path MTU : 9186
Delivery : MPLS
Far End : 0.0.3.234:42
Tunnel Far End : n/a LSP Types : MPLSTP
Hash Label : Disabled Hash Lbl Sig Cap : Disabled
Oper Hash Label : Disabled

Admin State : Up Oper State : Up
Acct. Pol : None Collect Stats : Disabled
Ingress Label : 16416 Egress Label : 16416
```

```

Ingr Mac Fltr-Id : n/a
Ingr IP Fltr-Id : n/a
Ingr IPv6 Fltr-Id : n/a
Admin ControlWord : Preferred
Admin BW(Kbps) : 0
Last Status Change : 12/03/2012 15:26:20
Last Mgmt Change : 12/03/2012 15:24:57
Endpoint : N/A
PW Status Sig : Enabled
Class Fwding State : Down
Flags : None
Local Pw Bits : None
Peer Pw Bits : None
Peer Fault Ip : None
Peer Vccv CV Bits : None
Peer Vccv CC Bits : None
Application Profile : None
Standby Sig Slave : False
Block On Peer Fault : False

```

```

Egr Mac Fltr-Id : n/a
Egr IP Fltr-Id : n/a
Egr IPv6 Fltr-Id : n/a
Oper ControlWord : True
Oper BW(Kbps) : 0
Signaling : None
Force Vlan-Vc : Disabled
Precedence : 4

```

```

Ingress Qos Policy : (none)
Ingress FP QGrp : (none)
Ing FP QGrp Inst : (none)

```

```

Egress Qos Policy : (none)
Egress Port QGrp : (none)
Egr Port QGrp Inst : (none)

```

```

Statistics :
I. Fwd. Pkts. : 272969957
E. Fwd. Pkts. : 273017433

```

```

I. Dro. Pkts. : 0
E. Fwd. Octets : 16381033352

```

-----  
Control Channel Status  
-----

```

PW Status : enabled
Peer Status Expire : false

```

```

Refresh Timer : 66 secs
Clear On Timeout : true

```

-----  
SDP-BIND PW Path Information  
-----

```

AGI : 1:1
SAII Type2 : 42:0.0.3.234:1
TAII Type2 : 42:0.0.3.233:1

```

-----  
RSVP/Static LSPs  
-----

Associated LSP List :

```

Lsp Name : lsp-32
Admin State : Up

```

```

Oper State : Up

```

```

*A:mlstp-dutA# show service id [1..4] all | match "Control Channel" pre-lines 1 post-lines
5

```

-----  
Control Channel Status  
-----

```

PW Status : enabled
Peer Status Expire : false

```

```

Refresh Timer : 66 secs
Clear On Timeout : true

```

## MPLS-TP DEBUG COMMANDS

```

Control Channel Status

PW Status : enabled Refresh Timer : 66 secs
Peer Status Expire : false Clear On Timeout : true

Control Channel Status

PW Status : enabled Refresh Timer : 66 secs
Peer Status Expire : false Clear On Timeout : true

Control Channel Status

PW Status : enabled Refresh Timer : 66 secs
Peer Status Expire : false Clear On Timeout : true

*A:mlstp-dutA# show service id [1..4] all | match SDP-BIND pre-lines 1 post-lines 5

SDP-BIND PW Path Information

AGI : 1:1
SAII Type2 : 42:0.0.3.234:1
TAII Type2 : 42:0.0.3.233:1

SDP-BIND PW Path Information

AGI : 1:2
SAII Type2 : 42:0.0.3.234:2
TAII Type2 : 42:0.0.3.233:2

SDP-BIND PW Path Information

AGI : 1:3
SAII Type2 : 42:0.0.3.234:3
TAII Type2 : 42:0.0.3.233:3

SDP-BIND PW Path Information

AGI : 1:4
SAII Type2 : 42:0.0.3.234:4
TAII Type2 : 42:0.0.3.233:4

```

---

## MPLS-TP DEBUG COMMANDS

The following command provides the debug command for an MPLS-TP tunnel:



**tools>dump>router>mpls>tp-tunnel <lsp-name> [clear]**

The following is a sample output:

```
A:mlstp-dutA# tools dump router mpls tp-tunnel
- tp-tunnel <lsp-name> [clear]
- tp-tunnel id <tunnel-id> [clear]
<lsp-name> : [32 chars max]
<tunnel-id> : [1..61440]
<clear> : keyword - clear stats after reading
*A:mlstp-dutA# tools dump router mpls tp-tunnel "lsp-
"lsp-32" "lsp-33" "lsp-34" "lsp-35" "lsp-36" "lsp-37" "lsp-38" "lsp-39"
"lsp-40" "lsp-41"
*A:mlstp-dutA# tools dump router mpls tp-tunnel "lsp-32"
Idx: 1-32 (Up/Up): pgId 4, paths 2, operChg 1, Active: Protect
TunnelId: 42::0.0.3.233::32-42::0.0.3.234::32
PgState: Dn, Cnt/Tm: Dn 1/000 04:00:48.160 Up:3/000 00:01:25.840
MplsMsg: tpDn 0/000 00:00:00.000, tunDn 0/000 00:00:00.000
wpDn 0/000 00:00:00.000, ppDn 0/000 00:00:00.000
wpDel 0/000 00:00:00.000, ppDel 0/000 00:00:00.000
tunUp 1/000 00:00:02.070
Paths:
Work (Up/Dn): Lsp 1, Lbl 32/32, If 2/128 (1/2/3 : 0.0.0.0)
Tmpl: ptc: , oam: privatebed-oam-template (bfd: privatebed-bfd-template(np)-10 ms)
Bfd: Mode CC state Dn/Up handle 160005/0
Bfd-CC (Cnt/Tm): Dn 1/000 04:00:48.160 Up:1/000 00:01:23.970
State: Admin Up (1::1::1) port Up , if Dn , operChg 2
DnReasons: ccFault ifDn
Protect (Up/Up): Lsp 2, Lbl 2080/2080, If 3/127 (5/1/1 : 0.0.0.0)
Tmpl: ptc: privatebed-protection-template, oam: privatebed-oam-template (bfd: privatebed-
bfd-template(np)-10 ms)
Bfd: Mode CC state Up/Up handle 160006/0
Bfd-CC (Cnt/Tm): Dn 0/000 00:00:00.000 Up:1/000 00:01:25.410
State: Admin Up (1::1::1) port Up , if Up , operChg 1
Aps: Rx - 5, raw 3616, nok 0(), txRaw - 3636, revert Y
Pdu: Rx - 0x1a-21::0101 (SF), Tx - 0x1a-21::0101 (SF)
State: PF:W:L LastEvt pdu (L-SFw/R-SFw)
Tmrs: slow
Defects: None Now: 000 05:02:19.130
Seq Event state TxPdu RxPdu Dir Act Time
=== =====
000 start UA:P:L SF (0,0) NR (0,0) Tx--> Work 000 00:00:02.080
001 pdu UA:P:L SF (0,0) SF (0,0) Rx<-- Work 000 00:01:24.860
002 pdu UA:P:L SF (0,0) NR (0,0) Rx<-- Work 000 00:01:26.860
003 pUp NR NR (0,0) NR (0,0) Tx--> Work 000 00:01:27.440
004 pdu NR NR (0,0) NR (0,0) Rx<-- Work 000 00:01:28.760
005 wDn PF:W:L SF (1,1) NR (0,0) Tx--> Prot 000 04:00:48.160
006 pdu PF:W:L SF (1,1) NR (0,1) Rx<-- Prot 000 04:00:48.160
007 pdu PF:W:L SF (1,1) SF (1,1) Rx<-- Prot 000 04:00:51.080
```

The following command shows the free MPLS tunnel ID's

```
A:SASR1# /tools dump router mpls mpls-tp check-lbl-range
- mpls-tp check-lbl-range <range1> <range2>

<check-lbl-range> : keyword
```

## MPLS-TP DEBUG COMMANDS

```
<range1> : [32..65520]
<range2> : [32..65520]
```

The following command provides a debug tool to view control-channel-status signaling packets.

```
*A:7210SAS# /debug service id 700 sdp 200:700 event-type ?{config-change|oper-status-
change|neighbor-discovery|control-channel-status}
```

```
*A:7210SAS# /debug service id 700 sdp 200:700 event-type control-channel-status
```

```
*A:7210SAS#
```

```
1 2012/08/31 09:56:12.09 EST MINOR: DEBUG #2001 Base PW STATUS SIG PKT (RX):
```

```
"PW STATUS SIG PKT (RX)::
```

```
Sdp Bind 200:700 Instance 3
```

```
Version : 0x0
PW OAM Msg Type : 0x27
Refresh Time : 0xa
Total TLV Length : 0x8
Flags : 0x0
TLV Type : 0x96a
TLV Len : 0x4
PW Status Bits : 0x0
```

```
"
```

```
2 2012/08/31 09:56:22.09 EST MINOR: DEBUG #2001 Base PW STATUS SIG PKT (RX):
```

```
"PW STATUS SIG PKT (RX)::
```

```
Sdp Bind 200:700 Instance 3
```

```
Version : 0x0
PW OAM Msg Type : 0x27
Refresh Time : 0xa
Total TLV Length : 0x8
Flags : 0x0
TLV Type : 0x96a
TLV Len : 0x4
PW Status Bits : 0x0
```

```
"
```

```
3 2012/08/31 09:56:29.44 EST MINOR: DEBUG #2001 Base PW STATUS SIG PKT (TX):
```

```
"PW STATUS SIG PKT (TX)::
```

```
Sdp Bind 200:700 Instance 3
```

```
Version : 0x0
PW OAM Msg Type : 0x27
Refresh Time : 0x1e
Total TLV Length : 0x8
Flags : 0x0
TLV Type : 0x96a
TLV Len : 0x4
PW Status Bits : 0x0
```

## IP Performance Monitoring (IP PM)

The 7210 SAS OS supports Two-Way Active Measurement Protocol (TWAMP) and Two-Way Active Measurement Protocol Light (TWAMP Light).

This section describes the following topics:

- [Two-Way Active Measurement Protocol \(TWAMP\) on page 131](#)
- [Two-Way Active Measurement Protocol Light \(TWAMP Light\) on page 132](#)

### Two-Way Active Measurement Protocol (TWAMP)

Two-Way Active Measurement Protocol (TWAMP) provides a standards-based method for measuring the round-trip IP performance (packet loss, delay and jitter) between two devices. TWAMP uses the methodology and architecture of One-Way Active Measurement Protocol (OWAMP) to define a way to measure two-way or round-trip metrics.

There are four logical entities in TWAMP:

- The control-client
- The session-sender
- The server
- The session-reflector.

The control-client and session-sender are typically implemented in one physical device (the “client”) and the server and session-reflector in a second physical device (the “server”) with which the two-way measurements are being performed. The 7210 SAS acts as the server. The control-client and server establishes a TCP connection and exchange TWAMP-Control messages over this connection. When the control-client requires to start testing, the client communicates the test parameters to the server. If the server corresponds to conduct the described tests, the test begins as soon as the client sends a Start-Sessions message. As part of a test, the session sender sends a stream of UDP-based test packets to the session-reflector, and the session reflector responds to each received packet with a response UDP-based test packet. When the session-sender receives the response packets from the session-reflector, the information is used to calculate two-way delay, packet loss, and packet delay variation between the two devices.

### Configuration Notes

The following are the configuration notes:

## Two-Way Active Measurement Protocol Light (TWAMP Light)

- Unauthenticated mode is supported. Encrypted and Authenticated modes are not supported.
- TWAMP is supported only in the base router instance.
- By default, the 7210 SAS uses TCP port number 862 to listen for TWAMP control connections. This is not user configurable.

## Two-Way Active Measurement Protocol Light (TWAMP Light)

TWAMP Light is an optional model included in the TWAMP standard RFC5357 that uses standard TWAMP test packets but provides a lightweight approach to gathering ongoing IP delay performance data for base router and per-VPRN statistics. Full details are described in Appendix I of RFC 5357 (Active Two Way Measurement Protocol). The current implementation supports the TWAMP Light model for gathering delay statistics.

For TWAMP Light, the TWAMP Client/Server model is replaced with the Session Controller/Responder model. In general terms, the Session Controller is the launch point for the test packets and the Responder performs the reflection function.

TWAMP Light maintains the TWAMP test packet exchange but eliminates the TWAMP TCP control connection with local configurations; however, not all negotiated control parameters are replaced with local configuration. For example, CoS parameters communicated over the TWAMP control channel are replaced with a reply-in-kind approach. The reply-in-kind model reflects back the received CoS parameters, which are influenced by the reflector's QoS policies.

The responder function is configured under the **config>router>twamp-light** command hierarchy for base router reflection, and under the **config>service>vprn>twamp-light** command hierarchy for per VPRN reflection. The TWAMP Light reflector function is configured per context and must be activated before reflection can occur; the function is not enabled by default for any context. The reflector requires the operator to define the TWAMP Light UDP listening port that identifies the TWAMP Light protocol and the prefixes that the reflector will accept as valid sources for a TWAMP Light request. If the configured TWAMP Light listening UDP port is in use by another application on the system, a Minor OAM message will be presented indicating that the port is unavailable and that the activation of the reflector is not allowed. If the source IP address in the TWAMP Light packet arriving on the responder does not match a configured IP address prefix, the packet is dropped. Multiple prefix entries may be configured per context on the responder. An inactivity timeout under the **config>oam-test>twamp>twamp-light** hierarchy defines the amount of time the reflector will keep the individual reflector sessions active in the absence of test packets. A responder requires CPM3 or better hardware.

TWAMP Light test packet launching is controlled by the OAM Performance Monitoring (OAM-PM) architecture and adheres to those rules; this includes the assignment of a "Test-Id". TWAMP Light does not carry the 4-byte test ID in the packet in order to remain locally significant and

uniform with other protocols under the control of the OAM-PM architecture. The OAM-PM construct allow the various test parameters to be defined. These test parameters include the IP session-specific information which allocates the test to the specific routing instance, the source and destination IP address, the destination UDP port (which must match the listening UDP port on the reflector) and a number of other options that allow the operator to influence the packet handling. The probe interval and padding size can be configured under the specific session. The size of the all “0” padding can be included to ensure that the TWAMP packet is the same size in both directions. The TWAMP PDU definition does not accomplish symmetry by default. A pad size of 27 bytes will accomplish symmetrical TWAMP frame sizing in each direction.

The OAM-PM architecture does not perform any validation of the session information. The test will be allowed to be activated regardless of the validity of this information. For example, if the configured source IP address is not local within the router instance to which the test is allocated, the test will start sending TWAMP Light packets but will not receive any responses.

The OAM-PM section of this guide provides more information describing the integration of TWAMP Light and the OAM-PM architecture, including hardware dependencies.

The following is a summary of supported TWAMP Light functions:

- Base router instances for network interfaces and IES services
  - IPv6 addresses are supported only with IP interfaces that support IPv6, such as an IES IP interface in access-uplink and network mode. IPv6 is not supported for IP interfaces that do not support IPv6; for example, routed VPLS services in access-uplink mode do not support IPv6, and therefore TWAMP Light IPv6 sessions are not supported with it.
- Per-VPRN service context
- IPv4 and IPv6
  - Must be unicast
  - For IPv6, addresses cannot be a reserved or link local address
- Reflector prefix definition for acceptable TWAMP Light sources
  - Prefix list may be added and removed without shutting down the reflector function
  - If no prefixes are defined, the reflector will drop all TWAMP Light packets
- Integration with OAM-PM architecture capturing delay and loss measurement statistics
  - Not available from interactive CLI
  - Multiple test sessions can be configured between the same source and destination IP endpoints. The tuple of Source IP, Destination IP, Source UDP, and Destination UDP provide a unique index for each test.

The example below shows a basic configuration using TWAMP Light to monitor two IP endpoints in a VPRN, including the default TWAMP Light values that were not overridden with configuration entries.

## Two-Way Active Measurement Protocol Light (TWAMP Light)

The following example shows a reflector configuration:

```
config>test-oam>twamp>twamp-light# info detail

(default) inactivity-timeout 100

config>service>vprn# info

route-distinguisher 65535:500
auto-bind ldp
vrf-target target:65535:500
interface "to-cpe31" create
 address 10.1.1.1/30
 sap 1/1/2:500 create
 exit
exit
static-route 192.168.1.0/24 next-hop 10.1.1.2
bgp
 no shutdown
exit
twamp-light
 reflector udp-port 15000 create
 description "TWAMP Light reflector VPRN 500"
 prefix 10.2.1.1/32 create
 description "Process only 10.2.1.1 TWAMP Light Packets"
 exit
 prefix 172.16.1.0/24 create
 description "Process all 172.16.1.0 TWAMP Light packets"
 exit
 no shutdown
 exit
exit
no shutdown

```

The following example shows a session controller configuration:

```
config>service>vprn# info

route-distinguisher 65535:500
auto-bind ldp
vrf-target target:65535:500
interface "to-cpe28" create
 address 10.2.1.1/30
 sap 1/1/4:500 create
 exit
exit
static-route 192.168.2.0/24 next-hop 10.2.1.2
no shutdown

config>oam-pm>session# info detail

bin-group 2
meas-interval 15-mins create

```

```
 intervals-stored 8
 exit
 ip
 dest-udp-port 15000
 destination 10.1.1.1
 fc "12"
 (default) no forwarding
 profile in
 router 500
 source 10.2.1.1
 (default) ttl 255
 twamp-light test-id 500 create
 (default) interval 100
 (default) pad-size 0
 (default) no test-duration
 no shutdown
 exit
 exit
```

-----

---

## Ethernet Connectivity Fault Management (ETH-CFM)

The IEEE and the ITU-T have cooperated to define the protocols, procedures and managed objects to support service based fault management. Both IEEE 802.1ag standard and the ITU-T Y.1731 recommendation support a common set of tools that allow operators to deploy the necessary administrative constructs, management entities and functionality, Ethernet Connectivity Fault Management (ETH-CFM). The ITU-T has also implemented a set of advanced ETH-CFM and performance management functions and features that build on the proactive and on demand troubleshooting tools.

CFM uses Ethernet frames and is distinguishable by ether-type 0x8902. In certain cases the different functions will use a reserved multicast address that could also be used to identify specific functions at the MAC layer. However, the multicast MAC addressing is not used for every function or in every case. The Operational Code (OpCode) in the common CFM header is used to identify the type of function carried in the CFM packet. CFM frames are only processed by IEEE MAC bridges. With CFM, interoperability can be achieved between different vendor equipment in the service provider network upto and including customer premises bridges. The following table lists CFM-related acronyms used in this section.

IEEE 802.1ag and ITU-T Y.1731 functions that are implemented are available on the 7210 SAS platforms.

<b>Acronym</b>	<b>Callout</b>
IDM	One way Delay Measurement (Y.1731)
AIS	Alarm Indication Signal
CCM	Continuity check message
CFM	Connectivity fault management
DMM	Delay Measurement Message (Y.1731)
DMR	Delay Measurement Reply (Y.1731)
LBM	Loopback message
LBR	Loopback reply
LTM	Linktrace message
LTR	Linktrace reply
ME	Maintenance entity



<b>Acronym</b>	<b>Callout (Continued)</b>
MA	Maintenance association
MA-ID	Maintenance association identifier
MD	Maintenance domain
MEP	Maintenance association end point
MEP-ID	Maintenance association end point identifier
MHF	MIP half function
MIP	Maintenance domain intermediate point
OpCode	Operational Code
RDI	Remote Defect Indication
TST	Ethernet Test (Y.1731)

## ETH-CFM Building Blocks

The IEEE and the ITU-T use their own nomenclature when describing administrative contexts and functions. This introduces a level of complexity to configuration, discussion and different vendors naming conventions. The 7210 SAS OS CLI has chosen to standardize on the IEEE 802.1ag naming where overlap exists. ITU-T naming is used when no equivalent is available in the IEEE standard. In the following definitions, both the IEEE name and ITU-T names are provided for completeness, using the format IEEE Name/ITU-T Name.

Maintenance Domain (MD)/Maintenance Entity (ME) is the administrative container that defines the scope, reach and boundary for faults. It is typically the area of ownership and management responsibility. The IEEE allows for various formats to name the domain, allowing upto 45 characters, depending on the format selected. ITU-T supports only a format of “none” and does not accept the IEEE naming conventions.

0 — Undefined and reserved by the IEEE.

1 — No domain name. It is the only format supported by Y.1731 as the ITU-T specification does not use the domain name. This is supported in the IEEE 802.1ag standard but not in currently implemented for 802.1ag defined contexts.

2,3,4 — Provides the ability to input various different textual formats, upto 45 characters. The string format (2) is the default and therefore the keyword is not shown when looking at the configuration.

Maintenance Association (MA)/Maintenance Entity Group (MEG) is the construct where the different management entities will be contained. Each MA is uniquely identified by its MA-ID. The MA-ID is comprised of the by the MD level and MA name and associated format. This is another administrative context where the linkage is made between the domain and the service using the **bridging-identifier** configuration option. The IEEE and the ITU-T use their own specific formats. The MA short name formats (0-255) have been divided between the IEEE (0-31, 64-255) and the ITU-T (32-63), with five currently defined (1-4, 32). Even though the different standards bodies do not have specific support for the others formats a Y.1731 context can be configured using the IEEE format options.

1 (Primary VID) — Values 0 — 4094

2 (String) — Raw ASCII, excluding 0-31 decimal/0-1F hex (which are control characters) form the ASCII table

3 (2-octet integer) — 0 — 65535

4 (VPN ID) — Hex value as described in RFC 2685, *Virtual Private Networks Identifier*

32 (icc-format) — Exactly 13 characters from the ITU-T recommendation T.50.

**Note:** When a VID is used as the short MA name, 802.1ag will not support VLAN translation because the MA-ID must match all the MEPs. The default format for a short MA name is an integer. Integer value 0 means the MA is not attached to a VID. This is useful for VPLS services on 7210 SAS platforms because the VID is locally significant.

Maintenance Domain Level (MD Level)/Maintenance Entity Group Level (MEG Level) is the numerical value (0-7) representing the width of the domain. The wider the domain, higher the numerical value, the farther the ETH-CFM packets can travel. It is important to understand that the level establishes the processing boundary for the packets. Strict rules control the flow of ETH-CFM packets and are used to ensure proper handling, forwarding, processing and dropping of these packets. To keep it simple ETH-CFM packets with higher numerical level values will flow through MEPs on MIPs on SAPs configured with lower level values. This allows the operator to implement different areas of responsibility and nest domains within each other. Maintenance association (MA) includes a set of MEPs, each configured with the same MA-ID and MD level used verify the integrity of a single service instance.

Maintenance Endpoint (MEP)/MEG Endpoint (MEP) are the workhorses of ETH-CFM. A MEP is the unique identification within the association (0-8191). Each MEP is uniquely identified by the MA-ID, MEPID tuple. This management entity is responsible for initiating, processing and terminating ETH-CFM functions, following the nesting rules. MEPs form the boundaries which prevent the ETH-CFM packets from flowing beyond the specific scope of responsibility. A MEP has direction, up or down. Each indicates the directions packets will be generated; UP toward the switch fabric, down toward the SAP away from the fabric. Each MEP has an active and passive side. Packets that enter the active point of the MEP will be compared to the existing level and processed accordingly. Packets that enter the passive side of the MEP are passed transparently through the MEP. Each MEP contained within the same maintenance association and with the same level (MA-ID) represents points within a single service. MEP creation on a SAP is allowed only for Ethernet ports with NULL, q-tags, q-in-q encapsulations. MEPs may also be created on SDP bindings.

Maintenance Intermediate Point (MIP)/MEG Intermediate Point (MIP) are management entities between the terminating MEPs along the service path. These provide insight into the service path connecting the MEPs. MIPs only respond to Loopback Messages (LBM) and Linktrace Messages (LTM). All other CFM functions are transparent to these entities. Only one MIP is allowed per SAP or SDP. The creation of the MIPs can be done when the lower level domain is created (explicit). This is controlled by the use of the mhf-creation mode within the association under the bridge-identifier. MIP creation is supported on a SAP and SDP, not including Mesh SDP bindings. By default, no MIPs are created.

---

**NOTE:** The 7210 SAS platforms supports either only Ingress MIPs in some services or bi-directional MIPs (that is, ingress and egress MIPs) in some services. The table below lists the MIP and MEP support for different services on different platforms.

There are two locations in the configuration where ETH-CFM is defined. The domains, associations (including linkage to the service id), MIP creation method, common ETH-CFM functions and remote MEPs are defined under the top level **eth-cfm** command. It is important to note, when Y.1731 functions are required the context under which the MEPs are configured must follow the Y.1731 specific formats (domain format of none, MA format icc-format). Once these parameters have been entered, the MEP and possibly the MIP can be defined within the service under the SAP or SDP.

This is a general table that indicates the ETH-CFM support for the different services and endpoints. It is not meant to indicate the services that are supported or the requirements for those services on the individual platforms.

**Table 6: ETH-CFM Support Matrix for 7210 SAS-M Network Mode**

Service	Ethernet Connection Type	Down MEP	UP MEP	MIP	Primary VLAN
Epipe	SAP	Yes	Yes	Ingress MIP and Egress MIP	Not Supported
	SDP	Yes	Yes	Ingress MIP and Egress MIP	Not Supported
VPLS	SAP	Yes	Yes	Ingress MIP only	Not Supported
	Spoke-SDP	Yes	Yes	Ingress MIP only	Not Supported
	Mesh-SDP	Yes	Yes	Not supported	Not Supported
RVPLS	SAP	Not Supported	Not Supported	Not Supported	Not Supported
IES	IES IPv4 interface	Not Supported	Not Supported	Not Supported	Not Supported
PBB Epipe	I-SAP	Not Supported	Yes	Not Supported	Not Supported

Service	Ethernet Connection Type	Down MEP	UP MEP	MIP	Primary VLAN
PBB VPLS	I-SAP	Not Supported	Not Supported	Not Supported	Not Supported
PBB B-VPLS	B-SAP	Not Supported	Not Supported	Not Supported	Not Supported
IES	SAP	Not Supported	Not Supported	Not Supported	Not Supported
VPRN	SAP	Not Supported	Not Supported	Not Supported	Not Supported

**Table 7: ETH-CFM Support Matrix for 7210 SAS-M Access-Uplink Mode**

Service	Ethernet Connection Type	Down MEP	UP MEP	MIP	Primary VLAN
Epipe	SAP (Access and Access-uplink SAP)	Yes	Yes	Ingress MIP and Egress MIP	Not Supported
VPLS	SAP (Access and Access-uplink SAP)	Yes	Yes	Ingress MIP only	Not Supported
RVPLS	SAP	Not Supported	Not Supported	Not Supported	Not Supported
IES	IES IPv4 interface	Not Supported	Not Supported	Not Supported	Not Supported
	SAP	Not Supported	Not Supported	Not supported	Not Supported

**Table 8: ETH-CFM Support Matrix for 7210 SAS-T Access-Uplink Mode**

Service	Ethernet Connection Type	Down MEP	UP MEP	MIP	Primary VLAN
Epipe	SAP (Access and Access-uplink SAP)	Yes	Yes	Ingress MIP and Egress MIP	Not Supported
VPLS	SAP (Access and Access-uplink SAP)	Yes	Yes	Ingress MIP only	Not Supported
RVPLS	SAP	Not Supported	Not Supported	Not Supported	Not Supported
IES	IES IPv4 interface	Not Supported	Not Supported	Not Supported	Not Supported
	SAP	Not Supported	Not Supported	Not supported	Not Supported

**Table 9: ETH-CFM Support Matrix for 7210 SAS-T Network Mode**

Service	Ethernet Connection Type	Down MEP	UP MEP	MIP	Primary VLAN
Epipe	SAP	Yes	Yes	Ingress MIP and Egress MIP	Yes (only on Down MEP)
	SDP	Yes	Yes	Ingress MIP and Egress MIP	Not Supported
VPLS	SAP	Yes	Yes	Ingress MIP only	Yes (only on Down MEP)
	Spoke-SDP	Yes	Yes	Ingress MIP only	Not Supported
	Mesh-SDP	Yes	Yes	Not supported	Not Supported

Service	Ethernet Connection Type	Down MEP	UP MEP	MIP	Primary VLAN
RVPLS	SAP	Not Supported	Not Supported	Not Supported	Not Supported
IES	IES IPv4 interface	Not Supported	Not Supported	Not Supported	Not Supported
PBB Epipe	I-SAP	Not Supported	Yes	Not Supported	Not Supported
PBB VPLS	I-SAP	Not Supported	Not Supported	Not Supported	Not Supported
PBB B-VPLS	B-SAP	Not Supported	Not Supported	Not Supported	Not Supported
IES	SAP	Not Supported	Not Supported	Not Supported	Not Supported
VPRN	SAP	Not Supported	Not Supported	Not Supported	Not Supported

**Table 10: ETH-CFM Support Matrix for 7210 SAS-X devices**

Service	Ethernet Connection Type	Down MEP	UP MEP	MIP	Primary VLAN
Epipe	SAP	Yes	Yes	Ingress MIP and Egress MIP	Yes (only on Down MEP)
	SDP	Yes	Yes	Ingress MIP and Egress MIP	Not Supported

VPLS	SAP	Yes	Yes	Ingress MIP only	Yes (only on Down MEP)
	Spoke-SDP	Yes	Yes	Ingress MIP only	Not Supported
	Mesh-SDP	Yes	Yes	Not supported	Not Supported
RVPLS	SAP	Not Supported	Not Supported	Not Supported	Not Supported
IES	IES IPv4 interface	Not Supported	Not Supported	Not Supported	Not Supported
PBB Epipe	I-SAP	Not Supported	Yes	Not Supported	Not Supported
PBB VPLS	I-SAP	Not Supported	Not Supported	Not Supported	Not Supported
PBB B-VPLS	B-SAP	Not Supported	Not Supported	Not Supported	Not Supported
IES	SAP	Not Supported	Not Supported	Not Supported	Not Supported
VPRN	SAP	Not Supported	Not Supported	Not Supported	Not Supported



**Table 11: ETH-CFM Support Matrix for 7210 SAS-R6 and 7210 SAS-R12 devices**

Service	Ethernet Connection Type	Down MEP	UP MEP	MIP	Primary VLAN
Epipe	SAP	Yes	Yes	Ingress MIP and Egress MIP	Yes (only on Down MEP)
	SDP	Yes	Yes	Ingress MIP and Egress MIP	Not Supported
VPLS	SAP	Yes	Yes (only for IMMv2)	Ingress MIP only for IMMv1 Ingress MIP and Egress MIP for IMMv2	Yes only on IMMv2 (on Down MEP and MIP)
	Spoke-SDP	Yes	Yes (only for IMMv2)	Ingress MIP only	Not Supported
	Mesh-SDP	Not Supported	Yes (only for IMMv2)	Not Supported	Not Supported
RVPLS	SAP	Not Supported	Not Supported	Not Supported	Not Supported
IES	IES IPv4 interface	Not Supported	Not Supported	Not Supported	Not Supported
PBB Epipe	I-SAP	Not Supported	Not Supported	Not Supported	Not Supported
PBB VPLS	I-SAP	Not Supported	Not Supported	Not Supported	Not Supported

ETH-CFM Building Blocks

PBB B-VPLS	B-SAP	Not Supported	Not Supported	Not Supported	Not Supported
IES	SAP	Not Supported	Not Supported	Not Supported	Not Supported
VPRN	SAP	Not Supported	Not Supported	Not Supported	Not Supported

**Table 12: ETH-CFM Support Matrix for 7210 SAS-Mxp devices**

Service	Ethernet Connection Type	Down MEP	UP MEP	MIP	Primary VLAN
Epipe	SAP	Yes	Yes	Ingress MIP and Egress MIP	Yes (only on Down MEP)
	SDP	Yes	Yes	Ingress MIP and Egress MIP	Not Supported
VPLS	SAP	Yes	Yes	Ingress MIP	Yes (only on Down MEP)
	Spoke-SDP	Yes	Yes	Ingress MIP	Not Supported
	Mesh-SDP	Yes	Yes	Not Supported	Not Supported
RVPLS	SAP	Not Supported	Not Supported	Not Supported	Not Supported
IES	IES IPv4 interface	Not Supported	Not Supported	Not Supported	Not Supported
PBB Epipe	I-SAP	Not Supported	Not Supported	Not Supported	Not Supported
PBB VPLS	I-SAP	Not Supported	Not Supported	Not Supported	Not Supported
PBB B-VPLS	B-SAP	Not Supported	Not Supported	Not Supported	Not Supported
IES	SAP	Not Supported	Not Supported	Not Supported	Not Supported
VPRN	SAP	Not Supported	Not Supported	Not Supported	Not Supported

**Table 13: ETH-CFM Support Matrix for 7210 SAS-Sx/S 1/10GE devices**

Service	Ethernet Connection Type	Down MEP	UP MEP	MIP	Primary VLAN
Epipe	SAP	Yes	Yes	Ingress MIP and Egress MIP	Yes (only on Down MEP)
	SDP	Yes	Yes	Ingress MIP and Egress MIP	Not Supported
VPLS	SAP	Yes	Yes	Ingress MIP only	Yes (only on Down MEP)
	Spoke-SDP	Yes	Yes	Ingress MIP only	Not Supported
	Mesh-SDP	Yes	Yes	Not supported	Not Supported
RVPLS	SAP	Not Supported	Not Supported	Not Supported	Not Supported
IES	IES IPv4 interface	Not Supported	Not Supported	Not Supported	Not Supported
PBB Epipe	I-SAP	Not Supported	Not Supported	Not Supported	Not Supported
PBB VPLS	I-SAP	Not Supported	Not Supported	Not Supported	Not Supported

**Table 13: ETH-CFM Support Matrix for 7210 SAS-Sx/S 1/10GE devices (Continued)**

Service	Ethernet Connection Type	Down MEP	UP MEP	MIP	Primary VLAN
PBB B-VPLS	B-SAP	Not Supported	Not Supported	Not Supported	Not Supported
IES	SAP	Not Supported	Not Supported	Not Supported	Not Supported
VPRN	SAP	Not Supported	Not Supported	Not Supported	Not Supported

**Table 14: ETH-CFM Support Matrix for 7210 SAS-Sx 10/100GE devices**


Service	Ethernet Connection Type	Down MEP	UP MEP	MIP	Primary VLAN
Epipe	SAP	Yes	Yes	Ingress MIP and Egress MIP	Yes (only on Down MEP)
	SDP	Yes	Yes	Ingress MIP and Egress MIP	Not Supported
VPLS	SAP	Yes	Yes	Ingress MIP only	Yes (only on Down MEP)
	Spoke-SDP	Yes	Yes	Ingress MIP only	Not Supported
	Mesh-SDP	Yes	Yes	Not supported	Not Supported
RVPLS	SAP	Not Supported	Not Supported	Not Supported	Not Supported

<b>Service</b>	<b>Ethernet Connection Type</b>	<b>Down MEP</b>	<b>UP MEP</b>	<b>MIP</b>	<b>Primary VLAN</b>
IES	IES IPv4 interface	Not Supported	Not Supported	Not Supported	Not Supported
PBB Epipe	I-SAP	Not Supported	Not Supported	Not Supported	Not Supported
PBB VPLS	I-SAP	Not Supported	Not Supported	Not Supported	Not Supported
PBB B-VPLS	B-SAP	Not Supported	Not Supported	Not Supported	Not Supported
IES	SAP	Not Supported	Not Supported	Not Supported	Not Supported
VPRN	SAP	Not Supported	Not Supported	Not Supported	Not Supported

**NOTES:**

- Ethernet-Rings are not configurable under all service types. Any service restrictions for MEP direction or MIP support will override the generic capability of the Ethernet-Ring MPs. For more information on Ethernet-Rings, refer to the 7210 SAS Interfaces Guide.
  - On 7210 platforms in some services only ingress MIP functionality is supported. An ingress MIP or a Down MIP processes messages in the ingress direction when the OAM message is received on ingress of the SAP/Port (subject to the other checks). An egress MIP or an UP MIP refers to a MIP that processes message in the egress direction when the OAM message is being sent out of the SAP/port. For more information for service entities support ingress MIP or egress MIP or both, see the tables above.
  - On 7210 SAS devices, when two bidirectional MIPs are configured in an Epipe service on both the service entities/endpoints (For example: on both the SAP and the SDP configured in the Epipe service), only the MIP ingressing to the direction of linktrace messages responds. This is applicable to 7210 SAS platforms that support both ingress and egress MIPs (also referred to as Bi-directional MIPs).
-

### Configuration Notes on CCM timer support



**Figure 13: MEP and MIP**

Figure 14 shows the detailed IEEE representation of MEPs, MIPs, levels and associations, using the standards defined icons.



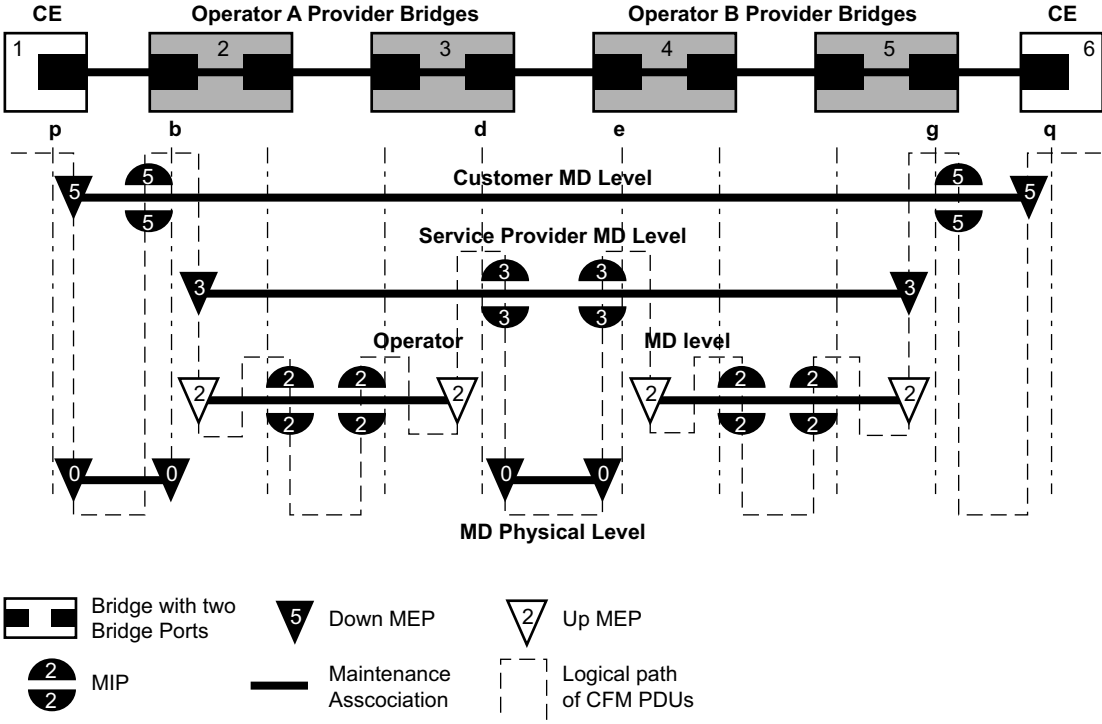



Figure 14: MEP, MIP and MD Levels

## Loopback

A loopback message is generated by an MEP to its peer MEP (Figure 15). The functions are similar to an IP ping to verify Ethernet connectivity between the nodes.



**Figure 15: CFM Loopback**

The following loopback-related functions are supported:


- Loopback message functionality on an MEP or MIP can be enabled or disabled.
- MEP — Supports generating loopback messages and responding to loopback messages with loopback reply messages.
- MIP — Supports responding to loopback messages with loopback reply messages when loopback messages are targeted to self.
- Displays the loopback test results on the originating MEP. There is a limit of ten outstanding tests per node.

## Linktrace

A linktrace message is originated by an MEP and targeted to a peer MEP in the same MA and within the same MD level (Figure 16). Its function is similar to IP traceroute. Traces a specific MAC address through the service. The peer MEP responds with a linktrace reply message after successful inspection of the linktrace message. The MIPs along the path also process the linktrace message and respond with linktrace replies to the originating MEP if the received linktrace message that has a TTL greater than 1 and forward the linktrace message if a look up of the target MAC address in the Layer 2 FIB is successful. The originating MEP shall expect to receive multiple linktrace replies and from processing the linktrace replies, it can put together the route to the target bridge.

A traced MAC address is carried in the payload of the linktrace message, the target MAC. Each MIP and MEP receiving the linktrace message checks whether it has learned the target MAC address. In order to use linktrace the target MAC address must have been learned by the nodes in the network. If so, a linktrace message is sent back to the originating MEP. Also, a MIP forwards the linktrace message out of the port where the target MAC address was learned.

The linktrace message itself has a multicast destination address. On a broadcast LAN, it can be received by multiple nodes connected to that LAN. But, at most, one node will send a reply.



Fig\_13

**Figure 16: CFM Linktrace**

The following linktrace related functions are supported:

- Enable or disables linktrace functions on an MEP.

- MEP — Supports generating linktrace messages and responding with linktrace reply messages.
- MIP — Supports responding to linktrace messages with linktrace reply messages when encoded TTL is greater than 1, and forward the linktrace messages accordingly if a lookup of the target MAC address in the Layer 2 FIB is successful.
- Displays linktrace test results on the originating MEP. There is a limit of ten outstanding tests per node. Storage is provided for upto ten MEPs and for the last ten responses. If more than ten responses are received older entries will be overwritten.

## Continuity Check (CC)

A Continuity Check Message (CCM) is a multicast frame that is generated by a MEP and multicast to all other MEPs in the same MA. The CCM does not require a reply message. To identify faults, the receiving MEP maintains an internal list of remote MEPs it should be receiving CCM messages from.

This list is based off of the remote-mepid configuration within the association the MEP is created in. When the local MEP does not receive a CCM from one of the configured remote MEPs within a pre-configured period, the local MEP raises an alarm.




Figure 17: CFM Continuity Check




Figure 18: CFM CC Failure Scenario

The following functions are supported:

- Enable and disable CC for an MEP
- Configure and delete the MEP entries in the CC MEP monitoring database manually. It is only required to provision remote MEPs. Local MEPs shall be automatically put into the database when they are created.
- CCM transmit interval configuration: 100ms(Supported only on 7210 SAS-M).  
When configuring MEPs with sub-second CCM intervals bandwidth consumption must be taken into consideration. Each CCM PDU is 100 bytes (800 bits). Taken individually this is a small value. However, the bandwidth consumption increases rapidly as multiple MEPs are configured with 10ms timers, 100 packets per second. Sub-second enabled MEPs are supported on the following:
  - On 7210 SAS-M, for Down MEPs configured on Ethernet SAPs.
  - 7210 SAS-M can have sub-second CCM configured on SDP.
  - Lowest MD-level, when multiple MEPs exist on same Ethernet SAP.
  - 7210 SAS-M can have multiple MEPs with sub-second intervals on the same SAP.
- CCM will declare a fault, when:
  - The CCM stops hearing from one of the remote MEPs for 3.5 times CC interval
  - Hears from a MEP with a LOWER MD level
  - Hears from a MEP that is not part of the local MEPs MA
  - Hears from a MEP that is in the same MA but not in the configured MEP list
  - Hears from a MEP in the same MA with the same MEP id as the receiving MEP
  - The CC interval of the remote MEP does not match the local configured CC interval
  - The remote MEP is declaring a fault
- An alarm is raised and a trap is sent if the defect is greater than or equal to the configured low-priority-defect value.
- Remote Defect Indication (RDI) is supported but by default is not recognized as a defect condition because the low-priority-defect setting default does not include RDI.

## Alarm Indication Signal (ETH-AIS Y.1731)

Alarm Indication Signal (AIS) provides an Y.1731 capable MEP the ability to signal a fault condition in the reverse direction of the MEP, out the passive side. When a fault condition is detected the MEP will generate AIS packets at the configured client levels and at the specified AIS interval until the condition is cleared. Currently a MEP configured to generate AIS must do so at a level higher than its own. The MEP configured on the service receiving the AIS packets is required to have the active side facing the receipt of the AIS packet and must be at the same level the AIS, The absence of an AIS packet for 3.5 times the AIS interval set by the sending node will clear the condition on the receiving MEP.

It is important to note that AIS generation is not supported to an explicitly configured endpoint. An explicitly configured endpoint is an object that contains multiple individual endpoints, as in PW redundancy.

---

## Test (ETH-TST Y.1731)

Ethernet test affords operators an Y.1731 capable MEP the ability to send an in service on demand function to test connectivity between two MEPs. The test is generated on the local MEP and the results are verified on the destination MEP. Any ETH-TST packet generated that exceeds the MTU will be silently dropped by the lower level processing of the node.

---

## Y.1731 Time Stamp Capability

Accurate results for one-way and two-way delay measurement tests using Y.1731 messages are obtained if the nodes are capable of time stamping packets in hardware.

- 7210 SAS-M and 7210 SAS-X support is as follows:
  - Y.1731 2-DM messages for both Down MEPs and UP MEPs, 1-DM for both Down MEPs and UP MEPs, and 2-SLM for both Down MEPs and UP MEPs use software based timestamps on Tx and hardware based timestamp on Rx. It uses the system clock (free-running or synchronized to NTP or synchronized to PTP, if the user configures “*configure system time ntp server ptp prefer*” command) to obtain the timestamps.
- 7210 SAS-Sx 10/100GE support is as follows:
  - Y.1731 2-DM messages for both Down MEPs and UP MEPs, 1-DM for both Down MEPs and UP MEPs, and 2-SLM for both Down MEPs and UP MEPs use software based timestamps on Tx and hardware based timestamp on Rx. It uses the system clock (free-running or synchronized to NTP) to obtain the timestamps.

- 7210 SAS-T (network mode), 7210 SAS-Sx 1/10GE, 7210 SAS-Mxp, 7210 SAS-R6 and 7210 SAS-R12 support is as follows:
  - Y.1731 2-DM messages for both Down MEPs and UP MEPs, 1-DM for both Down MEPs and UP MEPs, and 2-SLM for both Down MEPs and UP MEPs use software based timestamps on Tx and hardware based timestamp on Rx. The timestamps are obtained as given below:
    - From NTP, when NTP is enabled and PTP is disabled
    - From PTP, when PTP is enabled (irrespective of whether NTP is disabled or enabled)
    - From PTP, when PTP is enabled and NTP is configured to use PTP for system time
    - From free-running system time, when both NTP and PTP are disabled.
- The 7210 SAS-T (access-uplink mode) support is as follows:
  - Y.1731 2-DM messages for Down MEPs uses hardware timestamps for both Rx (packets received by the node) and Tx (packets sent out of the node). The timestamps is obtained from a free-running hardware clock. It provides accurate 2-way delay measurements and it is not recommended to use it for computing 1-way delay.
  - Y.1731 2-DM messages for UP MEPs, 1-DM for both Down MEPs and UP MEPs, and 2-SLM for both Down MEPs and UP MEPs use software based timestamps on Tx and hardware based timestamp on Rx. The timestamps are obtained as given below:
    - From NTP, when NTP is enabled and PTP is disabled
    - From PTP, when PTP is enabled (irrespective of whether NTP is disabled or enabled)
    - From PTP, when PTP is enabled and NTP is configured to use PTP for system time
    - From free-running system time, when both NTP and PTP are disabled.

**NOTE:** On 7210 SAS-T, Mxp, R6, R12 after PTP is enabled once, if the user needs to go back to NTP time scale, or system free-run time scale, a node reboot is required.

---



## Synthetic Loss Measurement (ETH-SL)

Nokia applied pre-standard OpCodes 53 (Synthetic Loss Reply) and 54 (Synthetic Loss Message) for the purpose of measuring loss using synthetic packets.

Notes: These will be changes to the assigned standard values in a future release. This means that the Release 4.0R6 is pre-standard and will not interoperate with future releases of SLM or SLR that supports the standard OpCode values.

This synthetic loss measurement approach is a single-ended feature that allows the operator to run on-demand and proactive tests to determine “in”, “out” loss and “unacknowledged” packets. This approach can be used between peer MEPs in both point to point and multipoint services. Only remote MEP peers within the association and matching the unicast destination will respond to the SLM packet.

The specification uses various sequence numbers in order to determine in which direction the loss occurred. Alcatel-Lucent has implemented the required counters to determine loss in each direction. In order to properly use the information that is gathered the following terms are defined;

- Count — The number of probes that are sent when the last frame is not lost. When the last frame(s) is or are lost, the count and unacknowledged equals the number of probes sent.
- Out-Loss (Far-end) — Packets lost on the way to the remote node, from test initiator to the test destination.
- In-Loss (Near-end) — Packet loss on the way back from the remote node to the test initiator.
- Unacknowledged — Number of packets at the end of the test that were not responded to.

The per probe specific loss indicators are available when looking at the on-demand test runs, or the individual probe information stored in the MIB. When tests are scheduled by Service Assurance Application (SAA) the per probe data is summarized and per probe information is not maintained. Any “unacknowledged” packets will be recorded as “in-loss” when summarized.

The on-demand function can be executed from CLI or SNMP. The on demand tests are meant to provide the carrier a way to perform on the spot testing. However, this approach is not meant as a method for storing archived data for later processing. The probe count for on demand SLM has a range of one to 100 with configurable probe spacing between one second and ten seconds. This means it is possible that a single test run can be upto 1000 seconds in length. Although possible, it is more likely the majority of on demand case can increase to 100 probes or less at a one second interval. A node may only initiate and maintain a single active on demand SLM test at any given time. A maximum of one storage entry per remote MEP is maintained in the results table.

Subsequent runs to the same peer can overwrite the results for that peer. This means, when using on demand testing the test should be run and the results checked prior to starting another test.

The proactive measurement functions are linked to SAA. This backend provides the scheduling, storage and summarization capabilities. Scheduling may be either continuous or periodic. It also allows for the interpretation and representation of data that may enhance the specification. As an example, an optional TVL has been included to allow for the measurement of both loss and delay or jitter with a single test. The implementation does not cause any interoperability because the optional TVL is ignored by equipment that does not support this. In mixed vendor environments loss measurement continues to be tracked but delay and jitter can only report round trip times. It is important to point out that the round trip times in this mixed vendor environments include the remote nodes processing time because only two time stamps will be included in the packet. In an environment where both nodes support the optional TLV to include time stamps unidirectional and round trip times is reported. Since all four time stamps are included in the packet the round trip time in this case does not include remote node processing time. Of course, those operators that wish to run delay measurement and loss measurement at different frequencies are free to run both ETH-SL and ETH-DM functions. ETH-SL is not replacing ETH-DM. Service Assurance is only briefly discussed here to provide some background on the basic functionality. To know more about SAA functions see [Service Assurance Agent Overview on page 176](#).

The ETH-SL packet format contains a test-id that is internally generated and not configurable. The test-id is visible for the on demand test in the display summary. It is possible for a remote node processing the SLM frames receives overlapping test-ids as a result of multiple MEPs measuring loss between the same remote MEP. For this reason, the uniqueness of the test is based on remote MEP-ID, test-id and Source MAC of the packet.

ETH-SL is applicable to up and down MEPs and as per the recommendation transparent to MIPs. There is no coordination between various fault conditions that could impact loss measurement. This is also true for conditions where MEPs are placed in shutdown state as a result of linkage to a redundancy scheme like MC-LAG. Loss measurement is based on the ETH-SL and not coordinated across different functional aspects on the network element. ETH-SL is supported on service based MEPs.

It is possible that two MEPs may be configured with the same MAC on different remote nodes. This causes various issues in the FDB for multipoint services and is considered a misconfiguration for most services. It is possible to have a valid configuration where multiple MEPs on the same remote node have the same MAC. In fact, this is likely to happen. In this release, only the first responder is used to measure packet loss. The second responder is dropped. Since the same MAC for multiple MEPs is only truly valid on the same remote node this should be an acceptable approach.

There is no way for the responding node to understand when a test is completed. For this reason a configurable “inactivity-timer” determines the length of time a test is valid. The timer will maintain an active test as long as it is receiving packets for that specific test, defined by the test-id, remote MEP Id and source MAC. When there is a gap between the packets that exceeds the inactivity-timer the responding node responds with a sequence number of one regardless of what the sequence number was the instantiating node sent. This means the remote MEP accepts that the

previous test has expired and these probes are part of a new test. The default for the inactivity timer is 100 second and has a range of 10 to 100 seconds.

The responding node is limited to a fixed number of SLM tests per platform. Any test that attempts to involve a node that is already actively processing more than the system limit of the SLM tests shows up as “out loss” or “unacknowledged” packets on the node that instantiated the test because the packets are silently discarded at the responder. It is important for the operator to understand this is silent and no log entries or alarms is raised. It is also important to keep in mind that these packets are ETH-CFM based and the different platforms stated receive rate for ETH-CFM must not be exceeded. ETH-SL provides a mechanism for operators to pro-actively trend packet loss for service based MEPs.

## Configuration Example

The following illustration, shows the configuration required for proactive SLM test using SAA.




Figure 19: SLM Example

The output from the MIB is shown below as an example of an on-demand test. Node1 is tested for this example. The SAA configuration does not include the accounting policy required to collect the statistics before they are overwritten. NODE2 does not have an SAA configuration. NODE2 includes the configuration to build the MEP in the VPLS service context.

```
config>eth-cfm# info

 domain 3 format none level 3
 association 1 format icc-based name "03-0000000100"
 bridge-identifier 100
 exit
 ccm-interval 1
 remote-mepid 101
 exit
 exit
```

## Configuration Example

```

*A:7210SAS>config>service>vpls# info

 stp
 shutdown
 exit
 sap 1/1/3:100.100 create
 exit
 sap lag-1:100.100 create
 eth-cfm
 mep 100 domain 3 association 1 direction down
 ccm-enable
 mac-address d0:0d:1e:00:01:00
 no shutdown
 exit
 exit
 exit
no shutdown

*A:7210SAS>config>service>vpls
*A:7210SAS>config>saa# info detail

 test "SLM" owner "TiMOS CLI"
 no description
 type
 eth-cfm-two-way-slm 00:01:22:22:33:34 mep 1 domain 1 association 1 size 0
 fc "nc" count 100 timeout 1 interval 1
 exit
 trap-gen
 no probe-fail-enable
 probe-fail-threshold 1
 no test-completion-enable
 no test-fail-enable
 test-fail-threshold 1
 exit
 continuous
 no shutdown
 exit

*A:7210SAS>config>saa#
```

The following sample output is meant to demonstrate the different loss conditions that an operator may see. The total number of attempts is "100" is because the final probe in the test was not acknowledged.

```
*A:7210SAS# show saa SLM42
```

```
=====
SAA Test Information
=====
Test name : SLM42
Owner name : TiMOS CLI
Description : N/A
Accounting policy : None
Continuous : Yes
Administrative status : Enabled
```

```

Test type : eth-cfm-two-way-slm 00:25:ba:02:a6:50 mep 4
 : domain 1 association 1 fc "h1" count 100
 : timeout 1 interval 1
Trap generation : None
Test runs since last clear : 117
Number of failed test runs : 1
Last test result : Success

```

```

Threshold
Type Direction Threshold Value Last Event Run #

Jitter-in Rising None None Never None
 Falling None None Never None
Jitter-out Rising None None Never None
 Falling None None Never None
Jitter-rt Rising None None Never None
 Falling None None Never None
Latency-in Rising None None Never None
 Falling None None Never None
Latency-out Rising None None Never None
 Falling None None Never None
Latency-rt Rising None None Never None
 Falling None None Never None
Loss-in Rising None None Never None
 Falling None None Never None
Loss-out Rising None None Never None
 Falling None None Never None
Loss-rt Rising None None Never None
 Falling None None Never None

```

=====

```

Test Run: 116
Total number of attempts: 100
Number of requests that failed to be sent out: 0
Number of responses that were received: 100
Number of requests that did not receive any response: 0
Total number of failures: 0, Percentage: 0

```

(in ms)	Min	Max	Average	Jitter
Outbound :	8.07	8.18	8.10	0.014
Inbound :	-7.84	-5.46	-7.77	0.016
Roundtrip :	0.245	2.65	0.334	0.025

```

Per test packet:
Sequence Outbound Inbound RoundTrip Result

1 8.12 -7.82 0.306 Response Received
2 8.09 -7.81 0.272 Response Received
3 8.08 -7.81 0.266 Response Received
4 8.09 -7.82 0.270 Response Received
5 8.10 -7.82 0.286 Response Received
6 8.09 -7.81 0.275 Response Received
7 8.09 -7.81 0.271 Response Received
8 8.09 -7.82 0.277 Response Received
9 8.11 -7.81 0.293 Response Received
10 8.10 -7.82 0.280 Response Received
11 8.11 -7.82 0.293 Response Received
12 8.10 -7.82 0.287 Response Received
13 8.10 -7.82 0.286 Response Received
14 8.09 -7.82 0.276 Response Received

```

## Configuration Example

```
15 8.10 -7.82 0.284 Response Received
16 8.09 -7.82 0.271 Response Received
17 8.11 -7.81 0.292 Response Received
```

```
=====
The following is an example of an on demand tests that and the associated output. Only
single test runs are stored and can be viewed after the fact.
#oam eth-cfm two-way-slm-test 00:25:ba:04:39:0c mep 4 domain 1 association 1 send-count
10 interval 1 timeout 1
Sending 10 packets to 00:25:ba:04:39:0c from MEP 4/1/1 (Test-id: 143)
Sent 10 packets, 10 packets received from MEP ID 3, (Test-id: 143)
(0 out-loss, 0 in-loss, 0 unacknowledged)
```

```
*A:7210SAS>show# eth-cfm mep 4 domain 1 association 1 two-way-slm-test
```

```
=====
Eth CFM Two-way SLM Test Result Table (Test-id: 143)
=====
```

Peer Mac Addr	Remote MEP	Count	In Loss	Out Loss	Unack
00:25:ba:04:39:0c	3	10	0	0	0

```
=====
*A:7210SAS>show#
```

## ETH-CFM QoS Considerations

UP MEPs and Down MEPs have been aligned as of this release to better emulate service data. When an UP MEP or DOWN MEP is the source of the ETH-CFM PDU the priority value configured, as part of the configuration of the MEP or specific test, will be treated as the Forwarding Class (FC) by the egress QoS policy. If there is no egress QoS policy the priority value will be mapped to the CoS values in the frame. However, egress QoS Policy may overwrite this original value. The Service Assurance Agent (SAA) uses [fc {fc-name}] to accomplish similar functionality.

UP MEPs and DOWN MEPs terminating an ETH-CFM PDU will use the received FC as the return priority for the appropriate response, again feeding into the egress QoS policy as the FC.

ETH-CFM PDUs received on the MPLS-SDP bindings will now properly pass the EXP bit values to the ETH-CFM application to be used in the response.

These are default behavioral changes without CLI options.

---

## Configuration guidelines

- On 7210 SAS platforms Ethernet Linktrace Response (ETH-LTR) is always sent out with priority 7.
- 7210 SAS platforms, send out all CFM packets as in-profile. Currently, there is no mechanism in the SAA tools to specify the profile of the packet.
- On 7210 SAS-M only UP MEPs behavior has been modified. Down MEPs are not supported.
- Other than as noted above, this behavior is applicable to all 7210 platforms where the feature supported has been added in both access-uplink mode and network mode of operation. Please see the Release Notes to know which platform and MEPs/MIPs support this feature.

## OAM Mapping

---

OAM mapping is a mechanism that enables a way of deploying OAM end-to-end in a network where different OAM tools are used in different segments. For instance, an Epipe service could span across the network using Ethernet access (CFM used for OAM), pseudowire (T-LDP status signaling used for OAM), and Ethernet access (CFM used for OAM).

In the 7210 SAS implementation, the Service Manager (SMGR) is used as the central point of OAM mapping. It receives and processes the events from different OAM components, then decides the actions to take, including triggering OAM events to remote peers.

Fault propagation for CFM is by default disabled at the MEP level to maintain backward compatibility. When required, it can be explicitly enabled by configuration.

Fault propagation for a MEP can only be enabled when the MA is comprised of no more than two MEPs (point-to-point).

---

## CFM Connectivity Fault Conditions

CFM MEP declares a connectivity fault when its defect flag is equal to or higher than its configured lowest defect priority. The defect can be any of the following depending on configuration:

- DefRDICCM
- DefMACstatus
- DefRemoteCCM
- DefErrorCCM
- DefXconCCM

The following additional fault condition applies to Y.1731 MEPs:

- Reception of AIS for the local MEP level

Setting the lowest defect priority to allDef may cause problems when fault propagation is enabled in the MEP. In this scenario, when MEP A sends CCM to MEP B with interface status down, MEP B will respond with a CCM with RDI set. If MEP A is configured to accept RDI as a fault, then it gets into a dead lock state, where both MEPs will declare fault and never be able to recover. The default lowest defect priority is DefMACstatus, which will not be a problem when interface status TLV is used. It is also very important that different Ethernet OAM strategies should not overlap the span of each other. In some cases, independent functions attempting to perform their normal fault



handling can negatively impact the other. This interaction can lead to fault propagation in the direction toward the original fault, a false positive, or worse, a deadlock condition that may require the operator to modify the configuration to escape the condition. For example, overlapping Link Loss Forwarding (LLF) and ETH-CFM fault propagation could cause these issues.

For the DefRemoteCCM fault, it is raised when any remote MEP is down. So whenever a remote MEP fails and fault propagation is enabled, a fault is propagated to SMGR.

---

## CFM Fault Propagation Methods

When CFM is the OAM module at the other end, it is required to use any of the following methods (depending on local configuration) to notify the remote peer:

- Generating AIS for certain MEP levels
- Sending CCM with interface status TLV “down”
- Stopping CCM transmission

---

**NOTE:** 7210 platforms expect that the fault notified using interface status TLV, is cleared explicitly by the remote MEP when the fault is no longer present on the remote node. On 7210 SAS, use of CCM with interface status TLV Down is not recommended to be configured with a Down MEP, unless it is known that the remote MEP clears the fault explicitly.

User can configure UP MEPs to use Interface Status TLV with fault propagation. Special considerations apply only to Down MEPs.

---

When a fault is propagated by the service manager, if AIS is enabled on the SAP/SDP-binding, then AIS messages are generated for all the MEPs configured on the SAP/SDP-binding using the configured levels.

Note that the existing AIS procedure still applies even when fault propagation is disabled for the service or the MEP. For example, when a MEP loses connectivity to a configured remote MEP, it generates AIS if it is enabled. The new procedure that is defined in this document introduces a new fault condition for AIS generation, fault propagated from SMGR, that is used when fault propagation is enabled for the service and the MEP.

The transmission of CCM with interface status TLV must be done instantly without waiting for the next CCM transmit interval. This rule applies to CFM fault notification for all services.

Notifications from SMGR to the CFM MEPs for fault propagation should include a direction for the propagation (up or down: up means in the direction of coming into the SAP/SDP-binding;

## Configuration guidelines

down means in the direction of going out of the SAP/SDP-binding), so that the MEP knows what method to use. For instance, an up fault propagation notification to a down MEP will trigger an AIS, while a down fault propagation to the same MEP can trigger a CCM with interface TLV with status down.

For a specific SAP/SDP-binding, CFM and SMGR can only propagate one single fault to each other for each direction (up or down).

When there are multiple MEPs (at different levels) on a single SAP/SDP-binding, the fault reported from CFM to SMGR will be the logical OR of results from all MEPs. Basically, the first fault from any MEP will be reported, and the fault will not be cleared as long as there is a fault in any local MEP on the SAP/SDP-binding.

## Epipe Services

Down and up MEPs are supported for Epipe services as well as fault propagation. When there are both up and down MEPs configured in the same SAP/SDP-binding and both MEPs have fault propagation enabled, a fault detected by one of them will be propagated to the other, which in turn will propagate fault in its own direction.

---

### CFM Detected Fault

When a MEP detects a fault and fault propagation is enabled for the MEP, CFM needs to communicate the fault to SMGR, so SMGR will mark the SAP/SDP-binding faulty but still oper-up. CFM traffic can still be transmitted to or received from the SAP/SDP-binding to ensure when the fault is cleared, the SAP will go back to normal operational state. Since the operational status of the SAP/SDP-binding is not affected by the fault, no fault handling is performed. For example, applications relying on the operational status are not affected.

If the MEP is an up MEP, the fault is propagated to the OAM components on the same SAP/SDP-binding; if the MEP is a down MEP, the fault is propagated to the OAM components on the mate SAP/SDP-binding at the other side of the service.

### SAP/SDP-Binding Failure (Including Pseudowire Status)

When a SAP/SDP-binding becomes faulty (oper-down, admin-down, or pseudowire status faulty), SMGR needs to propagate the fault to up MEP(s) on the same SAP/SDP-bindings about the fault, as well as to OAM components (such as down MEPs) on the mate SAP/SDP-binding.

---

### Service Down

This section describes procedures for the scenario where an Epipe service is down due to the following:

- Service is administratively shutdown. When service is administratively shutdown, the fault is propagated to the SAP/SDP-bindings in the service.
- If the Epipe service is used as a PBB tunnel into a B-VPLS, the Epipe service is also considered operationally down when the B-VPLS service is administratively shutdown or operationally down. If this is the case, fault is propagated to the Epipe SAP.

In addition, one or more SAPs/SDP-bindings in the B-VPLS can be configured to propagate fault to this Epipe (see fault-propagation-bmac below). If the B-VPLS is operationally up but all of these entities have detected fault or are down, the fault is propagated to this Epipe's SAP.

## Interaction with Pseudowire Redundancy

When a fault occurs on the SAP side, the pseudowire status bit is set for both active and standby pseudowires. When only one of the pseudowire is faulty, SMGR does not notify CFM. The notification occurs only when both pseudowire becomes faulty. The SMGR propagates the fault to CFM.

Since there is no fault handling in the pipe service, any CFM fault detected on an SDP binding is not used in the pseudowire redundancy's algorithm to choose the most suitable SDP binding to transmit on.

---

## LLF and CFM Fault Propagation

LLF and CFM fault propagation are mutually exclusive. CLI protection is in place to prevent enabling both LLF and CFM fault propagation in the same service, on the same node and at the same time. However, there are still instances where irresolvable fault loops can occur when the two schemes are deployed within the same service on different nodes. This is not preventable by the CLI. At no time should these two fault propagation schemes be enabled within the same service.

## 802.3ah EFM OAM Mapping and Interaction with Service Manager

802.3ah EFM OAM declares a link fault when any of the following occurs:

- Loss of OAMPDU for a certain period of time
- Receiving OAMPDU with link fault flags from the peer

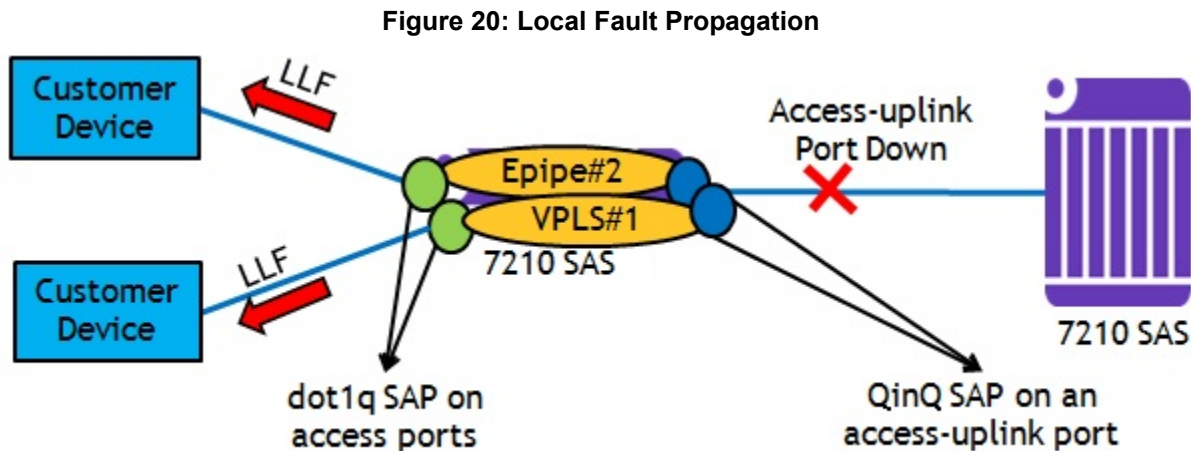
When 802.3ah EFM OAM declares a fault, the port goes into operation state down. The SMGR communicates the fault to CFM MEPs in the service. OAM fault propagation in the opposite direction (SMGR to EFM OAM) is not supported.

---

## Fault Propagation to access dot1q/QinQ ports in Access-uplink mode

The Fault on the access uplink port brings down all access ports with services independent of the encap type of the access port (NULL or dot1q or QinQ), that is, support Link Loss Forwarding (LLF). The Fault propagated from the access-uplink port to access ports is based on configuration. The Fault is propagated only in a single direction, that is access-uplink port to access port.

The Fault on access-uplink port is detected using LoS and EFM-OAM.



## Configuring Fault propagation:

To allow faults to be detected on access-uplink ports and be propagated to all the other interested access ports irrespective of their encapsulation, the concept of oper-group is used. In 7210, ports can be associated with oper-groups. Following bullets describe the use of oper-group functionality for fault-detection on a port and monitor-oper-group to track the status of the oper-group and propagate fault based on the operational state of the oper-group.

- Create an oper-group (For example: “uplink-to-7210”)
- Configure access-uplink port whose operational state needs to be tracked (For example: 1/1/20) and associate it with the oper-group created above (that is, uplink-to-7210)
- Configure dot1q access ports whose operational state needs to be driven by operational state of access-uplink port (For example: 1/1/1 and 1/1/5), as the monitor-oper-group
- To detect fault on access-uplink port and change the operational state, either Loss of Signal or EFM OAM can be used.
- When the operational state of access-uplink port changes from UP to DOWN changes the state of all access ports configured to monitor the group to “Down”. Similarly, a change in state from Down to UP, changes the operational state of the access port to UP. When the operational state of access port is brought down, the laser of the port is also shutdown. The command "*Hold timers*" are supported to avoid flapping of links.

## Configuration Example for fault propagation using oper-group

The following example displays an “oper-group” system configuration:

```
*A:7210SAS>config>system>oper-group# info detail

 hold-time
 group-down 0
 group-up 4
 exit

*A:7210SAS>config>system>oper-group#
```

**Note:** For more details on the CLI, see the “7210 SAS System Basics User Guide”.

---

## Service Assurance Agent Overview

In the last few years, service delivery to customers has drastically changed. Services such as VPLS and VPRN are offered. The introduction of Broadband Service Termination Architecture (BSTA) applications such as Voice over IP (VoIP), TV delivery, video and high speed Internet services force carriers to produce services where the health and quality of Service Level Agreement (SLA) commitments are verifiable to the customer and internally within the carrier.

SAA is a feature that monitors network operations using statistics such as jitter, latency, response time, and packet loss. The information can be used to troubleshoot network problems, problem prevention, and network topology planning.

The results are saved in SNMP tables are queried by either the CLI or a management system. Threshold monitors allow for both rising and falling threshold events to alert the provider if SLA performance statistics deviate from the required parameters.

SAA allows two-way timing for several applications. This provides the carrier and their customers with data to verify that the SLA agreements are being properly enforced.

---

## Traceroute Implementation

The 7210 SAS devices insert the timestamp in software (by control CPU).

When interpreting these timestamps care must be taken that some nodes are not capable of providing timestamps, as such timestamps must be associated with the same IP-address that is being returned to the originator to indicate what hop is being measured.

---

## NTP

Because NTP precision can vary (+/- 1.5ms between nodes even under best case conditions), SAA one-way latency measurements might display negative values, especially when testing network segments with very low latencies. The one-way time measurement relies on the accuracy of NTP between the sending and responding nodes.



## Writing SAA Results to Accounting Files

SAA statistics enables writing statistics to an accounting file. When results are calculated an accounting record is generated.

In order to write the SAA results to an accounting file in a compressed XML format at the termination of every test, the results must be collected, and, in addition to creating the entry in the appropriate MIB table for this SAA test, a record must be generated in the appropriate accounting file.

---

## Accounting File Management

Because the SAA accounting files have a similar role to existing accounting files that are used for billing purposes, existing file management information is leveraged for these accounting (billing) files.

---

## Assigning SAA to an Accounting File ID

Once an accounting file has been created, accounting information can be specified and will be collected by the **config>log>acct-policy>** to file *log-file-id* context.

## Continuous Testing

When you configure a test, use the **config>saa>test>continuous** command to make the test run continuously. Use the **no continuous** command to disable continuous testing and **shutdown** to disable the test completely. Once you have configured a test as continuous, you cannot start or stop it by using the **saa test-name [owner test-owner] {start | stop} [no-accounting]** command.

## Configuring SAA Test Parameters

The following example displays an SAA configuration:

```
*A:Dut-A>config>saa# info

....
 test "Dut-A:1413:1501" owner "TiMOS"
 description "Dut-A:1413:1501"
 type
 vccv-ping 1413:1501 fc "nc" timeout 10 size 200 count 2
 exit
 loss-event rising-threshold 2
```

## Configuring SAA Test Parameters

```
latency-event rising-threshold 100
no jitter-event
no shutdown
exit
```

```
.....
```

```

*A:Dut-A#
```


## Y.1564 Testhead OAM tool

**Note:** Port loopback with mac-swap and Y.1564 testhead is supported only for Epipe and VPLS services.

ITU-T Y.1564 defines the out-of-service test methodology to be used and parameters to be measured to test service SLA conformance during service turn up. It primarily defines 2 test phases. The first test phase defines service configuration test, which consists of validating whether the service is configured properly. As part of this test the throughput, Frame Delay, Frame Delay Variation (FDV), and Frame Loss Ratio (FLR) is measured for each service. This test is typically run for a short duration. The second test phase consists of validating the quality of services delivered to the end customer and is referred to as the service performance test. These tests are typically run for a longer duration and all traffic is generated up to the configured CIR for all the services simultaneously and the service performance parameters are measured for each the service.

The 7210 SAS supports service configuration test for user configured rate and measurement of delay, delay variation and frame loss ratio with the testhead OAM tool. 7210 testhead OAM tool supports bi-directional measurement and it can generate test traffic for only one service at a given time. It can validate if the user specified rate is available and compute the delay, delay variation and frame loss ratio for the service under test at the specified rate. It is capable of generating traffic upto 1G rate. On some 7210 SAS devices, the user needs to configure the resources of the front-panel port for use with this feature and some other 7210 SAS platforms resources needed for this feature is automatically allocated by software from the internal ports. For more information, see [Configuration Guidelines on page 186](#), below to which 7210 SAS platforms need user configuration and on which 7210 platforms software allocates it automatically.

[Figure 21](#) illustrates the remote loopback required and the flow of the frame through the network generated by the testhead tool.



**Figure 21: 7210 acting as traffic generator and traffic analyzer**

The tool allows the user to specify the frame payload header parameters independent of the test SAP configuration parameters to allow the user flexibility to test for different possible frame header encapsulations. This allows user to specify the appropriate VLAN tags, ethertype, and dotIp's, independent of the SAP configuration like with actual service testing. In other words, the software does not use the parameters (For example: SAP ID, Source MAC, and Destination MAC) during the invocation of the testhead tool to build the test frames. Instead it uses the parameters specified using the frame-payload CLI command tree. The software does not verify that the parameters specified match the service configuration used for testing, for example, software does not match if the VLAN tags specified matches the SAP tags, the ethertype specified matches the user configured port ethertype, and so on. It is expected that the user configures the frame-payload appropriately so that the traffic matches the SAP configuration.

7210 SAS supports Y.1564 testhead for performing CIR or PIR tests in color-aware mode. With this functionality, users can perform service turn-up tests to validate the performance characteristics (delay, jitter, and loss) for committed rate (CIR) and excess rate above CIR (that is, PIR rate). The testhead OAM tool uses the in-profile packet marking value and out-of-profile packet marking value, to differentiate between committed traffic and PIR traffic in excess of CIR traffic. Traffic within CIR (that is, committed traffic) is expected to be treated as in-profile traffic in the network and traffic in excess of CIR (that is, PIR traffic) is expected to be treated as out-of-profile traffic in the network, allowing the network to prioritize committed traffic over PIR traffic. The testhead OAM tool allows the user to configure individual thresholds for green or in-profile packets and out-of-profile or yellow packets. It is used by the testhead OAM tool to compare the

measured value for green or in-profile packets and out-of-profile or yellow packets against the configured thresholds and report success or failure.

The following functionality is supported by the testhead OAM tool:

- Supports configuration of only access SAPs as the test measurement point.
- Supports all port encapsulation supported on all svc-sap-types.
- Supported for only VPLS and Epipe service.
- Supports two-way measurement of service performance metrics. The tests must measure throughput, frame delay, frame delay variation, and frame loss ratio.
- Fortwo-way measurement of the service performance metrics, such as frame delay and frame delay variation, test frames (also called as marker packets) are injected at a low rate at periodic intervals. Frame delay and Frame delay variation is computed for these frames. Hardware based timestamps is used for delay computation.
- 7210 SAS supports configuration of PIR rate and provides an option to measure the performance metrics for the frames. The testhead OAM tool, generates traffic upto the specified rate and measures service performance metrics such as, delay, jitter, loss for in-profile, and out-of-profile traffic.
- Testhead tool can generate traffic up to about 1G rate. CIR and PIR rate can be specified by the user and is rounded off the nearest rate the hardware supports by using the adaptation rule configured by the user.
- Allows the user to specify the different frame-sizes from 64 bytes - 9212 bytes.
- User can configure the following frame payload types- L2 payload, IP payload, and IP/TCP/UDP payload. Testhead tool will use the configured values for the IP header fields and TCP header fields based on the payload type configured. User is provided with an option to specify the data pattern to used in the payload field of the frame/packet.
- Allows the user to configure the duration of the test up to a maximum of 24 hours, 60 minutes, and 60 seconds. The test performance measurements by are done after the specified rate is achieved. At any time user can probe the system to know the current status and progress of the test.
- Supports configuration of the Forwarding Class (FC). It is expected that user will define consistent QoS classification policies to map the packet header fields to the FC specified on the test SAP ingress on the local node, in the network on the nodes through which the service transits, and on the SAP ingress in the remote node.
- Allows the user to configure a test-profile, also known as, a policy template that defines the test configuration parameters. User can start a test using a pre-configured test policy for a specific SAP and service. The test profile allows the user to configure the acceptance criteria. The acceptance criteria allows user to configure the thresholds that indicates the acceptable range for the service performance metrics. An event is generated if the test results exceed the configured thresholds. For more information, see the CLI section below.

At the end of the test, the measured values for FD, FDV, and FLR are compared against the configured thresholds to determine the PASS or FAIL criteria and to generate a trap to the management station. If the acceptance criteria is not configured, the test result is declared to be PASS, if the throughput is achieved and frame-loss is 0 (zero).

- ITU-T Y.1564 specifies different test procedures as follows. CIR and PIR configuration tests are supported by the testhead tool.
  - CIR and PIR configuration test (color-aware and non-color aware).
  - Traffic policing test (color-aware and non-color aware) is supported. Traffic policing tests can be executed by the user by specifying a PIR to be 125% of the desired PIR. Traffic policing test can be executed in either color-aware mode or color-blind (non-color-aware) mode.
- ITU-T Y.1564 specifies separate test methodology for color-aware and non-color-aware tests. The standard requires a single test to provide the capability to generate both green-color/in-profile traffic for rates within CIR and yellow-color or out-of-profile traffic for rates above CIR and within EIR. The 7210 SAS testhead marks test packets appropriately when generating the traffic, as SAP ingress does not support color-aware metering, it is not possible to support EIR color-aware, and traffic policing color-aware tests end-to-end in a network (that is, from test SAP to test SAP). Instead, It is possible to use the tests to measure the performance parameters from the other endpoint (example Access-uplink SAP) in the service, through the network, to the remote test SAP, and back again to the local test SAP.
- The 7210 SAS Y.1564 testhead is applicable only for VPLS and Epipe services.

## Pre-requisites for using the Testhead Tool

---

### Generic pre-requisites for use of testhead tool (applicable for all 7210 SAS platforms):

- It is expected that the user will configure the appropriate ACL and QoS policies to ensure that the testhead traffic is processed as desired by the local and remote node/SAP. In particular, QoS policies in use must ensure that the rate in use for the SAP ingress meters exceed or are equal to the user configured rate for testhead tests and the classification policies map the testhead packets to the appropriate FCs/queues (the FC classification must match the FC specified in the CLI command testhead-test) using the packet header fields configured in the frame-payload. Similarly, ACL policies must ensure that testhead traffic is not blocked.
- The testhead OAM tool does not check the state of the service or the SAPs on the local endpoint before initiating the tests. The operator must ensure that the service and SAPs used for the test are UP before the tests are started. If they are not, the testhead tool will report a failure.
- The port configuration of the ports used for validation (For example: access port on which the test SAP is configured and access-uplink port) must not be modified after the testhead tool is invoked. Any modifications can be made only when the testhead tool is not running.
- Testhead tool can be used to test only unicast traffic flows. It must not be used to test BUM traffic flows.
- Only out-of-service performance metrics can be measured using testhead OAM tool. For in-service performance metrics, user has the option to use SAA based Y.1731/CFM tools.

This section describes some pre-requisites to use the testhead tool .

- The configuration guidelines and pre-requisites that are to be followed when the port loopback with mac-swap feature is used standalone, applies to its use along with testhead tool. For more information, see the description in the “7210 SAS-M, T, X, R6, R12, Mxp, Sx Interfaces User Guide”.
- User must configure resources for ACL MAC criteria in ingress-internal-tcam using the command `config>system> resource-profile> ingress-internal-tcam> acl-sap-ingress> mac-match-enable`. Additionally they must allocate resources to egress ACL MAC or IPv4 or IPv6 64-bit criteria (using the command `config>system> resource-profile> egress-internal-tcam> acl-sap-egress> mac-ipv4-match-enable` or `mac-ipv6-64bit-enable` or `mac-ipv4-match-enable`). Testhead tool uses resources from these resource pools. If no resources are allocated to these pools or no resources are available for use in these pools, then testhead fails to function. Testhead needs a minimum of about 6 entries from the

ingress-internal-tcam pool and 2 entries from the egress-internal-tcam pool. If user allocates resources to egress ACLs IPv6 128-bit match criteria (using the command `config> system> resource-profile> egress-internal-tcam> acl-sap-egress> ipv6-128bit-match-enable`), then testhead fails to function.

- For both Epipe and VPLS service, the test can be used to perform only a point-to-point test between the given source and destination MAC address. Port loopback mac-swap functionality must be used for both Epipe and VPLS services. The configured source and destination MAC address is associated with the two SAPs configured in the service and used as the two endpoints. In other words, the user configured source MAC and destination MAC addresses are used by the testhead tool on the local node to identify the packets as belonging to testhead application and are processed appropriately at the local end and at the remote end these packets are processed by the port loopback with mac-swap application.
- Configure the MACs (source and destination) statically for VPLS service.
- Port loopback must be in use on both the endpoints (that is, the local node, the port on which the test SAP is configured and the remote node, the port on which the remote SAP is configured for both Epipe and VPLS services. Port loopback with mac-swap must be setup by the user on both the local end and the remote end before invoking the testhead tool. These must match appropriately for traffic to flow, else there will be no traffic flow and the testhead tool reports a failure at the end of the completion of the test run.
- Additionally, port loopback with mac-swap must be used at both the ends and if any services/SAPs are configured on the test port, they need to be shutdown to avoid packets being dropped on the non-test SAP. The frames generated by the testhead tool will egress the access SAP and ingress back on the same port, using the resources of the 2 internal loopback ports (one for testhead and another for mac-swap functionality), before being sent out to the network side (typically an access-uplink SAP) to the remote end". At the remote end, it is expected that the frames will egress the SAP under test and ingress back in again through the same port, going through another loopback (with mac-swap) before being sent back to the local node where the testhead application is running.
- The FC specified is used to determine the queue to enqueue the marker packets generated by testhead application on the egress of the test SAP on the local node.
- The use of port loopback is service affecting. It affects all the services configured on the port. It is not recommended to use a SAP, if the port on which they are configured, is used to transport the service packets towards the core. As, a port loopback is required for the testhead to function correctly, doing so might result in loss of connectivity to the node when in-band management is in use. Additionally, all services being transported to the core will be affected.
- It also affects service being delivered on that SAP. Only out-of-service performance metrics can be measured using testhead OAM tool. For in-service performance metrics, user has the option to use SAA based Y.1731/CFM tools.



- Testhead tool uses marker packets with special header values. The QoS policies and ACL policies need to ensure that same treatment as accorded to testhead traffic is given to marker packets. Marker packets are IPv4 packet with IP option set and IP protocol set to 252. It uses the src and dst MAC addresses, Dot1p, IP ToS, IP DSCP, IP TTL, IP source address and destination address as configured in the frame-payload. It does not use the IP protocol and TCP/UDP port numbers from the frame-payload configured. If the payload-type is “l2”, IP addresses are set to 0.0.0.0, IP TTL is set to 0, IP TOS is set to 0 and DSCP is set to be, if these values are not explicitly configured in the frame-payload. Ether type configured in the frame-payload is not used for marker packets, it is always set to ether type = 0x0800 (ether type for IPv4) as marker packets are IPv4 packets. QoS policies applied in the network needs to be configured such that the classification for marker packets is similar to service packets. An easy way to do this is by using the header fields that are common across marker packets and service packets, such as MAC (src and dst) addresses, VLAN ID, Dot1p, IPv4 (src and dst) addresses, IP DSCP, and IP ToS. Use of other fields which are different for marker packets and service packets is not recommended. ACL policies in the network must ensure that marker packets are not dropped.
  - The mac-swap loopback port, the testhead loopback port and the uplink port must not be modified after the testhead tool is invoked. Any modifications can be made only when the testhead tool is not running.
  - Link-level protocols (For example: LLDP, EFM, and other protocols) must not be enabled on the port on which the test SAP is configured. In general, no other traffic must be sent out of the test SAP when the testhead tool is running.
  - The frame payload must be configured such that number of tags match the number of SAP tags. For example: For 0.\* SAP, the frame payload must be untagged or priority tagged and it cannot contain another tag following the priority tag.
-

## Configuration Guidelines

This section provide the configuration guidelines for this feature (that is, testhead OAM tool). It is applicable to all the platforms discussed in this guide unless a specific platform is called out explicitly.

- SAPs configured on LAG cannot be configured for testing with testhead tool. Other than the test SAP, other service endpoints (For example: SAPs/SDP-Bindings) configured in the service can be over a LAG.
- The user needs to configure the resources of another port for use with the testhead OAM tool. The user can configure the resources of the internal virtual ports (if available) or configure the resources of a front-panel port for use with testhead OAM tool. Please read the details provided in the CLI command `config> system> loopback-no-svc-port` in the 7210 SAS Interfaces user guide to know if front-panel port resources are needed and use the command `show>system>internal-loopback-ports [detail]` to know if internal port resources are available in use by other applications. The port configured for testhead tool use cannot be shared with other applications that need the loopback port. The resources of the loopback port are used by the testhead tool for traffic generation.  
**NOTE:** On 7210 SAS-R6 and 7210 SAS-R12, the ports allocated for testhead OAM tool, the mac-swap OAM tool and the test SAP must be on the same line card and cannot be on different line cards.
- Port loopback with mac-swap is used at both ends and all services on the port on which the test SAP is configured and SAPs in the VPLS service, other than the test SAP should be shutdown or should not receive any traffic.
- The configured CIR/PIR rate is rounded off to the nearest available hardware rates. User is provided with an option to select the adaptation rule to use (similar to support available for QoS policies).
- 7210 SAS supports all port speeds (that is, up to 1G rate). User needs to configure the appropriate loopback port to achieve a desired rate. For example, user must dedicate the resources of a 1G port, if intended to test rates to go up to 1G.
- ITU-T Y.1564 recommends to provide an option to configure the CIR step-size and the step-duration for the service configuration tests. This is not supported directly in 7210 SAS. It can be achieved by SAM or a third-party NMS system or an application with configuration of the desired rate and duration to correspond to the CIR step-size and step duration and repeating the test a second time, with a different value of the rate (that is, CIR step size) and duration (that is, step duration) and so on.
- Testhead waits for about 5 seconds at the end of the configured test duration before collecting statistics. This allows for all in-flight packets to be received by the node and accounted for in the test measurements. User cannot start another test during this period.
- When using testhead to test bandwidth available between SAPs configured in a VPLS service, operators must ensure that no other SAPs in the VPLS service are exchanging any

traffic, particularly BUM traffic and unicast traffic destined to either the local test SAP or the remote SAP. BUM traffic eats into the network resources which is also used by testhead traffic.

- It is possible that test packets (both data and/or marker packets) remain in the loop created for testing when the tests are killed. This is highly probably when using QoS policies with very less shaper rates resulting in high latency for packets flowing through the network loop. User must remove the loop at both ends once the test is complete or when the test is stopped and wait for a suitable time before starting the next test for the same service, to ensure that packets drain out of the network for that service. If this is not done, then the subsequent tests might process and account these stale packets, resulting in incorrect results. Software cannot detect stale packets in the loop as it does not associate or check each and every packet with a test session
- Traffic received from the remote node and looped back into the test port (where the test SAP is configured) on the local end (that is, the end where the testhead tool is invoked) is dropped by hardware after processing (and is not sent back to the remote end). The SAP ingress QoS policies and SAP ingress filter policies must match the packet header fields specified by the user in the testhead profile, except that the source/destination MAC addresses are swapped.
- Latency is not be computed if marker packets are not received by the local node where the test is generated and printed as 0 (zero), in such cases. If jitter = 0 and latency > 0, it means that jitter calculated is less than the precision used for measurement. There is also a small chance that jitter was not actually calculated, that is, only one value of latency has been computed. This typically indicates a network issue, rather than a testhead issue.
- When the throughput is not met, FLR cannot be calculated. If the measured throughput is approximately +/-10% of the user configured rate, FLR value is displayed; else software prints “Not Applicable”. The percentage of variance of measured bandwidth depends on the packet size in use and the configured rate.
- User must not use the CLI command to clear statistics of the test SAP port, testhead loopback port and MAC swap loopback port when the testhead tool is running. The port statistics are used by the tool to determine the Tx/Rx frame count.
- Testhead tool generates traffic at a rate slightly above the CIR. The additional bandwidth is attributable to the marker packets used for latency measurements. This is not expected to affect the latency measurement or the test results in a significant way.
- If the operational throughput is 1kbps and it is achieved in the test loop, the throughput computed could still be printed as 0 if it is < 1Kbps (0.99 kbps, for example). Under such cases, if FLR is PASS, the tool indicates that the throughput has been achieved.
- The testhead tool displays a failure result if the received count of frames is less than the injected count of frames, even though the FLR might be displayed as 0. This happens due to truncation of FLR results to 6 decimal places and can happen when the loss is very less.
- As the rate approaches 1Gbps or the maximum bandwidth achievable in the loop, user needs to account for the marker packet rate and the meter behavior while configuring the

CIR rate. In other words, if the user wants to test 1Gbps for 512 bytes frame size, then they will need to configure about 962396Kbps, instead of 962406Kbps, the maximum rate that can be achieved for this frame-size. In general, they would need to configure about 98%-99% (based on packet size) of the maximum possible rate to account for marker packets when they need to test at rates which are closer to bandwidth available in the network. The reason for this is that at the maximum rate, injection of marker packets by CPU will result in drops of either the injected data traffic or the marker packets themselves, as the net rate exceeds the capacity. These drops cause the testhead to always report a failure, unless the rate is marginally reduced.

- Testhead works with L2 rate, that is, the rate after subtracting the L1 overhead. The L1 overhead is due to IFG and Preamble added to every Ethernet frame and is typically about 20 bytes (IFG = 12 bytes and Preamble = 8 bytes). Depending on the frame size configured by the user, testhead tool computes the L2 rate and does not allow the user to configure a value greater than it. For 512 bytes Ethernet frame, L2 rate is 962406Kbps and L1 rate is 1Gbps.
- It is not expected that the operator will use the testhead tool to measure the throughput or other performance parameters of the network during the course of network event. The network events could be affecting the other SAP/SDP-Binding/PW configured in the service. Examples are transition of a SAP due to G8032 ring failure, transition of active/standby SDP-Binding/PW due to link or node failures.
- The 2-way delay (also known as “latency”) values measured by Testhead tool is more accurate than obtained using OAM tools, as the timestamps are generated in hardware.
- 7210 SAS does not support color-aware metering on access SAP ingress, therefore, any color-aware packets generated by the testhead is ignored on access SAP ingress. 7210 SAS service access port, access-uplink port, or network port can mark the packets appropriately on egress to allow the subsequent nodes in the network to differentiate the in-profile and out-of-profile packets and provide them with appropriate QoS treatment. 7210 SAS access-uplink ingress and network port ingress is capable of providing appropriate QoS treatment to in-profile and out-of-profile packets.
- The marker packets are sent over and above the configured CIR or PIR rate, the tool cannot determine the number of green packets injected and the number of yellow packets injected individually. Therefore, marker packets are not accounted in the injected or received green or in-profile and yellow or out-of-profile packet counts. They are only accounted for the Total Injected and the Total Received counts. So, the FLR metric accounts for marker packet loss (if any), while green or yellow FLR metric does not account for any marker packet loss.
- Marker packets are used to measure green or in-profile packets latency and jitter and the yellow or out-of-profile packets latency and jitter. These marker packets are identified as green or yellow based on the packet marking (Example: dot1p). The latency values can be different for green and yellow packets based on the treatment provided to the packets by the network QoS configuration.

- The following table provides details of SAP encapsulation that are supported for Testhead

**Table 15: SAP Encapsulations supported for testhead**

Epipse service configured with svc-sap-type	Test SAP Encapsulations
null-star	Null, :*, 0.*, Q.*
Any	Null, :0, :Q, :Q1.Q2
dot1q-preserve	:Q

- The following combination of 1G and 10G port can be used, as long as the rate validated is less than or equal to 1Gbps.
    - The test SAP is a 10G port, uplink is 1G port and other ports (that is, uplink, mac-swap and testhead) are 1G port.
    - The test SAP is a 10G port, uplink is a 10G port and other ports (that is, mac-swap and testhead) are 1G port.
    - The test SAP is a 1G port, uplink is a 10G port and other ports (that is, mac-swap and testhead) are 1G port.
-

## Configuring testhead tool parameters

The following example displays a port loopback mac-swap using the service and SAP:

```
configure> system> loopback-no-svc-port testhead <port-id>
*A:7210SAS>config>system# info

.....
resource-profile
 ingress-internal-tcam
 qos-sap-ingress-resource 5
 exit
 acl-sap-ingress 5
 exit
 exit
 egress-internal-tcam
 exit
exit
loopback-no-svc-port mac-swap 1/1/8 testhead 1/1/11
.....
```

The following example displays a port loopback with mac-swap on the remote end:

```
*A:7210SAS# configure system loopback-no-svc-port mac-swap 1/1/8
*A:7210SAS# configure system
*A:7210SAS>config>system# info

alarm-contact-input 1
 shutdown
exit
alarm-contact-input 2
 shutdown
exit
alarm-contact-input 3
 shutdown
exit
alarm-contact-input 4
 shutdown
exit
resource-profile
 ingress-internal-tcam
 qos-sap-ingress-resource 5
 exit
 acl-sap-ingress 5
 exit
 exit
 egress-internal-tcam
 exit
exit
loopback-no-svc-port mac-swap 1/1/8 testhead 1/1/11
```

.....

The following example displays a testhead profile:

```
*A:7210SAS# configure test-oam testhead-profile 1
*A:7210SAS>config>test-oam>testhd-prof# info

description "Testhead_Profile_1"
frame-size 512
rate cir 100 adaptation-rule max pir 200
dot1p in-profile 2 out-profile 4
frame-payload 1 payload-type tcp-ipv4 create
description "Frame_Payload_1"
dscp "af11"
dst-ip ipv4 2.2.2.2
dst-mac 00:00:00:00:00:02
src-mac 00:00:00:00:00:01
dst-port 50
src-port 40
ip-proto 6
ip-tos 8
ip-ttl 64
src-ip ipv4 1.1.1.1
exit
acceptance-criteria 1 create
jitter-rising-threshold 100
jitter-rising-threshold-in 100
jitter-rising-threshold-out 100
latency-rising-threshold 100
latency-rising-threshold-in 100
latency-rising-threshold-out 100
loss-rising-threshold 100
loss-rising-threshold-in 100
loss-rising-threshold-out 100
cir-threshold 1000
pir-threshold 2000
exit

*A:7210SAS>config>test-oam>testhd-prof#
```

The following command is used to execute the testhead profile:

```
*A:7210SAS# oam testhead testhead-profile 1 frame-payload 1 sap 1/1/2 test-me owner ownerme
color-aware enable
```

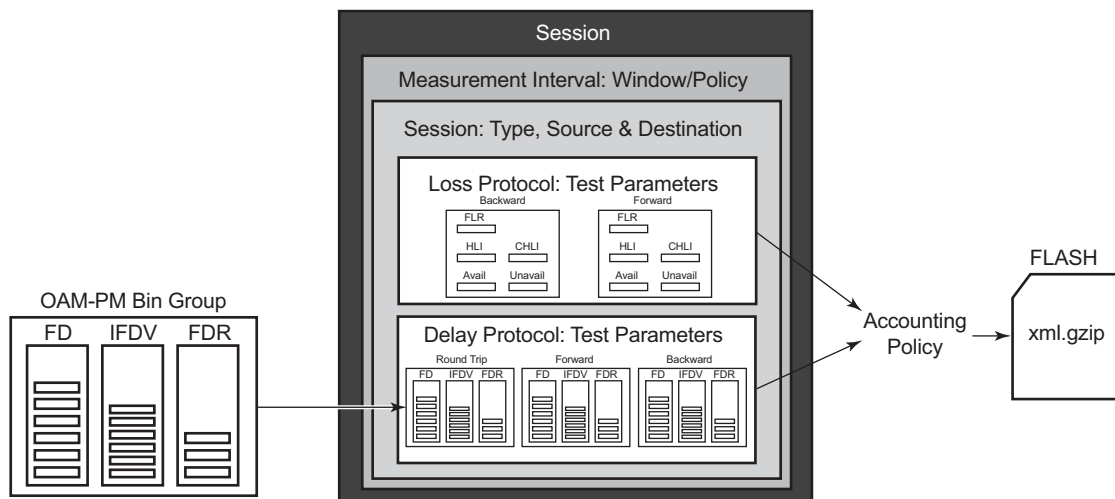
## OAM Performance Monitoring (OAM-PM)

OAM-PM provides an architecture for gathering and computing key performance indicators (KPIs) using standard protocols and a robust collection model. The architecture is comprised of the following foundational components.

- Session — The overall collection of different tests, test parameters, measurement intervals, and mappings to configured storage models. It is the overall container that defines the attributes of the session.
- Standard PM Packets — The protocols defined by various standards bodies which contains the necessary fields to collect statistical data for the performance attribute they represent. OAM-PM leverages single-ended protocols. Single-ended protocols typically follow a message response model, message sent by a launch point, response updated and reflected by a responder.
- Measurement Intervals (MI) — Time-based non-overlapping windows that capture all results that are received in that window of time
- Data Structures — The unique counters and measurement results that represent the specific protocol
- Bin group — Ranges in microseconds that count the results that fit into the range

Figure 22 shows the hierarchy of the architecture. This diagram is only meant to show the relationship between the components. It is not meant to depict all details of the required parameters.

**Figure 22: OAM-PM Architecture Hierarchy**



al\_0386



OAM-PM configurations are not dynamic environments. All aspects of the architecture must be carefully considered before configuring the various architectural components, making external references to other related components, or activating the OAM-PM architecture. No modifications are allowed to any components that are active or have any active sub-components. Any function being referenced by an active OAM-PM function or test cannot be modified or shut down. For example, to change any configuration element of a session, all active tests must be in a shutdown state. To change any bin group configuration (described later in this section) all sessions that reference the bin group must have every test shutdown. The description parameter is the only exception to this rule.

Session source and destination configuration parameters are not validated by the test that makes use of that information. Once the test is activated with a **no shutdown** command, the test engine will attempt to send the test packets even if the session source and destination information does not accurately represent the entity that must exist to successfully transmit packets. If the entity does not exist, the transmit count for the test will be zero.

OAM-PM is not a hitless operation. If a high availability event occurs that causes the backup CPM to become the active CPM, or when ISSU functions are performed, the test data will not be properly reported. There is no synchronization of state between the active and the backup control modules. All OAM-PM statistics stored in volatile memory will be lost. Once the reload or high availability event is completed and all services are operational then the OAM-PM functions will commence.

It is possible that during times of network convergence, high CPU utilizations, or contention for resources, OAM-PM may not be able to detect changes to an egress connection or allocate the necessary resources to perform its tasks.

This section describes the following topics:

- [Session on page 194](#)
- [Standard PM Packets on page 194](#)
- [Measurement Intervals on page 195](#)
- [Data Structures and Storage on page 201](#)
- [Bin Groups on page 204](#)
- [Relating the Components on page 205](#)
- [Monitoring on page 206](#)

## Session

This is the overall collection of different tests, the test parameters, measurement intervals, and mapping to configured storage models. It is the overall container that defines the attributes of the session.

- Session Type — The impetus of the test, which is either proactive (default) or on-demand. Individual test timing parameters are influenced by this setting. A proactive session will start immediately following the execution of a **no shutdown** command for the test. A proactive test will continue to execute until a manual shutdown stops the individual test. On-demand tests will also start immediately following the **no shutdown** command. However, the operator can override the **no test-duration** default and configure a fixed amount of time that the test will execute, up to 24 hours (86400 seconds). If an on-demand test is configured with a test-duration, it is important to shut down tests when they are completed. In the event of a high availability event causing the backup CPM to become the active CPM, all on-demand tests that have a test-duration statement will restart and run for the configured amount of time regardless of their progress on the previously active CPM.
- Test Family — The main branch of testing that addresses a specific technology. The available test for the session are based on the test family. The destination, source, and priority are common to all tests under the session and are defined separately from the individual test parameters.
- Test Parameters — The parameters included in individual tests, as well as the associated parameters including start and stop times and the ability to activate and deactivate the individual test.
- Measurement Interval — The assignment of collection windows to the session with the appropriate configuration parameters and accounting policy for that specific session.

The session can be viewed as the single container that brings all aspects of individual tests and the various OAM-PM components under a single umbrella. If any aspects of the session are incomplete, the individual test cannot be activated with a **no shutdown** command, and an "Invalid ethernet session parameters" error will occur.

## Standard PM Packets

A number of standards bodies define performance monitoring packets that can be sent from a source, processed, and responded to by a reflector. The protocols available to carry out the measurements are based on the test family type configured for the session.

Ethernet PM delay measurements are carried out using the Two Way Delay Measurement Protocol version 1 (DMMv1) defined in Y.1731 by the ITU-T. This allows for the collection of Frame Delay

(FD), InterFrame Delay Variation (IFDV), Frame Delay Range (FDR), and Mean Frame Delay (MFD) measurements for round trip, forward, and backward directions.

DMMv1 adds the following to the original DMM definition:

- The Flag Field (1 bit – LSB) is defined as the Type (Proactive=1 | On-Demand=0)
- The TestID TLV (32 bits) is carried in the Optional TLV portion of the PDU

DMMv1 and DMM are backwards compatible and the interaction is defined in Y.1731 ITU-T-2011 Section 11 "OAM PDU validation and versioning".

Ethernet PM loss measurements are carried out using Synthetic Loss Measurement (SLM), which is defined in Y.1731 by the ITU-T. This allows for the calculation of Frame Loss Ratio (FLR) and availability.

A session can be configured with one or more tests. Depending on the session test type family, one or more test configurations may need to be included in the session to gather both delay and loss performance information. Each test that is configured shares the common session parameters and the common measurement intervals. However, each test can be configured with unique per-test parameters. Using Ethernet as an example, both DMM and SLM would be required to capture both delay and loss performance data.

Each test must be configured with a *TestID* as part of the test parameters, which uniquely identifies the test within the specific protocol. A *TestID* must be unique within the same test protocol. Again using Ethernet as an example, DMM and SLM tests within the same session can use the same *TestID* because they are different protocols. However, if a *TestID* is applied to a test protocol (like DMM or SLM) in any session, it cannot be used for the same protocol in any other session. When a *TestID* is carried in the protocol, as it is with DMM and SLM, this value does not have global significance. When a responding entity must index for the purpose of maintaining sequence numbers, as in the case of SLM, the *TestID*, Source MAC, and Destination MAC are used to maintain the uniqueness of the responder. This means that the *TestID* has only local, and not global, significance.

## Measurement Intervals

A measurement interval is a window of time that compartmentalizes the gathered measurements for an individual test that have occurred during that time. Allocation of measurement intervals, which equates to system memory, is based on the metrics being collected. This means that when both delay and loss metrics are being collected, they allocate their own set of measurement intervals. If the operator is executing multiple delay and loss tests under a single session, then multiple measurement intervals will be allocated, with one interval allocated per criteria per test.

Measurement intervals can be 15 minutes (**15-min**), one hour (**1-hour**) and 1 day (**1-day**) in duration. The **boundary-type** defines the start of the measurement interval and can be aligned to the local time of day clock, with or without an optional offset. The **boundary-type** can be aligned using the **test-aligned** option, which means that the start of the measurement interval coincides with the activation of the test. By default the start boundary is clock-aligned without an offset. When this configuration is deployed, the measurement interval will start at zero, in relation to the length. When a boundary is clock-aligned and an offset is configured, the specified amount of time will be applied to the measurement interval. Offsets are configured on a per-measurement interval basis and only applicable to clock-aligned measurement intervals. Only offsets less than the measurement interval duration are allowed. The table below provides some examples of the start times of each measurement interval.

**Table 16: Measurement Interval Start Times**

Offset	15-min	1-hour	1-day
0 (default)	00, 15, 30, 45	00 (top of the hour)	midnight
10 minutes	10, 25, 40, 55	10 min after the hour	10 min after midnight
30 minutes	rejected	30 min after the hour	30 min after midnight
60 minutes	rejected	rejected	01:00 AM

Although test-aligned approaches may seem beneficial for simplicity, there are some drawbacks that need to be considered. The goal of the time-based and well defined collection windows allows for the comparison of measurements across common windows of time throughout the network and for relating different tests or sessions. It is suggested that proactive sessions use the default clock-aligned boundary type. On-demand sessions may make use of test-aligned boundaries. On-demand tests are typically used for troubleshooting or short term monitoring that does not require alignment or comparison to other PM data.

The statistical data collected and the computed results from each measurement interval are maintained in volatile system memory by default. The number of intervals stored is configurable per measurement interval. Different measurement intervals will have different defaults and ranges. The **interval-stored** parameter defines the number of completed individual test runs to store in volatile memory. There is an additional allocation to account for the active measurement interval. In order to look at the statistical information for the individual tests and a specific measurement interval stored in volatile memory, the **show oam-pm statistics ... interval-number** command can be used. If there is an active test, it can be viewed by using the interval number 1. In this case, the first completed record would be interval number 2, and previously completed records would increment up to the maximum intervals stored value plus one.

As new tests for the measurement interval are completed, the older entries are renumbered to maintain their relative position to the current test. If the retained test data for a measurement interval consumes the final entry, any subsequent entries cause the removal of the oldest data.

There are drawbacks to this storage model. Any high availability function that causes an active CPM switch will flush the results that are in volatile memory. Another consideration is the large amount of system memory consumed using this type of model. Given the risks and resource consumption this model incurs, an alternate method of storage is supported. An accounting policy can be applied to each measurement interval in order to write the completed data in system memory to non-volatile flash memory in an XML format. The amount of system memory consumed by historically completed test data must be balanced with an appropriate accounting policy. It is recommended that only necessary data be stored in non-volatile memory to avoid unacceptable risk and unnecessary resource consumption. It is further suggested that a large overlap between the data written to flash memory and stored in volatile memory is unnecessary.

The statistical information in system memory is also available through SNMP. If this method is chosen, a balance must be struck between the intervals retained and the times at which the SNMP queries collect the data. Determining the collection times through SNMP must be done with caution. If a file is completed while another file is being retrieved through SNMP, then the indexing will change to maintain the relative position to the current run. Proper spacing of the collection is key to ensuring data integrity.

The OAM-PM XML file contains the following keywords and MIB references.

**Table 17: OAM-PM XML Keywords and MIB Reference**

<b>XML File Keyword</b>	<b>Description</b>	<b>TIMETRA-OAM-PM-MIB Object</b>
oampm	—	None - header only
<b>Keywords Shared by all OAM-PM Protocols</b>		
sna	OAM-PM session name	tmnxOamPmCfgSessName
mi	Measurement interval record	None - header only
dur	Measurement interval duration (minutes)	tmnxOamPmCfgMeasIntvlDuration (enumerated)
ivl	Measurement interval number	tmnxOamPmStsIntvlNum
sta	Start timestamp	tmnxOamPmStsBaseStartTime
ela	Elapsed time (seconds)	tmnxOamPmStsBaseElapsedTime
ftx	Frames sent	tmnxOamPmStsBaseTestFramesTx
frx	Frames received	tmnxOamPmStsBaseTestFramesRx
sus	Suspect flag	tmnxOamPmStsBaseSuspect
dmm	Delay record	None - header only

**Table 17: OAM-PM XML Keywords and MIB Reference (Continued)**

<b>XML File Keyword</b>	<b>Description</b>	<b>TIMETRA-OAM-PM-MIB Object</b>
mdr	Minimum frame delay, round-trip	tmnxOamPmStsDelayDmm2wyMin
xdr	Maximum frame delay, round-trip	tmnxOamPmStsDelayDmm2wyMax
adr	Average frame delay, round-trip	tmnxOamPmStsDelayDmm2wyAvg
mdf	Minimum frame delay, forward	tmnxOamPmStsDelayDmmFwdMin
xdf	Maximum frame delay, forward	tmnxOamPmStsDelayDmmFwdMax
adf	Average frame delay, forward	tmnxOamPmStsDelayDmmFwdAvg
mdb	Minimum frame delay, backward	tmnxOamPmStsDelayDmmBwdMin
xdb	Maximum frame delay, backward	tmnxOamPmStsDelayDmmBwdMax
adb	Average frame delay, backward	tmnxOamPmStsDelayDmmBwdAvg
mvr	Minimum inter-frame delay variation, round-trip	tmnxOamPmStsDelayDmm2wyMin
xvr	Maximum inter-frame delay variation, round-trip	tmnxOamPmStsDelayDmm2wyMax
avr	Average inter-frame delay variation, round-trip	tmnxOamPmStsDelayDmm2wyAvg
mvf	Minimum inter-frame delay variation, forward	tmnxOamPmStsDelayDmmFwdMin
xvf	Maximum inter-frame delay variation, forward	tmnxOamPmStsDelayDmmFwdMax
avf	Average inter-frame delay variation, forward	tmnxOamPmStsDelayDmmFwdAvg
mvb	Minimum inter-frame delay variation, backward	tmnxOamPmStsDelayDmmBwdMin
xvb	Maximum inter-frame delay variation, backward	tmnxOamPmStsDelayDmmBwdMax
avb	Average inter-frame delay variation, backward	tmnxOamPmStsDelayDmmBwdAvg
mrr	Minimum frame delay range, round-trip	tmnxOamPmStsDelayDmm2wyMin
xrr	Maximum frame delay range, round-trip	tmnxOamPmStsDelayDmm2wyMax
arr	Average frame delay range, round-trip	tmnxOamPmStsDelayDmm2wyAvg
mrf	Minimum frame delay range, forward	tmnxOamPmStsDelayDmmFwdMin
xrf	Maximum frame delay range, forward	tmnxOamPmStsDelayDmmFwdMax
arf	Average frame delay range, forward	tmnxOamPmStsDelayDmmFwdAvg
mrb	Minimum frame delay range, backward	tmnxOamPmStsDelayDmmBwdMin

Table 17: OAM-PM XML Keywords and MIB Reference (Continued)

XML File Keyword	Description	TIMETRA-OAM-PM-MIB Object
xrb	Maximum frame delay range, backward	tmnxOamPmStsDelayDmmBwdMax
arb	Average frame delay range, backward	tmnxOamPmStsDelayDmmBwdAvg
fdr	Frame delay bin record, round-trip	None - header only
fdf	Frame delay bin record, forward	None - header only
fdb	Frame delay bin record, backward	None - header only
fvr	Inter-frame delay variation bin record, round-trip	None - header only
fvf	Inter-frame delay variation bin record, forward	None - header only
fvb	Inter-frame delay variation bin record, backward	None - header only
fr	Frame delay range bin record, round-trip	None - header only
frf	Frame delay range bin record, forward	None - header only
frb	Frame delay range bin record, backward	None - header only
lbo	Configured lower bound of the bin	tmnxOamPmCfgBinLowerBound
cnt	Number of measurements within the configured delay range <b>Note:</b> The session_name, interval_duration, interval_number, {fd, fdr, ifdv}, bin_number, and {forward, backward, round-trip} indices are provided by the surrounding XML context	tmnxOamPmStsDelayDmmBinFwdCount tmnxOamPmStsDelayDmmBinBwdCount tmnxOamPmStsDelayDmmBin2wyCount
slm	Synthetic loss measurement record	None - header only
txf	Transmitted frames in the forward direction	tmnxOamPmStsLossSlmTxFwd
rx	Received frames in the forward direction	tmnxOamPmStsLossSlmRxFwd
txb	Transmitted frames in the backward direction	tmnxOamPmStsLossSlmTxBwd
rx	Received frames in the backward direction	tmnxOamPmStsLossSlmRxBwd
avf	Available count in the forward direction	tmnxOamPmStsLossSlmAvailIndFwd
avb	Available count in the backward direction	tmnxOamPmStsLossSlmAvailIndBwd
uvf	Unavailable count in the forward direction	tmnxOamPmStsLossSlmUnavlIndFwd
uvb	Unavailable count in the backward direction	tmnxOamPmStsLossSlmUnavlIndBwd

**Table 17: OAM-PM XML Keywords and MIB Reference (Continued)**

XML File Keyword	Description	TIMETRA-OAM-PM-MIB Object
uaf	Undetermined available count in the forward direction	tmnxOamPmStsLossSImUndtAvlFwd
uab	Undetermined available count in the backward direction	tmnxOamPmStsLossSImUndtAvlBwd
uuf	Undetermined unavailable count in the forward direction	tmnxOamPmStsLossSImUndtUnavlFwd
uub	Undetermined unavailable count in the backward direction	tmnxOamPmStsLossSImUndtUnavlBwd
hlf	Count of HLIs in the forward direction	tmnxOamPmStsLossSImHliFwd
hlb	Count of HLIs in the backward direction	tmnxOamPmStsLossSImHliBwd
chf	Count of CHLIs in the forward direction	tmnxOamPmStsLossSImChliFwd
chb	Count of CHLIs in the backward direction	tmnxOamPmStsLossSImChliBwd
mff	Minimum FLR in the forward direction	tmnxOamPmStsLossSImMinFlrFwd
xff	Maximum FLR in the forward direction	tmnxOamPmStsLossSImMaxFlrFwd
aff	Average FLR in the forward direction	tmnxOamPmStsLossSImAvgFlrFwd
mfb	Minimum FLR in the backward direction	tmnxOamPmStsLossSImMinFlrBwd
xfb	Maximum FLR in the backward direction	tmnxOamPmStsLossSImMaxFlrBwd
afb	Average FLR in the backward direction	tmnxOamPmStsLossSImAvgFlrBwd

By default, the 15-min measurement interval stores 33 test runs (32+1) with a configurable range of 1 to 96, and the 1-hour measurement interval stores 9 test runs (8+1) with a configurable range of 1 to 24. The only storage for the 1-day measurement interval is 2 (1+1). This value for the 1-day measurement interval cannot be changed.

All three measurement intervals may be added to a single session if required. Each measurement interval that is included in a session is updated simultaneously for each test that is executing. If a measurement interval length is not required, it should not be configured. In addition to the three predetermined length measurement intervals, a fourth “always on” raw measurement interval is allocated at test creation. Data collection for the raw measurement interval commences immediately following the execution of a **no shutdown** command. It is a valuable tool for assisting in real-time troubleshooting as it maintains the same performance information and relates to the same bins as the fixed length collection windows. The operator may clear the contents of the raw



measurement interval and flush stale statistical data in order to look at current conditions. This measurement interval has no configuration options, cannot be written to flash memory, and cannot be disabled; It is a single never-ending window.

Memory allocation for the measurement intervals is performed when the test is configured. Volatile memory is not flushed until the test is deleted from the configuration, a high availability event causes the backup CPM to become the newly active CPM, or some other event clears the active CPM system memory. Shutting down a test does not release the allocated memory for the test.

Measurement intervals also include a suspect flag. The suspect flag is used to indicate that data collected in the measurement interval may not be representative. The flag will be set to true only under the following conditions.

- Time of day clock is adjusted by more than 10 seconds
- Test start does not align with the start boundary of the measurement interval. This would be common for the first execution for clock aligned tests.
- Test stopped before the end of the measurement interval boundary

The suspect flag is not set when there are times of service disruption, maintenance windows, discontinuity, low packet counts, or other such events. Higher level systems would be required to interpret and correlate those types of event for measurement intervals which executed during the time that relate to the specific interruption or condition. Since each measurement interval contains a start and stop time, the information is readily available for higher level systems to discount the specific windows of time.

## Data Structures and Storage

There are two main metrics that are the focus of OAM-PM: delay and loss. The different metrics have two unique storage structures and will allocate their own measurement intervals for these structures. This occurs regardless of whether the performance data is gathered with a single packet or multiple packet types.

Delay metrics include Frame Delay (FD), InterFrame Delay Variation (IFDV), Frame Delay Range (FDR) and Mean Frame Delay (MFD). Unidirectional and round trip results are stored for each metric.

- Frame Delay — The amount of time required to send and receive the packet
- InterFrame Delay Variation — The difference in the delay metrics between two adjacent packets

- **Frame Delay Range** — The difference between the minimum frame delay and the individual packet
- **Mean Frame Delay** — The mathematical average for the frame delay over the entire window.

FD, IFDV and FDR statistics are binnable results. FD, IFDV, FDR and MFD all include minimum, maximum, and average values. Unidirectional and round trip results are stored for each metric.

Unidirectional frame delay and frame delay range measurements require exceptional time of day clock synchronization. If the time of day clock does not exhibit extremely tight synchronization, unidirectional measurements will not be representative. In one direction, the measurement will be artificially increased by the difference in the clocks. In the other direction, the measurement will be artificially decreased by the difference in the clocks. This level of clocking accuracy is not available with NTP. In order to achieve this level of time of day clock synchronization, Precision Time Protocol (PTP) 1588v2 should be considered.

Round trip metrics do not require clock synchronization between peers, since the four timestamps allow for accurate representation of the round trip delay. The mathematical computation removes remote processing and any difference in time of day clocking. Round trip measurements do require stable local time of day clocks.

Any delay metric that is negative will be treated as zero and placed in bin 0, the lowest bin which has a lower boundary of 0 microseconds.

Delay results are mapped to the measurement interval that is active when the result arrives back at the source.

There are no supported log events based on delay metrics.

Loss metrics are only unidirectional and will report frame loss ratio (FLR) and availability information. Frame loss ratio is the computation of loss (lost/sent) over time. Loss measurements during periods of unavailability are not included in the FLR calculation as they are counted against the unavailability metric.

Availability requires relating three different functions. First, the individual probes are marked as available or unavailable based on sequence numbers in the protocol. A number of probes are rolled up into a small measurement window, typically 1 s. Frame loss ratio is computed over all the probes in a small window. If the resulting percentage is higher than the configured threshold, the small window is marked as unavailable. If the resulting percentage is lower than the threshold, the small window is marked as available. A sliding window is defined as some number of small windows, typically 10. The sliding window is used to determine availability and unavailability events. Switching from one state to the other requires every small window in the sliding window to be the same state and different from the current state.

Availability and unavailability counters are incremented based on the number of small windows that have occurred in all available and unavailable windows.

Availability and unavailability using synthetic loss measurements is meant to capture the loss behavior for the service. It is not meant to capture and report on service outages or communication failures. Communication failures of a bidirectional or unidirectional nature must be captured using some other means of connectivity verification, alarming, or continuity checking. During times of complete or extended failure periods it becomes necessary to timeout individual test probes. It is not possible to determine the direction of the loss because no response packets are being received back on the source. In this case, the statistics calculation engine maintains the previous state, updating the appropriate directional availability or unavailability counter. At the same time, an additional per-direction undetermined counter is updated. This undetermined counter is used to indicate that the availability or unavailability statistics could not be determined for a number of small windows.

During connectivity outages, the higher level systems can be used to discount the loss measurement interval, which covers the same span as the outage.

Availability and unavailability computations may delay the completion of a measurement interval. The declaration of a state change or the delay to a closing a measurement interval could be equal to the length of the sliding window and the timeout of the last packet. Closing of a measurement interval cannot occur until the sliding window has determined availability or unavailability. If the availability state is changing and the determination is crossing two measurement intervals, the measurement interval will not complete until the declaration has occurred. Typically, standard bodies indicate the timeout per packet. In the case of Ethernet, DMMv1, and SLM, timeout values are set at 5 s and cannot be configured.

There are no log events based on availability or unavailability state changes.

During times of availability, there can be times of high loss intervals (HLI) or consecutive high loss intervals (CHLI). These are indicators that the service was available but individual small windows or consecutive small windows experienced frame loss ratios exceeding the configured acceptable limit. A HLI is any single small window that exceeds the configured frame loss ratio. This could equate to a severely errored second, assuming the the small window is one second in length. A CHIL is a consecutive high loss interval that exceeds a consecutive threshold within the sliding window. Only one HLI will be counted for a window.

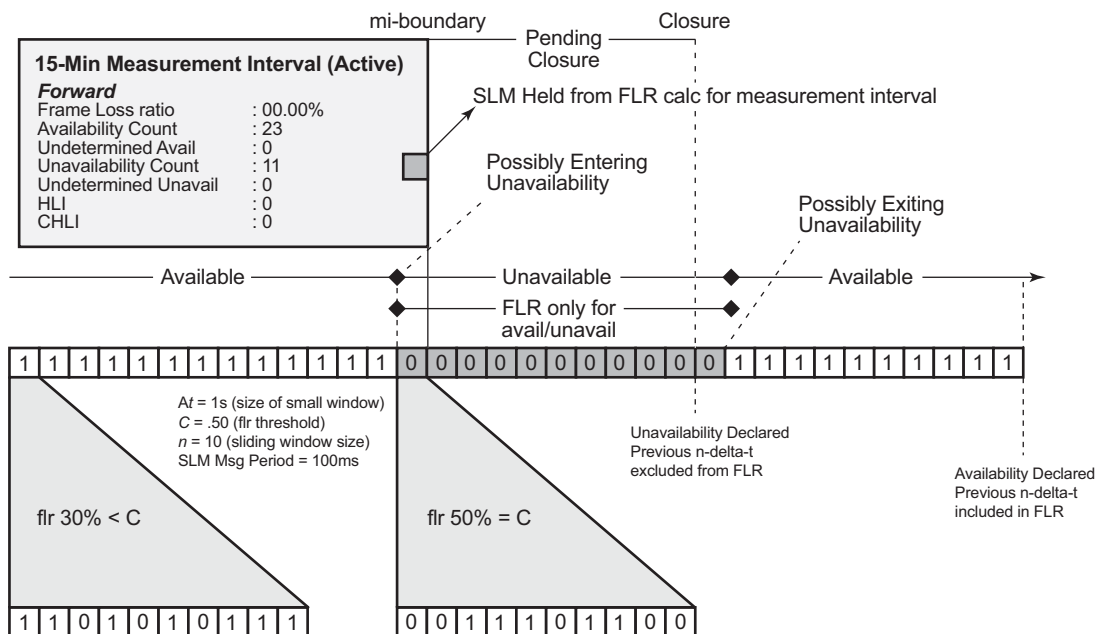
Availability can only be reasonably determined with synthetic packets. This is because because the synthetic packet is the packet being counted and provides a uniform packet flow that can be used for the computation. Transmit and receive counter-based approaches cannot reliably be used to determine availability because there is no guarantee that service data is on the wire, or the service data on the wire uniformity could make it difficult to make a declaration valid.

[Figure 23](#) shows loss in a single direction using synthetic packets, and demonstrates what happens when a possible unavailability event crosses a measurement interval boundary. In the diagram, the

## Bin Groups

first 13 small windows are all marked available (1), which means that the loss probes that fit into each of those small windows did not equal or exceed a frame loss ratio of 50%. The next 11 small windows are marked as unavailable, which means that the loss probes that fit into each of those small windows were equal to or above a frame loss ratio of 50%. After the 10th consecutive small window of unavailability, the state transitions from available to unavailable. The 25th small window is the start of the new available state which is declared following the 10th consecutive available small window. Notice that the frame loss ratio is 00.00%; this is because all the small windows that are marked as unavailable are counted towards unavailability, and as such are excluded from impacting the FLR. If there were any small windows of unavailability that were outside of an unavailability event, they would be marked as HLI or CHLI and be counted as part of the frame loss ratio.

**Figure 23: Evaluating and Computing Loss and Availability**



al\_0387

## Bin Groups

Bin groups are templates that are referenced by the session. Three types of binnable statistics are available: FD, IFDV, and FDR, all of which are available in forward, backward, and round trip directions. Each of these metrics can have up to ten bin groups configured to group the results. Bin groups are configured by indicating a lower boundary. Bin 0 has a lower boundary that is always zero and is not configurable. The microsecond range of the bins is the difference between the adjacent lower boundaries. For example, **bin-type fd bin 1** configured with **lower-bound 1000**

means that bin 0 will capture all frame delay statistics results between 0 and 1 ms. Bin 1 will capture all results above 1 ms and below the bin 2 lower boundary. The last bin to be configured would represent the bin that collects all the results at and above that value. Not all ten bins must be configured.

Each binnable delay metric type requires their own values for the bin groups. Each bin in a type is configurable for one value. It is not possible to configure a bin with different values for round trip, forward, and backward. Consideration must be given to the configuration of the boundaries that represent the important statistics for that specific service.

As stated earlier in this section, this is not a dynamic environment. If a bin group is being referenced by any active test the bin group cannot shutdown. In order to modify the bin group it must be shut down. If the configuration of a bin group must be changed, and a large number of sessions are referencing the bin group, migrating existing sessions to a new bin group with the new parameters can be considered in order to reduce the maintenance window. In order to modify any session parameter, every test in the session must be shut down.

Bin group 1 is the default bin group. Every session requires a bin group to be assigned. By default, bin group 1 is assigned to every OAM-PM session that does not have a bin group explicitly configured. Bin group 1 cannot be modified. The bin group 1 configuration parameters are as follows:

```

Configured Lower Bounds for Delay Measurement (DMM) Tests, in microseconds

Group Description Admin Bin FD(us) FDR(us) IPDV(us)

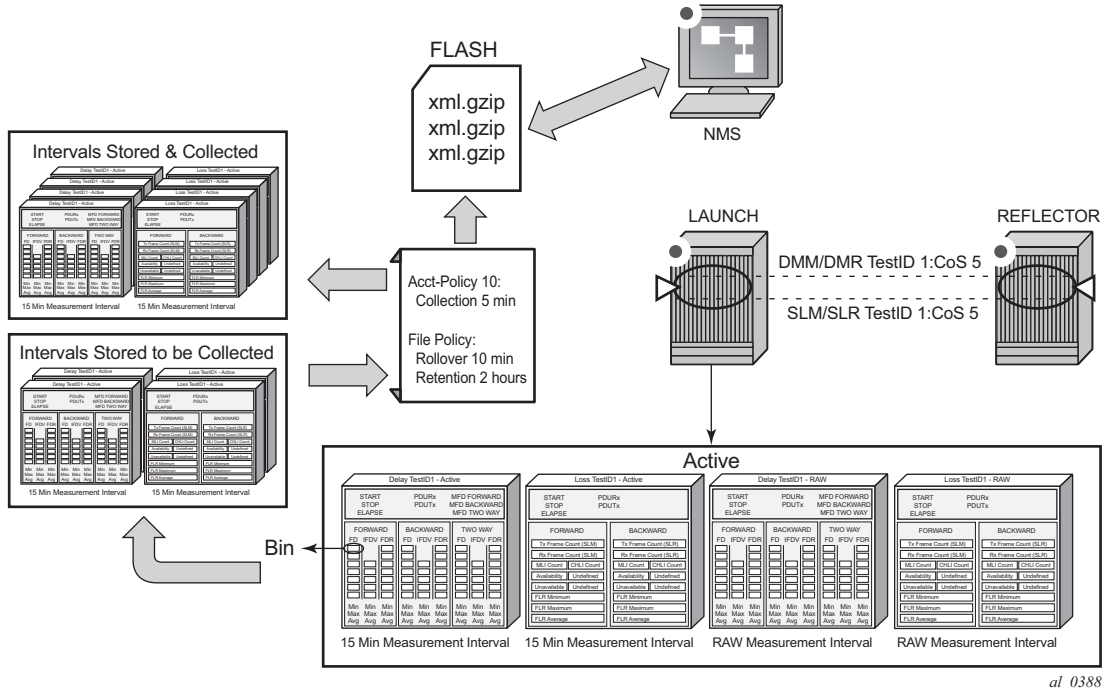
1 OAM PM default bin group (not* Up 0 0 0 0
 1 5000 5000 5000
 2 10000 - -

```

## Relating the Components

[Figure 24](#) shows the architecture of all of the OAM-PM concepts previously discussed. It shows a more detailed hierarchy than previously shown in the introduction. This shows the relationship between the tests, the measurement intervals, and the storage of the results.

Figure 24: Relating OAM-PM Components



al\_0388

## Monitoring

The following configuration examples are used to demonstrate the different show and monitoring commands available to check OAM-PM.

## Accounting Policy Configuration

```

config>log# info

file-id 1
 description "OAM PM XML file Paramaters"
 location cf2:
 rollover 10 retention 2
exit
accounting-policy 1
 description "Default OAM PM Collection Policy for 15-min Bins"
 record complete-pm
 collection-interval 5
 to file 1
 no shutdown

```

```

exit
log-id 1
exit

```

## ETH-CFM Configuration

```

config>eth-cfm# info

domain 12 format none level 2
 association 4 format string name "vpls4-0000001"
 bridge-identifier 4
 id-permission chassis
 exit
 ccm-interval 1
 remote-mepid 30
exit
exit

```

## Service Configuration

```

config>service>vpls# info

description "OAM PM Test Service to v30"
stp
 shutdown
exit
sap 1/1/10:4.* create
 eth-cfm
 mep 28 domain 12 association 4 direction up
 ccm-enable
 mac-address 00:00:00:00:00:28
 no shutdown
 exit
exit
exit
sap 1/2/1:4.* create
exit
no shutdown

```

## OAM-PM Configuration

```

config>oam-pm#info detail

bin-group 2 fd-bin-count 10 fdr-bin-count 2 ifdv-bin-count 10 create
no description
bin-type fd

```

## Monitoring

```
bin 1
 lower-bound 1000
exit
bin 2
 lower-bound 2000
exit
bin 3
 lower-bound 3000
exit
bin 4
 lower-bound 4000
exit
bin 5
 lower-bound 5000
exit
bin 6
 lower-bound 6000
exit
bin 7
 lower-bound 7000
exit
bin 8
 lower-bound 8000
exit
bin 9
 lower-bound 10000
exit
exit
bin-type fdr
 bin 1
 lower-bound 5000
 exit
exit
bin-type ifdv
 bin 1
 lower-bound 100
 exit
 bin 2
 lower-bound 200
 exit
 bin 3
 lower-bound 300
 exit
 bin 4
 lower-bound 400
 exit
 bin 5
 lower-bound 500
 exit
 bin 6
 lower-bound 600
 exit
 bin 7
 lower-bound 700
 exit
 bin 8
 lower-bound 800
```



```

 exit
 bin 9
 lower-bound 1000
 exit
 exit
 no shutdown
exit
session "eth-pm-service-4" test-family ethernet session-type proactive create
 bin-group 2
 no description
 meas-interval 15-mins create
 no accounting-policy
 boundary-type clock-aligned
 clock-offset 0
 intervals-stored 32
 exit
 ethernet
 dest-mac 00:00:00:00:00:30
 priority 0
 source mep 28 domain 12 association 4
 dmm test-id 10004 create
 data-tlv-size 1000
 interval 1000
 no test-duration
 no shutdown
 exit
 slm test-id 10004 create
 data-tlv-size 1000
 flr-threshold 50
 no test-duration
 timing frames-per-delta-t 10 consec-delta-t 10 interval 100
 chli-threshold 4
 no shutdown
 exit
 exit
exit
exit

```

## Show and Monitor Commands

```
show oam-pm bin-group
```

```

Configured Lower Bounds for Delay Measurement (DMM) Tests, in microseconds

Group Description Admin Bin FD(us) FDR(us) IFDV(us)

1 OAM PM default bin group (not* Up 0 0 0 0
 1 5000 5000 5000
 2 10000 - -

2 Up 0 0 0 0
 1 1000 5000 100
 2 2000 - 200
 3 3000 - 300
 4 4000 - 400

```

# Monitoring

5	5000	-	500
6	6000	-	600
7	7000	-	700
8	8000	-	800
9	10000	-	1000

-----  
 -----  
 \* indicates that the corresponding row element may have been truncated.

show oam-pm bin-group 2

-----  
 Configured Lower Bounds for Delay Measurement (DMM) Tests, in microseconds  
 -----

Group Description	Admin	Bin	FD(us)	FDR(us)	IFDV(us)
2	Up	0	0	0	0
		1	1000	5000	100
		2	2000	-	200
		3	3000	-	300
		4	4000	-	400
		5	5000	-	500
		6	6000	-	600
		7	7000	-	700
		8	8000	-	800
		9	10000	-	1000

-----  
 show oam-pm bin-group-using

=====  
 OAM Performance Monitoring Bin Group Configuration for Sessions  
 =====

Bin Group	Admin	Session	Session State
2	Up	eth-pm-service-4	Act

-----  
 show oam-pm bin-group-using bin-group 2

=====  
 OAM Performance Monitoring Bin Group Configuration for Sessions  
 =====

Bin Group	Admin	Session	Session State
2	Up	eth-pm-service-4	Act

-----  
 show oam-pm sessions test-family ethernet

=====  
 OAM Performance Monitoring Session Summary for the Ethernet Test Family  
 =====

Session	State	Bin Group	Sess Type	Test Types
eth-pm-service-4	Act	2	proactive	DMM SLM

-----  
 show oam-pm session "eth-pm-service-4" all

-----  
Basic Session Configuration  
-----

Session Name : eth-pm-service-4  
 Description : (Not Specified)  
 Test Family : ethernet                      Session Type : proactive  
 Bin Group : 2

-----  
Ethernet Configuration  
-----

Source MEP : 28                              Priority : 0  
 Source Domain : 12                          Dest MAC Address : 00:00:00:00:00:30  
 Source Assoc'n : 4

-----  
DMM Test Configuration and Status  
-----

Test ID : 10004                              Admin State : Up  
 Oper State : Up                              Data TLV Size : 1000 octets  
 On-Demand Duration: Not Applicable      On-Demand Remaining: Not Applicable  
 Interval : 1000 ms

-----  
SLM Test Configuration and Status  
-----

Test ID : 10004                              Admin State : Up  
 Oper State : Up                              Data TLV Size : 1000 octets  
 On-Demand Duration: Not Applicable      On-Demand Remaining: Not Applicable  
 Interval : 100 ms  
 CHLI Threshold : 4 HLIs                      Frames Per Delta-T : 10 SLM frames  
 Consec Delta-Ts : 10                          FLR Threshold : 50%

-----  
15-mins Measurement Interval Configuration  
-----

Duration : 15-mins                              Intervals Stored : 32  
 Boundary Type : clock-aligned              Clock Offset : 0 seconds  
 Accounting Policy : none

-----  
Configured Lower Bounds for Delay Measurement (DMM) Tests, in microseconds  
-----

Group Description	Admin Bin	FD(us)	FDR(us)	IFDV(us)
2	Up 0	0	0	0
	1	1000	5000	100
	2	2000	-	200
	3	3000	-	300
	4	4000	-	400
	5	5000	-	500
	6	6000	-	600

# Monitoring

```

7 7000 - 700
8 8000 - 800
9 10000 - 1000

```

```

show oam-pm statistics session "eth-pm-service-4" dmm meas-interval 15-mins interval-number
2 all

```

```

Start (UTC) : 2014/02/01 10:00:00 Status : completed
Elapsed (seconds) : 900 Suspect : no
Frames Sent : 900 Frames Received : 900

```

```

Bin Type Direction Minimum (us) Maximum (us) Average (us)

FD Forward 0 8330 712
FD Backward 143 11710 2605
FD Round Trip 1118 14902 3111
FDR Forward 0 8330 712
FDR Backward 143 11710 2605
FDR Round Trip 0 13784 1990
IFDV Forward 0 8330 431
IFDV Backward 1 10436 800
IFDV Round Trip 2 13542 1051

```

```

Frame Delay (FD) Bin Counts

```

```

Bin Lower Bound Forward Backward Round Trip

0 0 us 624 53 0
1 1000 us 229 266 135
2 2000 us 29 290 367
3 3000 us 4 195 246
4 4000 us 7 71 94
5 5000 us 5 12 28
6 6000 us 1 7 17
7 7000 us 0 1 5
8 8000 us 1 4 3
9 10000 us 0 1 5

```

```

Frame Delay Range (FDR) Bin Counts

```

```

Bin Lower Bound Forward Backward Round Trip

0 0 us 893 875 873
1 5000 us 7 25 27

```

```

Inter-Frame Delay Variation (IFDV) Bin Counts

```

```

Bin Lower Bound Forward Backward Round Trip

```

```

0 0 us 411 162 96
1 100 us 113 115 108
2 200 us 67 84 67
3 300 us 56 67 65
4 400 us 36 46 53
5 500 us 25 59 54
6 600 us 25 27 38
7 700 us 29 34 22
8 800 us 41 47 72
9 1000 us 97 259 325

```

```

show oam-pm statistics session "eth-pm-service-4" slm meas-interval 15-mins interval-num-
ber 2

```

```

Start (UTC) : 2014/02/01 10:00:00 Status : completed
Elapsed (seconds) : 900 Suspect : no
Frames Sent : 9000 Frames Received : 9000

```

```

 Frames Sent Frames Received

Forward 9000 9000
Backward 9000 9000

```

```

Frame Loss Ratios

```

```

 Minimum Maximum Average

Forward 0.000% 0.000% 0.000%
Backward 0.000% 0.000% 0.000%

```

```

Availability Counters (Und = Undetermined)

```

```

 Available Und-Avail Unavailable Und-Unavail HLI CHLI

Forward 900 0 0 0 0 0
Backward 900 0 0 0 0 0

```

```

show oam-pm statistics session "eth-pm-service-4" dmm meas-interval raw

```

```

Start (UTC) : 2014/02/01 09:43:58 Status : in-progress
Elapsed (seconds) : 2011 Suspect : yes
Frames Sent : 2011 Frames Received : 2011

```

```

Bin Type Direction Minimum (us) Maximum (us) Average (us)

FD Forward 0 11670 632

```

# Monitoring

FD	Backward	0	11710	2354
FD	Round Trip	1118	14902	2704
FDR	Forward	0	11670	611
FDR	Backward	0	11710	2353
FDR	Round Trip	0	13784	1543
IFDV	Forward	0	10027	410
IFDV	Backward	0	10436	784
IFDV	Round Trip	0	13542	1070

## Frame Delay (FD) Bin Counts

Bin	Lower Bound	Forward	Backward	Round Trip
0	0 us	1465	252	0
1	1000 us	454	628	657
2	2000 us	62	593	713
3	3000 us	8	375	402
4	4000 us	11	114	153
5	5000 us	7	26	41
6	6000 us	2	10	20
7	7000 us	0	2	8
8	8000 us	1	10	11
9	10000 us	1	1	6

## Frame Delay Range (FDR) Bin Counts

Bin	Lower Bound	Forward	Backward	Round Trip
0	0 us	2001	1963	1971
1	5000 us	11	49	41

## Inter-Frame Delay Variation (IFDV) Bin Counts

Bin	Lower Bound	Forward	Backward	Round Trip
0	0 us	954	429	197
1	100 us	196	246	197
2	200 us	138	168	145
3	300 us	115	172	154
4	400 us	89	96	136
5	500 us	63	91	108
6	600 us	64	53	89
7	700 us	61	55	63
8	800 us	112	82	151
9	1000 us	219	619	771

show oam-pm statistics session "eth-pm-service-4" slm meas-interval raw

Start (UTC)	: 2014/02/01 09:44:03	Status	: in-progress
Elapsed (seconds)	: 2047	Suspect	: yes







## Diagnostics Command Reference

- [OAM Commands on page 217](#)
- [OAM Performance Monitoring and Binning Commands on page 229](#)
- [SAA Commands on page 231](#)

---

### OAM Commands

#### Base Operational Commands

##### GLOBAL

- **ping** [*ip-address | dns-name*] [**rapid** | **detail**] [**ttl** *time-to-live*] [**tos** *type-of-service*] [**size** *bytes*] [**pattern pattern**] [**source** *ip-address | dns-name*] [**interval** *seconds*] [ { **next-hop** *ip-address* } | { **interface** *interface-name* } ] [**bypass-routing**] [**count** *requests*] [**do-not-fragment**] [**router** *router-instance* | **service-name** *service-name*] [**timeout** *timeout*] [**fc** *fc-name*]
  - **tracert** [*ip-address | dns-name*] [**tth** *tth*] [**wait** *milli-seconds*] [**no-dns**] [**source** *ip-address*] [**tos** *type-of-service*] [**router** [*router-instance* | **service-name** *service-name*]
  - **oam**
    - **dns target-addr** *dns-name* **name-server** *ip-address* [**source** *ip-address*] [**count** *send-count*] [**timeout** *timeout*] [**interval** *interval*] [**record-type** { *ipv4-a-record* | *ipv6-aaaa-record* }]
    - **saa test-name** [**owner** *test-owner*] { **start** | **stop** } [**no-accounting**]
- 

#### LSP Diagnostics

##### GLOBAL

- **oam**
  - **lsp-ping bgp-label-prefix** *ip-prefix/mask* [**path-destination** *ip-address*] [**interface** *if-name* | **next-hop** *ip-address*]
  - **lsp-ping** *lsp-name* [**path** *path-name*]
  - **lsp-ping prefix** *ip-prefix/mask* [**path-destination** *ip-address*] [**interface** *if-name* | **next-hop** *ip-address*]
  - **lsp-ping static** *lsp-name* [**assoc-channel** *ipv4* | **non-ip** | **none**] [**dest-global-id** *global-id* **dest-node-id** *node-id*] [**force**] [**path-type** *active* | *working* | *protect*]
  - **lsp-trace bgp-label-prefix** *ip-prefix/mask* [**path-destination** *ip-address*] [**interface** *if-name* | **next-hop** *ip-address*] [**downstream-map-tlv** *dsmap* | *ddmap* | *none*]
  - **lsp-trace** *lsp-name* [**path** *path-name*]
  - **lsp-trace prefix** *ip-prefix/mask* [**path-destination** *ip-address*] [**interface** *if-name* | **next-hop** *ip-address*]
  - **lsp-trace static** *lsp-name* [**assoc-channel** *ipv4* | **non-ip** | **none**] [**path-type** *active* | *working* | *protect*]
  - **p2mp-lsp-ping** { { *lsp-name* [**p2mp-instance** *instance-name*] [**s2l-dest-address** *ip-address*... [*ip-address*...*up-to-5*]] [**tth** *label-ttl*] } | { **ldp** *p2mp-id* [**sender-addr** *ip-address*] [**leaf-addr** *ip-address*... [*ip-address*...*up-to-5*]] } } [**fc** *fc-name*] [**size** *octets*] [**timeout** *timeout*] [**interval** *interval*] [**detail**]

## OAM Commands

- **p2mp-lsp-trace** {*lsp-name* [**p2mp-instance** *instance-name*] [**s2l-dest-address** *ip-address...*[*ip-address...up-to-5*]]} | {**ldp** *p2mp-id*} [**fc** *fc-name*] [**size** *octets*] [**max-fail** *no-response-count*] [**probe-count** *probes-per-hop*] [**min-ttl** *min-label-ttl*] [**max-ttl** *max-label-ttl*] [**timeout** *timeout*] [**interval** *interval*] [**detail**]

## LDP Diagnostics

### GLOBAL

- **oam**
  - **ldp-treetrace** {**prefix** *ip-prefix/mask*} [**max-ttl** *ttl-value*] [**max-path** *max-paths*] [**timeout** *timeout*] [**retry-count** *retry-count*] [**fc** *fc-name*] [**profile** *profile*] [**downstream-map-tlv** {**dsmap** | **ddmap**}]
- **config**
  - **test-oam**
    - [**no**] **ldp-treetrace**
      - **fc** *fc-name*
      - **no fc**
      - **path-discovery**
        - **interval** *minutes*
        - **no interval**
        - **max-path** *max-paths*
        - **no max-path**
        - **max-ttl** *ttl-value*
        - **no max-ttl**
        - **policy-statement** *policy-name*[...(up to 5 max)]
        - **no policy-statement**
        - **retry-count** *retry-count*
        - **no retry-count**
        - **timeout** *timeout*
        - **no timeout**
      - **path-probing**
        - **interval** *minutes*
        - **no interval**
        - **retry-count** *retry-count*
        - **no retry-count**
        - **timeout** *timeout*
        - **no timeout**
      - [**no**] **shutdown**
    - **mpls-echo-request-downstream-map** {**dsmap** | **ddmap**}
    - **no mpls-echo-request-downstream-map**

## TWAMP

### GLOBAL

- **configure**
  - **test-oam**
    - **twamp**
      - **server**
        - [**no**] **prefix** {*ip-prefix* | *mask*} [**create**]
          - [**no**] **description** *description string*
          - [**no**] **max-conn-prefix** *count*
          - [**no**] **max-sess-prefix** *count*

- [no] **shutdown**
- [no] **inactivity-timeout** *seconds*
- [no] **max-conn-server** *count*
- [no] **max-sess-server** *count*
- [no] **shutdown**

## TWAMP Light

```

configure
 — router
 — twamp-light
 — reflector [udp-port udp-port-number] [create]
 — no reflector
 — description description-string
 — no description
 — prefix ip-prefix/prefix-length [create]
 — no prefix ip-prefix/prefix-length
 — description description string
 — no description
 — [no] shutdown

configure
 — service
 — vprn
 — twamp-light
 — reflector [udp-port udp-port-number] [create]
 — no reflector
 — description description-string
 — no description
 — prefix ip-prefix/prefix-length [create]
 — no prefix ip-prefix/prefix-length
 — description description string
 — no description
 — [no] shutdown

configure
 — test-oam
 — twamp
 — twamp-light
 — inactivity-timeout seconds
 — no inactivity-timeout

monitor
 — oam-pm
 — session session-name [test-family {ethernet | ip} [session-type {proactive | on-demand}] create]
 — no session session-name
 — ip
 — destination ip-address
 — no destination
 — dest-udp-port udp-port-number
 — no dest-udp-port
 — fc fc-name
 — no fc
 — forwarding bypass-routing
 — forwarding interface interface-name
 — forwarding next-hop ip-address
 — no forwarding
 — profile {in | out}

```

- **no profile**
- **router** *router-instance*
- **router service-name** *service-name*
- **no router**
- **source** *ip-address*
- **no source**
- **source-udp-port** *udp-port-number*
- **no source-udp-port**
- **ttl** *time-to-live*
- **no ttl**
- **twamp-light** [**test-id** *test-id*] [**create**]
- **no twamp-light**
  - **interval** *milliseconds*
  - **no interval**
  - **pad-size** *octets*
  - **no pad-size**
  - [**no**] **shutdown**
  - **test-duration** *seconds*
  - **no test-duration**

## OAM Commands

### SDP Diagnostics

Note: SDP diagnostics commands are not applicable for 7210 SAS devices configured in Access uplink mode.

#### GLOBAL

##### — oam

— **sdp-mtu** *orig-sdp-id* **size-inc** *start-octets end-octets* [**step** *step-size*] [**timeout** *timeout*] [**interval** *interval*]

— **sdp-ping** *orig-sdp-id* [**resp-sdp** *resp-sdp-id*] [**fc** *fc-name*] [**timeout** *seconds*] [**interval** *seconds*] [**size** *octets*] [**count** *send-count*] [**interval** *<interval>*]

## Common Service Diagnostics

**Note:** Only DNS command is supported by 7210 SAS devices configured in access uplink mode.

## GLOBAL

## — oam

- **svc-ping** *{ip-addr}* **service** *service-id* [**local-sdp**] [**remote-sdp**]
  - **dns** **target-addr** *dns-name* **name-server** *ip-address* [**source** *ip-address*] [**count** *send-count*] [**time-out** *timeout*] [**interval** *interval*]
  - **vprn-ping** *service-id* **service** *svc-name* **source** *ip-address* **destination** *ip-address* [**fc** *fc-name*] [**size** *size*] [**tfl** *vc-label-ttl*] [**return-control**] [**interval** *interval*] [**count** *send-count*] [**timeout** *timeout*]
  - **vprn-trace** *service-id* **source** *src-ip* **destination** *ip-address* [**fc** *fc-name*] [**size** *size*] [**min-ttl** *vc-label-ttl*] [**max-ttl** *vc-label-ttl*] [**return-control**] [**probe-count** *sendcount*] [**interval** *interval*] [**time-out** *timeout*]
- 

## VLL Diagnostics

**Note:** VLL diagnostics commands are not applicable for 7210 SAS devices configured in access uplink mode.

## GLOBAL

## — oam

- **vccv-ping** *sdp-id:vc-id* [**src-ip-address** *ip-addr* **dst-ip-address** *ip-addr* **pw-id** *pw-id*] [**reply-mode** *{ip-routed | control-channel}*] [**fc** *fc-name*] [**size** *octets*] [**send-count** *send-count*] [**timeout** *timeout*] [**interval** *interval*] [**tfl** *vc-label-ttl*]
- **vccv-ping** **spoke-sdp** *spoke-sdp-id* [**reply-mode** *ip-routed| control-channel*] [**src-ip-address** *ip-addr* **dst-ip-address** *ip-addr*]
- **vccv-ping** **static** *sdp-id:vc-id* [**assoc-channel** *ipv4|non-ip*] [**dest-global-id** *global-id* **dest-node-id** *node-id*] [**src-ip-address** *ip-addr*]
- **vccv-ping** **saii-type2** *global-id:prefix:ac-id* **taii-type2** *global-id:prefix:ac-id* [**reply-mode** *ip-routed|control-channel*] [**src-ip-address** *ip-addr* **dst-ip-address** *ip-addr*]
- **vccv-ping** **spoke-sdp-fec** *spoke-sdp-fec-id* [**reply-mode** *ip-routed| control-channel*] [**saii-type2** *global-id:prefix:ac-id* **taii-type2** *global-id:prefix:ac-id*] [**src-ip-address** *ip-addr* **dst-ip-address** *ip-addr*]
- **options common to all vccv-ping cases:** [**count** *send-count*] [**fc** *fc-name*] [**profile** *in|out*] [**interval** *interval*] [**size** *octets*] [**timeout** *timeout*] [**tfl** *vc-label-ttl*]
- **vccv-trace** *sdp-id:vc-id* [**fc** *fc-name*] [**profile** *{in | out}*] [**size** *octets*] [**reply-mode** *ip-routed|control-channel*] [**probe-count** *probes-count*] [**timeout** *timeout*] [**interval** *interval*] [**min-ttl** *min-vc-label-ttl*] [**max-ttl** *max-vc-label-ttl*] [**max-fail** *no-response-count*] [**detail**]
- **vccv-trace** **static** *sdp-id:vc-id* [**assoc-channel** *ipv4|non-ip*] [**src-ip-address** *ipv4-address*]
- **vccv-trace** **saii-type2** *global-id:prefix:ac-id* **taii-type2** *global-id:prefix:ac-id* [**reply-mode** *ip-routed|control-channel*]
- **vccv-trace** **spoke-sdp-fec** *spoke-sdp-fec-id* [**reply-mode** *ip-routed| control-channel*] [**saii-type2** *global-id:prefix:ac-id* **taii-type2** *global-id:prefix:ac-id*]
- **options common to all vccv-trace cases:** [**detail**] [**fc** *fc-name*] [**profile** *in|out*] [**interval** *interval-value*] [**max-fail** *no-response-count*] [**max-ttl** *max-vc-label-ttl*] [**min-ttl** *min-vc-label-ttl*] [**probe-count** *probe-count*] [**size** *octets*] [**timeout** *timeout-value*]

---

## VPLS MAC Diagnostics

**Note:** VPLS diagnostics commands are not applicable for 7210 SAS devices configured in Access uplink mode.

### GLOBAL

- **oam**
    - **cpe-ping** *service service-id destination ip-address source ip-address* [*source-mac ieee-address*] [*fc fc-name*] [*ttl vc-label-ttl*] [*count send-count*] [*send-control*] [*return-control*] [*interval interval*]
    - **mac-ping** *service service-id destination dst-ieee-address* [*source src-ieee-address*] [*fc fc-name*] [*size octets*] [*fc fc-name*] [*ttl vc-label-ttl*] [*send-count send-count*] [*return-control*] [*interval interval*] [*timeout timeout*]
    - **mac-populate** *service-id mac ieee-address* [*flood*] [*age seconds*] [*force*] [*target-sap sap-id*]
    - **mac-purge** *service-id target ieee-address* [*flood*] [*register*]
    - **mac-trace** *service service-id destination ieee-address* [*source ieee-address*] [*fc fc-name*] [*size octets*] [*min-ttl vc-label-ttl*] [*max-ttl vc-label-ttl*] [*probe-count send-count*] [*return-control*] [*interval interval*] [*timeout timeout*]
- 

## Ethernet in the First Mile (EFM) Commands

### GLOBAL

- **oam**
    - **efm** *port-id local-loopback* {*start* | *stop*}
    - **efm** *port-id remote-loopback* {*start* | *stop*}
- 

## ETH-CFM OAM Commands

### oam

- **eth-cfm eth-test** *mac-address mep mep-id domain md-index association ma-index* [*priority priority*] [*data-length data-length*]
  - **eth-cfm linktrace** *mac-address mep mep-id domain md-index association ma-index* [*ttl ttl-value*]
  - **eth-cfm loopback** *mac-address mep mep-id domain md-index association ma-index* [*send-count send-count*] [*size data-size*] [*priority priority*]
  - **eth-cfm one-way-delay-test** *mac-address mep mep-id domain md-index association ma-index* [*priority priority*]
  - **eth-cfm two-way-delay-test** *mac-address mep mep-id domain md-index association ma-index* [*priority priority*]
  - **eth-cfm two-way-slm-test** *mac-address mep mep-id domain md-index association ma-index* [*fc {fc-name}*] [*profile {in|out}*] [*count send-count*] [*size data-size*] [*timeout timeout*] [*interval interval*]
-



## Y.1564 Testhead OAM commands

- ```

config
  — test-oam
    — testhead-profile profile-id create
      — [no] acceptance-criteria acceptance-criteria-id create
        — [no] cir-threshold threshold
        — [no] jitter-rising-threshold threshold
        — [no] jitter-rising-threshold-in threshold
        — [no] jitter-rising-threshold-out threshold
        — [no] latency-rising-threshold threshold
        — [no] latency-rising-threshold-in threshold
        — [no] latency-rising-threshold-out threshold
        — [no] loss-rising-threshold threshold
        — [no] loss-rising-threshold-in threshold
        — [no] loss-rising-threshold-out threshold
        — [no] pir-threshold threshold
      — [no] description description-string
      — dot1p in-profile dot1p-value out-profile dot1p-value
      — no dot1p
      — no frame-payload payload-id [payload-type [l2|tcp-ipv4|udp-ipv4|ipv4]] create
      — no frame-payload payload
        — [no] data-pattern data-pattern
        — [no] description description-string
        — [no] dscp dscp-name
        — [no] dst-ip ipv4 ipv4-address
        — [no] dst-mac ieee-address [ieee-address-mask]
        — [no] dst-port dst-port-number
        — [no] ethertype 0x0600..0xffff
        — [no] ip-proto ip-protocol-number
        — [no] ip-tos type-of-service
        — [no] ip-ttl ttl-value
        — [no] src-ip ipv4 ipv4-address
        — [no] src-mac ieee-address [ieee-address-mask]
        — [no] src-port src-port-number
        — [no] vlan-tag-1 vlan-id vlan-id [tpid tpid] [dot1p dot1p-value]
        — [no] vlan-tag-2 vlan-id vlan-id [tpid tpid] [dot1p dot1p-value]
      — [no] frame-size frame-size
      — [no] rate cir cir-rate-in-kbps [cir-adaptation-rule adaptation-rule] [pir cir-rate-i n-kbp]
      — [no] test-completion-trap-enable
      — [no] test-duration [hours hours|minutes minutes|seconds seconds]
      — [no] test-duration

```
-

OAM Commands

Y.1564 Testhead OAM Commands

- oam**
- **testhead** *test-name* **owner** *owner-name* **testhead-profile** *profile-id* [**frame-payload** *frame-payload-id*] **sap** *sap-id* [**fc** *fc-name*] [**acceptance-criteria** *acceptance-criteria-id*] [**color-aware** **enable**|**disable**]
 - **testhead** *test-name* **owner** *owner-name* **stop**

Show commands

```
show
  — test-oam
     — testhead-profile profile-id
show
  — testhead [test-name owner test-owner] [detail]
```

OAM Commands

Clear commands

- clear**
- **test-oam**
 - **twamp server**
- **testhead result** [*test-name*] **owner** [*test-owner*]
- **testhead testhead-profile** *profile-id*

OAM Performance Monitoring and Binning Commands

```

config
  — oam-pm
    — bin-group bin-group-number [fd-bin-count fd-bin-count fdr-bin-count fdr-bin-count ifdv-bin-
      count ifdv-bin-count create]
    — no bin-group bin-group-number
      — bin-type {fd | fdr | ifdv}
        — bin bin-number
          — lower-bound microseconds
          — no lower-bound
        — description description-string
        — no description
        — [no] shutdown
    — session session-name [test-family {ethernet | ip} [session-type {proactive | on-demand}] create]
    — no session session-name
      — bin-group bin-group-number
      — no bin-group
      — description description-string
      — no description
      — ethernet
        — dest-mac ieee-address
        — no dest-mac
        — dmm [test-id test-id] [create]
        — no dmm
        — data-tlv-size octets
        — no data-tlv-size
        — interval milliseconds
        — no interval
        — [no] shutdown
        — test-duration seconds
        — no test-duration
        — priority priority
        — no priority
        — slm [test-id test-id] [create]
        — data-tlv-size octets
        — no data-tlv-size
        — flr-threshold percentage
        — no flr-threshold
        — [no] shutdown
        — test-duration seconds
        — no test-duration
        — timing frames-per-delta-t frames consec-delta-t deltas interval
          milliseconds chli-threshold threshold
        — no timing
        — source mep mep-id domain md-index association ma-index
        — no source
      — meas-interval {5-mins | 15-mins | 1-hour | 1-day} [create]
        — accounting-policy acct-policy-id
        — no accounting-policy
        — boundary-type {clock-aligned | test-relative}
  
```

OAM Performance Monitoring and Binning Commands

- **no boundary-type**
- **clock-offset** *seconds*
- **no clock-offset**
- **intervals-stored** *intervals*
- **no intervals-stored**

SAA Commands

Note: The following commands are not supported by 7210 SAS devices configured in Access uplink mode:

- cpe-ping
- lsp-ping
- lsp-trace
- mac-ping
- mac-trace
- sdp-ping
- vccv-ping
- vccv-trace
- vprn-ping
- vprn-trace

config

- **saa**
 - [no] **test** *test-name* [owner *test-owner*]
 - **accounting-policy** *acct-policy-id*
 - **no accounting-policy**
 - [no] **continuous**
 - **description** *description-string*
 - **no description**
 - [no] **jitter-event** **rising-threshold** *threshold* [**falling-threshold** *threshold*] [*direction*]
 - [no] **latency-event** **rising-threshold** *threshold* [**falling-threshold** *threshold*] [*direction*]
 - [no] **loss-event** **rising-threshold** *threshold* [**falling-threshold** *threshold*] [*direction*]
 - **probe-history** {keep|drop|auto}
 - [no] **shutdown**
 - **trap-gen**
 - [no] **probe-fail-enable**
 - [no] **probe-fail-threshold** *0..15*
 - [no] **test-completion-enable**
 - [no] **test-fail-enable**
 - [no] **test-fail-threshold** *0..15*
 - [no] **type**

- **cpe-ping** *service service-id destination ip-address source ip-address* [**source-mac** *ieee-address*] [**fc** *fc-name*] [**ttl** *vc-label-ttl*] [**send-count** *send-count*] [**return-control**] [**interval** *interval*]
- **dns** *target-addr dns-name name-server ip-address* [**source** *ip-address*] [**send-count** *send-count*] [**timeout** *timeout*] [**interval** *interval*]
- **eth-cfm-linktrace** *mac-address mep mep-id domain md-index association ma-index* [**ttl** *ttl-value*] [**fc** *{fc-name}*] [**count** *send-count*] [**timeout** *timeout*] [**interval** *interval*] [**record-type** *{ipv4-a-record|ipv6-aaaa-record}*]
- **eth-cfm-loopback** *mac-address mep mep-id domain md-index association ma-index* [**size** *data-size*] [**fc** *{fc-name}*] [**count** *send-count*] [**timeout** *timeout*] [**interval** *interval*]
- **eth-cfm-two-way-delay** *mac-address mep mep-id domain md-index association ma-index* [**fc** *{fc-name}*] [**count** *send-count*] [**timeout** *timeout*] [**interval** *interval*]
- **eth-cfm-two-way-slm** *mac-address mep mep-id domain md-index association ma-index* [**fc** *fc-name*] [**count** *send-count*] [**size** *data-size*] [**timeout** *timeout*] [**interval** *interval*]
- **icmp-ping** [*ip-address | dns-name*] [**rapid** | **detail**] [**ttl** *time-to-live*] [**tos** *type-of-service*] [**size** *bytes*] [**pattern** *pattern*] [**source** *ip-address*] [**interval** *seconds*] [**{next-hop ip-address}**] [**{interface interface-name}**] [**bypass-routing**] [**count** *requests*] [**do-not-fragment**] [*router-instance | service-name service-name*] [**timeout** *timeout*] [**fc** *{fc-name}*]
- **icmp-trace** [*ip-address | dns-name*] [**ttl** *time-to-live*] [**wait** *milli-seconds*] [**source** *ip-address*] [**tos** *type-of-service*] [*router-instance | service-name service-name*]
- **lsp-ping** *bgp-label-prefix ip-prefix/mask* [**src-ip-address** *ip-address*] [**fc** *fc-name*] [**size** *octets*] [**ttl** *label-ttl*] [**send-count** *send-count*] [**timeout** *timeout*] [**interval** *interval*] [**path-destination** *ip-address*] [**interface** *if-name* | **next-hop** *ip-address*]
- **lsp-ping** *{lsp-name [path path-name]}* [**prefix** *ip-prefix/mask*] [**src-ip-address** *ip-address*] [**size** *octets*] [**ttl** *label-ttl*] [**timeout** *timeout*] [**interval** *interval*] [**fc** *fc-name*] [**send-count** *send-count*] [*lsp-name [path path-name]*] [**fc** *fc-name*] [**size** *octets*] [**ttl** *label-ttl*] [**send-count** *send-count*] [**timeout** *timeout*] [**interval** *interval*]
- **lsp-trace** *bgp-label-prefix ip-prefix/mask* [**src-ip-address** *ip-address*] [**fc** *fc-name*] [**max-fail** *no-response-count*] [**probe-count** *probes-per-hop*] [**size** *octets*] [**min-ttl** *min-label-ttl*] [**max-ttl** *max-label-ttl*] [**timeout** *timeout*] [**interval** *interval*] [**path-destination** *ip-address*] [**interface** *if-name* | **next-hop** *ip-address*] [**downstream-map-tlv** *dsmap* | *ddmap* | *none*] [**detail**]
- **lsp-trace** *{lsp-name [path path-name]}* [**fc** *fc-name*] [**max-fail** *no-response-count*] [**probe-count** *probes-per-hop*] [**size** *octets*] [**min-ttl** *min-label-ttl*] [**max-ttl** *max-label-ttl*] [**src-ip-address** *ip-address*] [**timeout** *timeout*] [**interval** *interval*]
- **mac-ping** *service service-id destination dst ieee-address* [**source** *src-ieee-address*] [**fc** *fc-name*] [**size** *octets*] [**ttl** *vc-label-ttl*] [**send-count** *send-count*] [**send-control**] [**return-control**] [**interval** *interval*] [**timeout** *timeout*]
- **mac-trace** *service service-id destination ieee-address* [**source** *src-ieee-address*] [**fc** *fc-name*] [**size** *octets*] [**min-ttl** *min-label-ttl*] [**max-ttl** *max-label-ttl*] [**send-count** *send-count*] [**send-control**] [**return-control**] [**interval** *interval*] [**timeout** *timeout*]

- **sdp-ping** *orig-sdp-id* [**resp-sdp** *resp-sdp-id*] [**fc** *fc-name*] [**size** *octets*] [**send-count** *send-count*][**timeout** *seconds*] [**interval** *seconds*]
- **vccv-ping** *sdp-id:vc-id* [**src-ip-address** *ip-addr* **dst-ip-address** *ip-addr* **pw-id** *pw-id*][**reply-mode** {**ip-routed** | **control-channel**}][**fc** *fc-name*] [**size** *octets*] [**send-count** *send-count*][**timeout** *timeout*] [**interval** *interval*][**ttl** *vc-label-ttl*]
- **vccv-trace** *sdp-id:vc-id* [**size** *octets*][**min-ttl** *vc-label-ttl*] [**max-ttl** *vc-label-ttl*][**max-fail** *no-response-count*][**probe-count** *probe-count*][**reply-mode** *ip-routed*]**control-channel**][**timeout** *timeout-value*][**interval** *interval-value*][**fc** *fc-name*] [**detail**]
- **vprn-ping** *service-id* **service** *svc-name* [**src-ip-address** *ip-addr* **dst-ip-address** *ip-addr*] [**fc** *fc-name*] [**profile** *in* | *out*] [**size** *size*] [**ttl** *vc-label-ttl*] [**count** *send-count*] [**return-control**] [**timeout** *timeout*] [**interval** *seconds*]
- **vprn-trace** *service-id* **source** *src-ip* **destination** *dst-ip* [**fc** *fc-name*] [**profile** *in* | *out*] [**size** *size*] [**min-ttl** *vc-label-ttl*] [**max-ttl** *vc-label-ttl*] [**count** *send-count*] [**return-control**] [**timeout** *timeout*] [**interval** *interval*]

Show Commands

show

- **eth-cfm**
 - **association** [*ma-index*] [**detail**]
 - **cfm-stack-table** [**port** [*port-id*] [**vlan** *qtag*[*.qtag*]]] **sdp** *sdp-id*[:*vc-id*][**level** 0..7] [**direction** *up* | *down*]
 - **domain** [*md-index*] [**association** *ma-index* | **all-associations**] [**detail**]
 - **mep** *mep-id* **domain** *md-index* **association** *ma-index* [**loopback**] [**linktrace**]
 - **mep** *mep-id* **domain** *md-index* **association** *ma-index* [**remote-mepid** *mep-id* | **all-remote-mepids**]
 - **mep** *mep-id* **domain** *md-index* **association** *ma-index* **eth-test-results** [**remote-peer** *mac-address*]
 - **mep** *mep-id* **domain** *md-index* **association** *ma-index* **one-way-delay-test** [**remote-peer** *mac-address*]
 - **mep** *mep-id* **domain** *md-index* **association** *ma-index* **two-way-delay-test** [**remote-peer** *mac-address*]
 - **mep** *mep-id* **domain** *md-index* **association** *ma-index* **two-way-slm-test** [**remote-peer** *mac-address*]
 - **mip**
 - **system-config**
 - **router**
 - **twamp-light**
 - **saa** [*test-name*] [**owner** *test-owner*]
 - **service** *service-id*
 - **twamp-light**
 - **test-oam**
 - **ldp-treetrace** [**prefix** *ip-prefix/mask*] [**detail**]
 - **twamp**
 - **twamp-light**
 - **reflectors**
 - **twamp server**
 - **server** *all*
 - **server** **prefix** *ip-prefix/mask*
- show
- **oam-pm**

SAA Commands

- **bin-group** [*bin-group-number*]
- **bin-group-using** [**bin-group** *bin-group-number*]
- **session** *session-name* [{**all** | **base** | **bin-group** | **event-mon** | **meas-interval**}]
- **sessions** [**test-family** {**ethernet** | **ip**}] [**event-mon**]
- **statistics**
 - **session** *session-name*
 - **dmm**
 - **meas-interval** **raw** [{**all** | **bins** | **summary**}]
 - **meas-interval** {**5-mins** | **15-mins** | **1-hour** | **1-day**} *interval-number* *interval-number* [{**all** | **bins** | **summary**}]
 - **slm**
 - **meas-interval** **raw**
 - **meas-interval** {**5-mins** | **15-mins** | **1-hour** | **1-day**} *interval-number* *interval-number*
 - **twamp-light**
 - **meas-interval** **raw** **delay** [{**all** | **bins** | **summary**}]
 - **meas-interval** **raw** [**loss**]
 - **meas-interval** {**5-mins** | **15-mins** | **1-hour** | **1-day**} *interval-number* *interval-number* **delay** [{**all** | **bins** | **summary**}]
 - **meas-interval** {**5-mins** | **15-mins** | **1-hour** | **1-day**} *interval-number* *interval-number* **loss**

Monitor Commands

- monitor**
 - **oam-pm**
 - **session** *session-name* [{**dmm** | **slm** | **twamp-light**}]

Clear Commands

- clear**
 - **saa** [*test-name* [**owner** *test-owner*]]
 - **test-oam**
 - **twamp server**
 - **server**

Clear Commands

- clear**
 - **oam-pm**
 - **session** *session-name* {**dmm** | **slm** | **twamp-light**}

OAM and SAA Command Hierarchies

Operational Commands

shutdown

Syntax [no] shutdown

Context config>saa>test

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

In order to modify an existing test it must first be shut down. When a test is created it will be in shutdown mode until a **no shutdown** command is executed.

A **shutdown** can only be performed if a test is not executing at the time the command is entered.

Use the **no** form of the command to set the state of the test to operational.

shutdown

Syntax [no] shutdown

Context config>test-oam>ldp-treetrace
config>test-oam>twamp>server
config>test-oam>twamp>server>prefix

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command suspends the background process running the LDP ECMP OAM tree discovery and path probing features. The configuration is not deleted.

Use the **no** form of the command to enable the background process.

dns

Syntax dns target-addr dns-name name-server ip-address [source ip-address] [count send-count] [timeout timeout] [interval interval] [record-type {ipv4-a-record|ipv6-aaaa-record}]

Context oam

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command performs DNS name resolution. If `ipv4-a-record` is specified, dns-names are queried for A-records only.

Parameters **send-count** *send-count* — The number of messages to send, expressed as a decimal integer. The **send-count** parameter is used to override the default number of message requests sent. Each message request must either timeout or receive a reply before the next message request is sent. The message **interval** value must be expired before the next message request is sent.

Default 1

Values 1 — 100

ip-address — The IP or IPv6 address of the primary DNS server.

Values
ipv4-address - a.b.c.d
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

timeout *timeout* — The **timeout** parameter in seconds, expressed as a decimal integer. This value is used to override the default **timeout** value and is the amount of time that the router will wait for a message reply after sending the message request. Upon the expiration of message timeout, the requesting router assumes that the message response will not be received. Any response received after the request times out will be silently discarded.

Default 5

Values 1 — 120

interval *interval* — The **interval** parameter in seconds, expressed as a decimal integer. This parameter is used to override the default request message send interval and defines the minimum amount of time that must expire before the next message request is sent.

If the **interval** is set to 1 second, and the **timeout** value is set to 10 seconds, then the maximum time between message requests is 10 seconds and the minimum is 1 second. This depends upon the receipt of a message reply corresponding to the outstanding message request.

Default 1

Values 1 — 10

record-type — Specifies a record type.

Values
`ipv4-a-record` — A record specific mapping a host name to an IPv4 address.
`ipv6-aaaa-record` — A record specific to the Internet class that stores a single IPv6 address.

ping

Syntax ping [*ip-address* | *dns-name*] [**rapid** | **detail**] [**ttl** *time-to-live*] [**tos** *type-of-service*] [**size** *bytes*] [**pattern** *pattern*] [**source** *ip-address* | *dns-name*] [**interval** *seconds*] [{**next-hop** *ip-address*} | {**interface** *interface-name*} | **bypass-routing**] [**count** *requests*] [**do-not-fragment**] [**router** *router-instance* | **service-name** *service-name*] [**timeout** *timeout*]

Context <GLOBAL>

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command verifies the reachability of a remote host.

Parameters *ip-address* — The far-end IP address to which to send the **sve-ping** request message in dotted decimal notation.

Note: IPv6 is supported only for "Management" instance of the router.

Values

| | |
|---------------|-------------------------------------|
| ipv4-address: | a.b.c.d |
| 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 |

dns-name — The DNS name of the far-end device to which to send the **sve-ping** request message, expressed as a character string.

rapid — Packets will be generated as fast as possible instead of the default 1 per second.

detail — Displays detailed information.

ttl *time-to-live* — The TTL value for the MPLS label, expressed as a decimal integer.

Values 1 — 128

tos *type-of-service* — Specifies the service type.

Values 0 — 255

size *bytes* — The request packet size in bytes, expressed as a decimal integer.

Values 0 — 16384

pattern *pattern* — The data portion in a ping packet will be filled with the pattern value specified. If not specified, position info will be filled instead.

Values 0 — 65535

source *ip-address* — Specifies the IP address to be used.

Note: IPv6 is supported only for "Management" instance of the router.

Values

| | |
|---------------|-------------------------------------|
| ipv4-address: | a.b.c.d |
| 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

router *router-instance* — Specifies the router name or service ID.

Values *router-name:* Base , management

Default Base

service-name *service-name* - Specifies the service name as an integer or string.

bypass-routing — Specifies whether to send the ping request to a host on a directly attached network bypassing the routing table.

interface *interface-name* — Specifies the name of an IP interface. The name must already exist in the **config>router>interface** context.

next-hop *ip-address* — Only displays static routes with the specified next hop IP address.

Note: IPv6 is supported only for "Management" instance of the router.

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]

interval *seconds* — The interval parameter in seconds, expressed as a decimal integer. This parameter is used to override the default request message send interval and defines the minimum amount of time that must expire before the next message request is sent.

If the interval is set to 1 second, and the timeout value is set to 10 seconds, then the maximum time between message requests is 10 seconds and the minimum is 1 second. This depends upon the receipt of a message reply corresponding to the outstanding message request.

Values 1 — 10000

count *requests* — Specifies the number of times to perform an OAM ping probe operation. Each OAM echo message request must either timeout or receive a reply before the next message request is sent.

Values 1 — 100000

Default 5

do-not-fragment — Sets the DF (Do Not Fragment) bit in the ICMP ping packet.

fc-name — The forwarding class of the MPLS echo request encapsulation.

Values be|l2|af|l1|h2|ef|h1|nc

Default nc

timeout *seconds* — Overrides the default **timeout** value and is the amount of time that the router will wait for a message reply after sending the message request. Upon the expiration of message timeout, the requesting router assumes that the message response will not be received. A 'request timeout' message is displayed by the CLI for each message request sent that expires. Any response received after the request times out will be silently discarded.

Default 5
Values 1 — 10

traceroute

Syntax **traceroute** [*ip-address* [*dns-name*] [**ttl** *ttl*] [**wait** *milli-seconds*] [**no-dns**] [**source** *ip-address*] [*tos type-of-service*] [**router** *router-instance* | **service-name** *service-name*]

Context Global

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

The TCP/IP traceroute utility determines the route to a destination address. DNS lookups of the responding hosts is enabled by default.

```
*A:ALA-1# traceroute 192.168.xx.xx4
traceroute to 192.168.xx.xx4, 30 hops max, 40 byte packets
 1 192.168.xx.xx4 0.000 ms 0.000 ms 0.000 ms
*A:ALA-1#
```

Parameters *ip-address* — The far-end IP address to which to send the traceroute request message in dotted decimal notation.

Note: IPv6 is supported only for "Management" instance of the router.

Values

| | |
|---------------|-------------------------------------|
| ipv4-address: | a.b.c.d |
| 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 |

dns-name — The DNS name of the far-end device to which to send the traceroute request message, expressed as a character string.

ttl *ttl* — The maximum Time-To-Live (TTL) value to include in the traceroute request, expressed as a decimal integer.

Values 1 — 255

wait *milliseconds* — The time in milliseconds to wait for a response to a probe, expressed as a decimal integer.

Default 5000

Values 10 — 60000

no-dns — When the **no-dns** keyword is specified, DNS lookups of the responding hosts will not be performed, only the IP addresses will be printed.

Default DNS lookups are performed

source *ip-address* — The source IP address to use as the source of the probe packets in dotted decimal notation. If the IP address is not one of the device's interfaces, an error is returned.

tos *type-of-service* — The type-of-service (TOS) bits in the IP header of the probe packets, expressed as a decimal integer.

Values 0 — 255

router *router-name* — Specify the alphanumeric character string up to 32 characters.

Values Base, Management

service-name *service-name* - Specifies the service name as an integer or string.

lsp-ping

Syntax **lsp-ping** *lsp-name* [**path** *path-name*]

Context **lsp-ping bgp-label prefix** *ip-prefix/mask* [**path-destination** *ip-address* [**interface** *if-name* | **next-hop** *ip-address*]]
lsp-ping prefix *ip-prefix/mask* [**path-destination** *ip-address* [**interface** *if-name* | **next-hop** *ip-address*]] [**src-ip-address** *ip-address*]
lsp-ping {{**lsp-name** [**path** *path-name*]}{**prefix** *ip-prefix/mask*}} [**src-ip-address** *ip-address*] [**size** *octets*] [**ttl** *label-ttl*] [**timeout** *timeout*] [**interval** *interval*] [**fc** *fc-name*] [**send-count** *send-count*] {**lsp-name** [**path** *path-name*]} [**fc** *fc-name*] [**size** *octets*][**ttl** *label-ttl*] [**send-count** *send-count*] [**timeout** *timeout*] [**interval** *interval*]
lsp-ping static *lsp-name* [**assoc-channel** *ipv4|none|non-ip*] [**force**] [**dest-global-id** *global-id*] [**dest-node-id** *node-id*] [**path-type** *active | working | protect*]

NOTE: Options common to all **lsp-ping** cases: [**detail**] [**fc** *fc-name*] [**interval** *interval*] [**send-count** *send-count*] [**size** *octets*] [**src-ip-address** *ip-address*] [**timeout** *timeout*] [**ttl** *label-ttl*]

Context oam
config>saa>test>type

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command performs in-band LSP connectivity tests.

The **lsp-ping** command performs an LSP ping using the protocol and data structures defined in the RFC 4379, Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures.

The LSP ping operation is modeled after the IP ping utility which uses ICMP echo request and reply packets to determine IP connectivity.

In an LSP ping, the originating device creates an MPLS echo request packet for the LSP and path to be tested. The MPLS echo request packet is sent through the data plane and awaits an MPLS echo reply packet from the device terminating the LSP. The status of the LSP is displayed when the MPLS echo reply packet is received.

This command, when used with the **static** option, performs in-band on-demand LSP connectivity verification tests for static MPLS-TP LSPs. For other LSP types, the **static** option should be excluded and these are described elsewhere in this user guide.

The **lsp-ping static** command performs an LSP ping using the protocol and data structures defined in the RFC 4379, Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures, as extended by RFC 6426, MPLS On-Demand Connectivity Verification and Route Tracing.

In MPLS-TP, the echo request and echo reply messages are always sent in-band over the LSP, either in a G-ACh channel or encapsulated as an IP packet below the LSP label.

The timestamp format to be sent, and to be expected when received in a PDU, is as configured by the **config>test-oam>mpls-time-stamp-format** command. If RFC 4379 is selected, then the timestamp is in seconds and microseconds since 1900, otherwise it is in seconds and microseconds since 1970.

Parameters

lsp-name — Name that identifies an LSP to ping. The LSP name can be up to 32 characters long.

dest-global-id *global-id* — The MPLS-TP global ID for the far end node of the LSP under test. If this is not entered, then the dest-global-id is taken from the LSP context.

dest-node-id *node-id* — The MPLS-TP global ID for the far end node of the LSP under test. If this is not entered, then the dest-global-id is taken from the LSP context.

control-channel {**none** | **non-ip**} — The encapsulation format to use for the LSP Ping echo request and echo reply packet.

Values none — IP encapsulation in an MPLS labeled packet

Values non-ip — MPLS-TP encapsulation without UDP/IP headers, in an MPLS-TP G-ACh on the LSP using channel type 0x025.

Default non-ip

force — Allows LSP Ping to test a path that is operationally down, including cases where MPLS-TP BFD CC/V is enabled and has taken a path down. This parameter is only allowed in the OAM context; it is not allowed for a test configured as a part of an SAA.

Default disabled

path-type {**active** | **working** | **protect**} — The LSP path to test.

Default active

Values active — The currently active path. If MPLS-TP linear protection is configured on the LSP, then this is the path that is selected by by MPLS-TP PSC protocol for sending user plane traffic. If MPLS-TP linear protection is not configured, then this will be the working path.

Values working — The working path of the MPLS-TP LSP.

Values protect — The protect path of the MPLS-TP LSP.

path *path-name* — The LSP path name along which to send the LSP ping request.

Values Any path name associated with the LSP.

Default The active LSP path.

bgp-label-prefix *ip-prefix/mask* — Specifies the address prefix and subnet mask of the target BGP IPv4 label route.

src-ip-address *ip-addr* — Specifies the source IP address. This option is used when an OAM packet must be generated from a different address than the node's system interface address. An example is when the OAM packet is sent over an LDP LSP and the LDP LSR-ID of the corresponding LDP session to the next-hop is set to an address other than the system interface address.

Values ipv4-address: a.b.c.d

fc *fc-name* — The fc parameter is used to indicate the forwarding class of the MPLS echo request packet.

When an MPLS echo request packet is generated in CPM and is forwarded to the outgoing interface, the packet is queued in the egress network queue corresponding to the specified fc parameter value. The marking of the packet's EXP is dictated by the LSP-EXP mappings on the outgoing interface.

When the MPLS echo request packet is received on the responding node, The fc parameter value is dictated by the LSP-EXP mappings of the incoming interface.

When an MPLS echo reply packet is generated in CPM and is forwarded to the outgoing interface, the packet is queued in the egress network queue corresponding to the fc parameter value determined by the classification of the echo request packet, which is being replied to, at the incoming interface. The marking of the packet's EXP is dictated by the LSP-EXP mappings on the outgoing interface. The TOS byte is not modified. The following table summarizes this behavior:

Table 18: Request Packet and Behavior

| | |
|-------------------------------------|--|
| cpm (sender node) | echo request packet: <ul style="list-style-type: none"> • packet{tos=1, fc1 • fc1 is as entered by user in OAM command or default values • tos1 as per mapping of fc1 to IP precedence in network egress QoS policy of outgoing interface |
| outgoing interface (sender node) | echo request packet: <ul style="list-style-type: none"> • pkt queued as fc1 • ToS field=tos1 not remarked • EXP=exp1, as per mapping of fc1 to EXP in network egress QoS policy of outgoing interface |
| Incoming interface (responder node) | echo request packet: <ul style="list-style-type: none"> • packet{tos1, exp1} • exp1 mapped to fc2 as per classification in network QoS policy of incoming interface |
| cpm (responder node) | echo reply packet: <ul style="list-style-type: none"> • packet{tos=1, fc2 |

Table 18: Request Packet and Behavior (Continued)

| | |
|-------------------------------------|--|
| outgoing interface (responder node) | echo reply packet: <ul style="list-style-type: none"> • pkt queued as fc2 • ToS filed= tos1 not remarked (reply inband or out-of-band) • EXP=exp2, if reply is inband, remarked as per mapping of fc2 to EXP in network egress QoS policy of outgoing interface |
| Incoming interface (sender node) | echo reply packet: <ul style="list-style-type: none"> • packet{tos1, exp2} • exp2 mapped to fc1 as per classification in network QoS policy of incoming interface |

The LSP-EXP mappings on the receive network interface controls the mapping of the message reply back at the originating router.

Values be, l2, af, l1, h2, ef, h1, nc

Default be

src-ip-address *ip-addr* — This parameter specifies the source IP address. This parameter is used when an OAM packet must be generated from a different address than the node's system interface address. For example, when the OAM packet is sent over an LDP LSP and the LDP LSR-ID of the corresponding LDP session to the next-hop is set to an address other than the system interface address.

Values ipv4-address: a.b.c.d

size *octets* — The MPLS echo request packet size in octets, expressed as a decimal integer. The request payload is padded with zeroes to the specified size.

Values 1 — 9198

Default 1

ttl *label-ttl* — The TTL value for the MPLS label, expressed as a decimal integer.

Values 1 — 255

Default 255

send-count *send-count* — The number of messages to send, expressed as a decimal integer. The **send-count** parameter is used to override the default number of message requests sent. Each message request must either timeout or receive a reply before the next message request is sent. The message **interval** value must be expired before the next message request is sent.

Values 1 — 100

Default 1

time-out *interval* — The time-out parameter in seconds, expressed as a decimal integer. This value is used to override the default timeout value and is the amount of time that the router will wait for a message

reply after sending the last probe for a particular test. Upon the expiration of timeout the test will be marked complete and no more packets will be processed for any of those request probes.

Values 1 — 10

Default 5

interval *interval* — The interval parameter in seconds, expressed as a decimal integer. This parameter is used to override the default request message send interval and defines the minimum amount of time that must expire before the next message request is sent.

path-destination *ip-address* — Specifies the IP address of the path destination from the range 127/8.

interface *interface-name* — Specifies the name of an IP interface. The name must already exist in the **config>router>interface** context.

next-hop *ip-address* — Only displays static routes with the specified next hop IP address.

ipv4-address: a.b.c.d (host bits must be 0)

prefix *ip-prefix/mask* — Specifies the address prefix and subnet mask of the target BGP IPv4 label route.

static *lsp-name* — Specifies an LSP ping route using the RFC 6426, *MPLS On-Demand Connectivity Verification and Route Tracing*, Target FEC Stack code point Static LSP.

assoc-channel **ipv4none|non-ip** — Specifies the launched echo request's usage of the Associated Channel (ACH) mechanism, when testing an MPLS-TP LSP.

Values **ipv4** — Use the

none — Use the Associated Channel mechanism described in RFC 6426, Section 3.3.

non-ip — Do not use an Associated Channel, as described in RFC 6426, Section 3.1.

dest-global-id *global-id* — Indicates the source MPLS-TP global identifier of the replying node. The value is copied from the reply's RFC 6426 Source Identifier TLV.

Values 0 — 4294967295

Default 0

dest-node-id *node-id* — Specifies the target MPLS-TP Node Identifier.

Values a.b.c.d | 1 — 4294967295>

Default 0

path-type **active | working | protect** — Specifies the type of an MPLS TP path.

Values **active** - test the currently-active path of the MPLS-TP LSP

working - test the primary path of the MPLS-TP LSP

protect - test the secondary path of the MPLS-TP LSP

Sample Output

```
A:DUTA# oam lsp-ping prefix 4.4.4.4/32 detail
LSP-PING 4.4.4.4/32: 80 bytes MPLS payload
Seq=1, send from intf dut1_to_dut3, reply from 4.4.4.4
      udp-data-len=32 ttl=255 rtt=5.23ms rc=3 (EgressRtr)
```

```

---- LSP 4.4.4.4/32 PING Statistics ----
1 packets sent, 1 packets received, 0.00% packet loss
round-trip min = 5.23ms, avg = 5.23ms, max = 5.23ms, stddev = 0.000ms

=====
LDP LSR ID: 1.1.1.1
=====
Legend: U - Label In Use, N - Label Not In Use, W - Label Withdrawn
       WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
=====
LDP Prefix Bindings
=====
Prefix          IngLbl      EgrLbl      EgrIntf/      EgrNextHop
  Peer
-----
4.4.4.4/32      131069N     131067      1/1/1         1.3.1.2
  3.3.3.3
4.4.4.4/32      131069U     131064      --            --
  6.6.6.6
-----
No. of Prefix Bindings: 2
=====
A:DUTA#

```

lsp-trace

Syntax

```

lsp-trace lsp-name [path path-name]
lsp-trace bgp-label prefix ip-prefix/mask [path-destination ip-address [interface if-name | next-hop ip-address]]
lsp-trace prefix ip-prefix/mask [path-destination ip-address [interface if-name | next-hop ip-address]]
lsp-trace static lsp-name [assoc-channel ipv4|none|non-ip] [path-type active | working | protect]

```

NOTE: Options common to all **lsp-trace** cases: [**detail**] [**downstream-map-tlv** {**dsmap** | **ddmap** | **none**}] [**fc** *fc-name*] [**interval** *interval*] [**max-fail** *no-response-count*] [**max-ttl** *max-label-ttl*] [**min-ttl** *min-label-ttl*] [**probe-count** *probes-per-hop*] [**size** *octets*] [**src-ip-address** *ip-address*] [**timeout** *timeout*]

Context

```

oam
config>saa>test>type

```

Description

Platforms Supported: 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command, when used with the **static** option, performs in-band on-demand LSP traceroute tests for static MPLS-TP LSPs. For other LSP types, the **static** option should be excluded and these are described elsewhere in this user guide.

The **lsp-trace static** command performs an LSP trace using the protocol and data structures defined in the RFC 4379, Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures, as extended by RFC 6426, MPLS On-Demand Connectivity Verification and Route Tracing.

The LSP trace operation is modeled after the IP traceroute utility which uses ICMP echo request and reply packets with increasing TTL values to determine the hop-by-hop route to a destination IP.

In an LSP trace, the originating device creates an MPLS echo request packet for the LSP to be tested with increasing values of the TTL in the outermost label. The MPLS echo request packet is sent through the data plane and awaits a TTL exceeded response or the MPLS echo reply packet from the device terminating the LSP. The devices that reply to the MPLS echo request packets with the TTL exceeded and the MPLS echo reply are displayed.

The downstream mapping TLV is used in **lsp-trace** to provide a mechanism for the sender and responder nodes to exchange and validate interface and label stack information for each downstream hop in the path of the LDP FEC or an RSVP LSP, or a BGP IPv4 label route.

Two downstream mapping TLVs are supported. The original Downstream Mapping (DSMAP) TLV defined in RFC 4379 and the new Downstream Detailed Mapping (DDMAP) TLV defined in RFC 6424. More details are provided in the DDMAP TLV sub-section below.

In addition, when the responder node has multiple equal cost next-hops for an LDP FEC or a BGP label IPv4 prefix, it replies in the Downstream Mapping TLV with the downstream information for each outgoing interface which is part of the ECMP next-hop set for the prefix. The downstream mapping TLV can further be used to exercise a specific path of the ECMP set using the path-destination option.

In MPLS-TP, the echo request and echo reply messages are always sent in-band over the LSP, either in a G-ACh channel or encapsulated as an IP packet below the LSP label.

Some caveats apply when using this feature on 7210 nodes. See “LSP Diagnostics: LSP Ping and Trace” on page 82.

Parameters

lsp-name — Name that identifies an LSP to ping. The LSP name can be up to 32 characters long.

path *path-name* — The LSP path name along which to send the LSP trace request.

Values Any path name associated with the LSP.

Default The active LSP path.

control-channel {**none** | **non-ip**} — The encapsulation format to use for the MPLS echo request and echo reply packet.

Values none — IP encapsulation in an MPLS labeled packet

Values non-ip — MPLS-TP encapsulation without UDP/IP headers, in an MPLS-TP G-ACh on the LSP using channel type 0x025.

Default non-ip

size *octets* — The size in octets, expressed as a decimal integer, of the MPLS echo request packet, including the IP header but not the label stack. The request pay-load is padded with zeroes to the specified size. Note that an OAM command is not failed if the user entered a size lower than the minimum required to build the packet for the echo request message. The payload is automatically padded to meet the minimum size.

Values [104..9198]

Default 1

src-ip-address *ip-addr* — Specifies the source IP address. This option is used when an OAM packet must be generated from a different address than the node's system interface address. An example is when the OAM packet is sent over an LDP LSP and the LDP LSR-ID of the corresponding LDP session to the next-hop is set to an address other than the system interface address.

Values ipv4-address: a.b.c.d

min-ttl *min-label-ttl* — The minimum TTL value in the MPLS label for the LSP trace test, expressed as a decimal integer.

Default 1

Values 1 — 255

max-ttl *max-label-ttl* — The maximum TTL value in the MPLS label for the LDP tree-trace test, expressed as a decimal integer.

Values 1 — 255

Default 30

max-fail *no-response-count* — The maximum number of consecutive MPLS echo requests, expressed as a decimal integer that do not receive a reply before the trace operation fails for a given TTL.

Values 1 — 255

Default 5

probes-per-hop — The number of messages to send, expressed as a decimal integer. The **send-count** parameter is used to override the default number of message requests sent. Each message request must either timeout or receive a reply before the next message request is sent. The message **interval** value must be expired before the next message request is sent.

Values 1 — 10

Default 1

timeout *timeout* — The **timeout** parameter in seconds, expressed as a decimal integer. This value is used to override the default **timeout** value and is the amount of time that the router will wait for a message reply after sending the message request. Upon the expiration of message timeout, the requesting router assumes that the message response will not be received. A 'request timeout' message is displayed by the CLI for each message request sent that expires. Any response received after the request times out will be silently discarded.

Values 1 — 60

Default 3

interval *interval* — The **interval** parameter in seconds, expressed as a decimal integer. This parameter is used to override the default request message send interval and defines the minimum amount of time that must expire before the next message request is sent.

If the **interval** is set to 1 second, and the **timeout** value is set to 10 seconds, then the maximum time between message requests is 10 seconds and the minimum is 1 second. This depends upon the receipt of

a message reply corresponding to the outstanding message request.

Values 1 — 10

Default 1

fc *fc-name* — The *fc* parameter is used to indicate the forwarding class of the MPLS echo request packet.

When an MPLS echo request packet is generated in CPM and is forwarded to the outgoing interface, the packet is queued in the egress network queue corresponding to the specified *fc* parameter value. The marking of the packet's EXP is dictated by the LSP-EXP mappings on the outgoing interface.

When the MPLS echo request packet is received on the responding node, The *fc* parameter value is dictated by the LSP-EXP mappings of the incoming interface.

When an MPLS echo reply packet is generated in CPM and is forwarded to the outgoing interface, the packet is queued in the egress network queue corresponding to the *fc* parameter value determined by the classification of the echo request packet, which is being replied to, at the incoming interface. The marking of the packet's EXP is dictated by the LSP-EXP mappings on the outgoing interface. The TOS byte is not modified. The following table summarizes this behavior:

Table 19: Request Packet and Behavior

| | |
|-------------------------------------|--|
| cpm (sender node) | echo request packet: <ul style="list-style-type: none"> packet{tos=1, fc1 fc1 is as entered by user in OAM command or default values tos1 as per mapping of fc1 to IP precedence in network egress QoS policy of outgoing interface |
| outgoing interface (sender node) | echo request packet: <ul style="list-style-type: none"> pkt queued as fc1 ToS field=tos1 not remarked EXP=exp1, as per mapping of fc1 to EXP in network egress QoS policy of outgoing interface |
| Incoming interface (responder node) | echo request packet: <ul style="list-style-type: none"> packet{tos1, exp1} exp1 mapped to fc2 as per classification in network QoS policy of incoming interface |
| cpm (responder node) | echo reply packet: <ul style="list-style-type: none"> packet{tos=1, fc2 |
| outgoing interface (responder node) | echo reply packet: <ul style="list-style-type: none"> pkt queued as fc2 ToS filed= tos1 not remarked (reply inband or out-of-band) EXP=exp2, if reply is inband, remarked as per mapping of fc2 to EXP in network egress QoS policy of outgoing interface |
| Incoming interface (sender node) | echo reply packet: <ul style="list-style-type: none"> packet{tos1, exp2} exp2 mapped to fc1 as per classification in network QoS policy of incoming interface |

Values be, l2, af, l1, h2, ef, h1, nc

Default be

path-destination *ip-address* — Specifies the IP address of the path destination from the range 127/8.

interface *interface-name* — Specifies the name of an IP interface. The name must already exist in the con-fig>router>interface context.

downstream-map-tlv {dsmap|ddmap|none} — Specifies which format of the downstream mapping TLV

to use in the LSP trace packet. The DSMAP TLV is the original format in RFC 4379. The DDMAP is the new enhanced format specified in RFC 6424. The user can also choose not to include the downstream mapping TLV by entering the value none. When lsp-trace is used on a MPLS-TP LSP (static option), it can only be executed if the control-channel is set to none. In addition, the DSMAP/DDMAP TLV is only included in the echo request message if the egress interface is either a numbered IP interface, or an unnumbered IP interface. The TLV will not be included if the egress interface is of type **unnumbered-mpls-tp**.

Default Inherited from global configuration of downstream mapping TLV in option **mpls-echo-request-downstream-map {dsmap | ddmmap}**.

Sample Output

```
*A:Dut-A# oam lsp-trace prefix 10.20.1.6/32 downstream-map-tlv ddmmap path-destination
127.0.0.1 detail lsp-trace to 10.20.1.6/32: 0 hops min, 0 hops max, 152 byte packets
1 10.20.1.2 rtt=3.44ms rc=8(DSRtrMatchLabel) rsc=1
   DS 1: ipaddr=127.0.0.1 ifaddr=0 iftype=ipv4Unnumbered MRU=1500
       label[1]=131070 protocol=3(LDP)
2 10.20.1.4 rtt=4.65ms rc=8(DSRtrMatchLabel) rsc=1
   DS 1: ipaddr=127.0.0.1 ifaddr=0 iftype=ipv4Unnumbered MRU=1500
       label[1]=131071 protocol=3(LDP)
3 10.20.1.6 rtt=7.63ms rc=3(EgressRtr) rsc=1
*A:Dut-A#
```

```
*A:Dut-C# oam lsp-trace "p_1" detail
lsp-trace to p_1: 0 hops min, 0 hops max, 116 byte packets
1 10.20.1.2 rtt=3.46ms rc=8(DSRtrMatchLabel)
   DS 1: ipaddr 10.20.1.4 ifaddr 3 iftype 'ipv4Unnumbered' MRU=1500
       label=131071 proto=4(RSVP-TE)
2 10.20.1.4 rtt=3.76ms rc=8(DSRtrMatchLabel)
   DS 1: ipaddr 10.20.1.6 ifaddr 3 iftype 'ipv4Unnumbered' MRU=1500
       label=131071 proto=4(RSVP-TE)
3 10.20.1.6 rtt=5.68ms rc=3(EgressRtr)
*A:Dut-C#
```

lsp-trace over a numbered IP interface

```
A:DUTA#
A:DUTA# oam lsp-trace prefix 5.5.5.5/32 detail
lsp-trace to 5.5.5.5/32: 0 hops min, 0 hops max, 104 byte packets
1 6.6.6.6 rtt=2.45ms rc=8(DSRtrMatchLabel)
   DS 1: ipaddr=5.6.5.1 ifaddr=5.6.5.1 iftype=ipv4Numbered MRU=1564
       label=131071 proto=3(LDP)
2 5.5.5.5 rtt=4.77ms rc=3(EgressRtr)
A:DUTA#
```

lsp-trace over an unnumbered IP interface

```
*A:Dut-A# oam lsp-trace prefix 10.20.1.6/32 downstream-map-tlv ddmmap path-destination
127.0.0.1 detail lsp-trace to 10.20.1.6/32: 0 hops min, 0 hops max, 152 byte packets
```

```

1 10.20.1.2 rtt=3.44ms rc=8(DSRtrMatchLabel) rsc=1
   DS 1: ipaddr=127.0.0.1 ifaddr=0 iftype=ipv4Unnumbered MRU=1500
       label[1]=131070 protocol=3(LDP)
2 10.20.1.4 rtt=4.65ms rc=8(DSRtrMatchLabel) rsc=1
   DS 1: ipaddr=127.0.0.1 ifaddr=0 iftype=ipv4Unnumbered MRU=1500
       label[1]=131071 protocol=3(LDP)
3 10.20.1.6 rtt=7.63ms rc=3(EgressRtr) rsc=1
*A:Dut-A#

*A:Dut-A# oam ldp-treetrace prefix 10.20.1.6/32

ldp-treetrace for Prefix 10.20.1.6/32:

      127.0.0.1, ttl = 3 dst = 127.1.0.255 rc = EgressRtr status = Done
Hops: 127.0.0.1          127.0.0.1

      127.0.0.1, ttl = 3 dst = 127.2.0.255 rc = EgressRtr status = Done
Hops: 127.0.0.1          127.0.0.1

ldp-treetrace discovery state: Done
ldp-treetrace discovery status: ' OK '
Total number of discovered paths: 2
Total number of failed traces: 0

```

p2mp-lsp-ping

Syntax `p2mp-lsp-ping` { *lsp-name* **p2mp-instance** *instance-name* **s2l-dest-addr** *ip-address...* [*ip-address...up-to-5*] [**ttl** *label-ttl*] } | { **ldp** *p2mp-id* [**sender-addr** *ip-address*] [**leaf-addr** *ip-address...* [*ip-address...up-to-5*]] } [**fc** *fc-name*] [**size** *octets*] [**timeout** *timeout*] [**interval** *interval*] [**detail**]

Context oam

Description **Platforms Supported:** 7210 SAS Mxp, 7210 SAS- R6, 7210 SAS-R12, and 7210 SAS-T.

This command performs in-band connectivity test for an RSVP P2MP LSP.

The echo request message is sent on the active P2MP instance and is replicated in the data path over all branches of the P2MP LSP instance. By default, all egress LER nodes which are leaves of the P2MP LSP instance will reply to the echo request message.

The user can reduce the scope of the echo reply messages by explicitly entering a list of addresses for the egress LER nodes that are required to reply. A maximum of 5 addresses can be specified in a single run of the `p2mp-lsp-ping` command. A 7x50 LER node is able to parse the list of Egress LER addresses and if its address is included, it will reply with an echo reply message.

The output of the command without the **detail** option provides a high-level summary of error codes and/or success codes received. The output of the command with the **detail** option shows a line for each replying node as in the output of the LSP Ping for a P2P LSP.

The display will be delayed until all responses are received or the timer configured in the timeout parameter expired. No other CLI commands can be entered while waiting for the display. A ^C will abort the ping operation.

Parameters **fc** *fc-name* — The fc parameter is used to indicate the forwarding class of the MPLS echo request packet.

When an MPLS echo request packet is generated in CPM and is forwarded to the outgoing interface, the packet is queued in the egress network queue corresponding to the specified fc parameter value. The marking of the packet's EXP is dictated by the LSP-EXP mappings on the outgoing interface.

When the MPLS echo request packet is received on the responding node, The fc parameter value is dictated by the LSP-EXP mappings of the incoming interface.

When an MPLS echo reply packet is generated in CPM and is forwarded to the outgoing interface, the packet is queued in the egress network queue corresponding to the fc parameter value determined by the classification of the echo request packet, which is being replied to, at the incoming interface. The marking of the packet's EXP is dictated by the LSP-EXP mappings on the outgoing interface.

Values be, l2, af, l1, h2, ef, h1, nc

Default be

p2mp-instance *instance-name* — Configures the name, up to 32 characters long, of the specific instance of the P2MP LSP to send the echo request.

p2mp-lsp-ping *lsp-name* — Name that identifies an P2MP LSP, to 32 characters long, to ping.

s2l-dest-addr *ip-address* — Specifies the egress LER system address of the S2L sub-LSP path which is being traced.

Values ipv4-address: a.b.c.d

size *octets* — The size in octets, expressed as a decimal integer, of the MPLS echo request packet, including the IP header but not the label stack. The request payload is padded with zeros to the specified size. Note that an OAM command is not failed if the user entered a size lower than the minimum required to build the packet for the echo request message. The payload is automatically padded to meet the minimum size.

Values 1 to 9198

Default 1

ttl *label-ttl* — The TTL value for the MPLS label, expressed as a decimal integer..

Values 1 to 255

Default 255

ldp *p2mp-id*— Identifier to specify a LDP P2MP LSP to ping.

Values Must be a 32 bit integer. Range : 1 to 4294967295

sender-addr *ip-address* — Optional parameter to specify root-node address. If not, specified then the system IP address is used.

leaf-addr *ip-address...[ip-address...up-to-5]* — Specifies the list of egress LER system addresses which are required to reply to an LSP ping echo request message.

Values ipv4-address: a.b.c.d

timeout *timeout* — The timeout parameter in seconds, expressed as a decimal integer. This value is used to override the default timeout value and is the amount of time that the router will wait for an echo reply message from all leaves of the P2MP LSP after sending the message request message. Upon the expiration of message timeout, the requesting router assumes that the missing replies will not be received. Any echo reply message received after the request times out will be silently discarded.

Values 1 to 60

Default 3

Sample Output

```
*A:Dut-C# oam lsp-ping A2F_3

LSP-PING A2F_3: 92 bytes MPLS payload
Total S2L configured/up/responded = 400/390/388, round-trip min/avg/max
= <10/10/11 ms
Responses based on return code:
EgressRtr(3) = 387
NoFecMapping(4) = 1
```

Note: Missing responses on UP S2Ls implies "request timeout"

```
*A:Dut-A#
```

```
*A:Dut-C# oam lsp-ping A2F_3 detail
```

```
LSP-PING A2F_3: 92 bytes MPLS payload

=====
S2L Info
=====
From           Rtt           ReturnCode
-----
10.20.1.2      <10ms        EgressRtr
10.20.1.3      10ms         EgressRtr
10.20.1.5      11ms         EgressRtr
10.20.1.6      <10ms        EgressRtr
10.20.1.7      10ms         NoFecMapping
:
:
=====

Total S2L configured/up/responded = 400/390/388, round-trip min/avg/max
= <10/10/11 ms

Responses based on return code:
```

```

EgressRtr(3) = 387
NoFecMapping(4) = 1

*A:Dut-A#

*A:Dut-C# oam lsp-ping A2F_3 > p2mp-instance "1" s2l-dest-address 10.20.1.5 10.20.2.6
10.20.3.7

LSP-PING A2F_3: 132 bytes MPLS payload
P2MP Instance 1, S2L Egress 10.20.1.5 S2L Egress 10.20.2.6 S2L Egress 10.20.3.7

Total S2L configured/up/responded = 400/390/4, round-trip min/avg/max
= <10/10/11 ms

Responses based on return code:
EgressRtr(3) = 3
NoFecMapping(4) = 1

Note: Missing responses on UP S2Ls implies "request timeout"

*A:Dut-A#

*A:Dut-C# oam lsp-ping A2F_3 > p2mp-instance "1" s2l-dest-address 10.20.1.5 10.20.2.6
10.20.3.7 detail

LSP-PING A2F_3: 132 bytes MPLS payload
P2MP Instance 1, S2L Egress 10.20.1.5 S2L Egress 10.20.2.6 S2L Egress 10.20.3.7

=====
S2L Info
=====
From           Rtt           ReturnCode
-----
10.20.1.2      <10ms         EgressRtr
10.20.1.3      10ms          EgressRtr
10.20.1.4      Timeout       N/A
10.20.1.5      11ms          NoFecMapping
10.20.1.6      <10ms         EgressRtr
=====

Total S2L configured/up/responded = 400/390/4, round-trip min/avg/max
= <10/10/11 ms

Responses based on return code:
EgressRtr(3) = 3
NoFecMapping (4) = 1

*A:Dut-A#

```

p2mp-lsp-trace

Syntax `p2mp-lsp-trace {lsp-name [p2mp-instance instance-name] [s2l-dest-address ip-address...[ip-address...up-to-5]] | {ldp p2mp-id}} [fc fc-name] [size octets] [max-fail no-response-count]`

[probe-count probes-per-hop] [min-ttl min-label-ttl] [max-ttl max-label-ttl] [timeout timeout] [interval interval] [detail]

Context oam

Description **Platforms Supported:** 7210 SAS Mxp, 7210 SAS-R6, 7210 SAS-R12, and 7210 SAS-T.

This command discovers and displays the hop-by-hop path for a source-to-leaf (S2L) sub-LSP of an RSVP P2MP LSP.

The LSP trace capability allows the user to trace the path of a single S2L path of a P2MP LSP. Its operation is similar to that of the `p2mp-lsp-ping`, but the sender of the echo reply request message includes the downstream mapping TLV to request the downstream branch information from a branch LSR or bud LSR. The branch LSR or bud LSR will then also include the downstream mapping TLV to report the information about the downstream branches of the P2MP LSP. An egress LER must not include this TLV in the echo response message.

The parameter **probe-count** operates in the same way as in LSP Trace on a P2P LSP. It represents the maximum number of probes sent per TTL value before giving up on receiving the echo reply message. If a response is received from the traced node before reaching maximum number of probes, then no more probes are sent for the same TTL. The sender of the echo request then increments the TTL and uses the information it received in the downstream mapping TLV to start sending probes to the node downstream of the last node which replied. This continues until the egress LER for the traced S2L path replied.

Similar to `p2mp-lsp-ping`, an LSP trace probe results on all egress LER nodes eventually receiving the echo request message but only the traced egress LER node will reply to the last probe.

Also any branch LSR node or bud LSR node in the P2MP LSP tree may receive a copy of the echo request message with the TTL in the outer label expiring at this node. However, only a branch LSR or bud LSR which has a downstream branch over which the traced egress LER is reachable will respond.

When a branch LSR or bud LSR responds, it sets the global return code in the echo response message to RC=14 - "See DDMAP TLV for Return Code and Return Sub-Code" and the return code in the DDMAP TLV corresponding to the outgoing interface of the branch used by the traced S2L path to RC=8 - "Label switched at stack-depth <RSC>".

Parameters **fc fc-name** — The `fc` parameter is used to indicate the forwarding class of the MPLS echo request packet.

When an MPLS echo request packet is generated in CPM and is forwarded to the outgoing interface, the packet is queued in the egress network queue corresponding to the specified `fc` parameter value. The marking of the packet's EXP is dictated by the LSP-EXP mappings on the outgoing interface.

When the MPLS echo request packet is received on the responding node, The `fc` parameter value is dictated by the LSP-EXP mappings of the incoming interface.

When an MPLS echo reply packet is generated in CPM and is forwarded to the outgoing interface, the packet is queued in the egress network queue corresponding to the `fc` parameter value determined by the classification of the echo request packet, which is being replied to, at the incoming interface. The marking of the packet's EXP is dictated by the LSP-EXP mappings on the outgoing interface.

Values be, l2, af, l1, h2, ef, h1, nc

Default be

interval *interval* — The interval parameter in seconds, expressed as a decimal integer. This parameter is used to override the default echo request message send interval and defines the minimum amount of time that must expire before the next echo request message is sent. If the interval is set to 1 second, and the timeout value is set to 10 seconds, then the maximum time between message requests is 10 seconds and the minimum is 1 second. This depends upon the receipt of an echo reply message corresponding to the outstanding message request.

Values 1 to 10

Default 1

ldp *p2mp-id* — Identifier to specify a LDP P2MP LSP to ping.

Values Must be a 32 bit integer. Range : 1 to 4294967295

p2mp-lsp-trace *lsp-name* — Name that identifies an P2MP LSP, to 32 characters long, to ping.

max-fail *no-response-count* — The maximum number of consecutive MPLS echo requests, expressed as a decimal integer that do not receive a reply before the trace operation fails for a given TTL.

Values 1 to 255

Default 5

max-ttl *max-label-ttl* — Specifies the maximum TTL value in the MPLS label for the LSP trace test, expressed as a decimal integer.

Values 1 to 255

Default 30

min-ttl *min-label-ttl* — Specifies the minimum TTL value in the MPLS label for the LSP trace test, expressed as a decimal integer.

Values 1 to 255

Default 1

p2mp-instance *instance-name* — Configures the name, up to 32 characters long, of the specific instance of the P2MP LSP to send the echo request.

probe-count *probes-per-hop* — Specifies the number of LSP trace echo request messages to send per TTL value.

Values 1 to 10

Default 1

s2l-dest-addr *ip-address* — Specifies the egress LER system address of the S2L sub-LSP path which is being traced.

Values ipv4-address: a.b.c.d

size *octets* — The size in octets, expressed as a decimal integer, of the MPLS echo request packet, including the IP header but not the label stack. The request payload is padded with zeros to the specified size. Note that an OAM command is not failed if the user entered a size lower than the minimum required to build the packet for the echo request message. The payload is automatically padded to meet the minimum size.

Values 1 to 9198

Default 1

timeout *timeout* — The timeout parameter in seconds, expressed as a decimal integer. This value is used to override the default timeout value and is the amount of time that the router will wait for an echo reply message from all leaves of the P2MP LSP after sending the message request message. Upon the expiration of message timeout, the requesting router assumes that the missing replies will not be received. Any echo reply message received after the request times out will be silently discarded.

Values 1 to 60

Default 3

Sample Output

```
*A:Dut-C# oam p2mp-lsp-trace "p2mp_1" p2mp-instance "1" s2l-dest-address 10.20.1.5 detail
P2MP LSP p2mp_1: 132 bytes MPLS payload
P2MP Instance 1, S2L Egress 10.20.1.5

  1  10.20.1.1  rtt=3.78 ms rc=8(DSRtrMatchLabel)
      DS 1: ipaddr 10.20.1.2 iftype 'ipv4Unnumbered' ifaddr 2 MRU=1500
label=131060 proto=4(RSVP-TE) B/E flags:0/0
  2  10.20.1.2  rtt=3.54 ms rc=8(DSRtrMatchLabel)
      DS 1: ipaddr 10.20.1.4 iftype 'ipv4Unnumbered' ifaddr 3 MRU=1500
label=131061 proto=4(RSVP-TE) B/E flags:0/0
  3  10.20.1.5  rtt=5.30 ms rc=5(DSMappingMismatched)

*A:Dut-A#
```

Service Diagnostics

sdp-mtu

Note: This command is not supported on 7210 SAS-M and 7210 SAS-T devices configured in Access uplink mode.

Syntax **sdp-mtu** *orig-sdp-id* **size-inc** *start-octets end-octets* [**step** *step-size*] [**timeout** *timeout*] [**interval** *interval*]

Context oam

Description Performs MTU Path tests on an SDP to determine the largest path-mtu supported on an SDP. The **size-inc** parameter can be used to easily determine the **path-mtu** of a given SDP-ID. The forwarding class is assumed to be Best-Effort Out-of-Profile. The message reply is returned with IP encapsulation from the far-end 7210 SAS M. OAM request messages sent within an IP SDP must have the 'DF' IP header bit set to 1 to prevent message fragmentation.

To terminate an **sdp-mtu** in progress, use the CLI break sequence <Ctrl-C>.

Special Cases **SDP Path MTU Tests** — SDP Path MTU tests can be performed using the **sdp-mtu size-inc** keyword to easily determine the **path-mtu** of a given SDP-ID. The forwarding class is assumed to be Best-Effort Out-of-Profile. The message reply is returned with IP encapsulation from the far-end 7210 SAS M.

With each OAM Echo Request sent using the **size-inc** parameter, a response line is displayed as message output. The path MTU test displays incrementing packet sizes, the number sent at each size until a reply is received and the response message.

As the request message is sent, its size value is displayed followed by a period for each request sent of that size. Up to three requests will be sent unless a valid response is received for one of the requests at that size. Once a response is received, the next size message is sent.

The response message indicates the result of the message request.

After the last reply has been received or response timeout, the maximum size message replied to indicates the largest size OAM Request message that received a valid reply.

Parameters *orig-sdp-id* — The *sdp-id* to be used by **sdp-ping**, expressed as a decimal integer. The far-end address of the specified *sdp-id* is the expected *responder-id* within each reply received. The specified *sdp-id* defines the encapsulation of the SDP tunnel encapsulation used to reach the far end. This can be IP MPLS. If *orig-sdp-id* is invalid or administratively down or unavailable for some reason, the SDP echo request message is not sent and an appropriate error message is displayed (once the **interval** timer expires, **sdp-ping** will attempt to send the next request if required).

Values 1 — 17407

size-inc *start-octets end-octets* — Indicates an incremental path MTU test will be performed with by sending a series of message requests with increasing MTU sizes. The *start-octets* and *end-octets* parameters are described below.

start-octets — The beginning size in octets of the first message sent for an incremental MTU test, expressed

as a decimal integer.

Values 40 — 9198

end-octets — The ending size in octets of the last message sent for an incremental MTU test, expressed as a decimal integer. The specified value must be greater than *start-octets*.

Values 40 — 9198

step *step-size* — The number of octets to increment the message size request for each message sent for an incremental MTU test, expressed as a decimal integer. The next size message will not be sent until a reply is received or three messages have timed out at the current size.

If the incremented size exceeds the *end-octets* value, no more messages will be sent.

Default 32

Values 1 — 512

timeout *timeout* — The **timeout** parameter in seconds, expressed as a decimal integer. This value is used to override the default **timeout** value and is the amount of time that the router will wait for a message reply after sending the message request. Upon the expiration of message timeout, the requesting router assumes that the message response will not be received. A 'request timeout' message is displayed by the CLI for each message request sent that expires. Any response received after the request times out will be silently discarded.

Default 5

Values 1 — 10

interval *interval* — The **interval** parameter in seconds, expressed as a decimal integer. This parameter is used to override the default request message send interval and defines the minimum amount of time that must expire before the next message request is sent.

If the **interval** is set to 1 second, and the **timeout** value is set to 10 seconds, then the maximum time between message requests is 10 seconds and the minimum is 1 second. This depends upon the receipt of a message reply corresponding to the outstanding message request.

Default 1

Values 1 — 10

Output Sample SDP MTU Path Test Sample Output

```
*A:Dut-A# oam sdp-mtu 1201 size-inc 512 3072 step 256
Size      Sent      Response
-----
512       .         Success
768       .         Success
1024      .         Success
1280      .         Success
1536      .         Success
1792      .         Success
2048      .         Success
2304      .         Success
2560      .         Success
2816      .         Success
```

Service Diagnostics

```
3072 . Success
Maximum Response Size: 3072
*A:Dut-A#
```

svc-ping

Note: This command is not supported on 7210 SAS-M and 7210 SAS-T devices configured in Access uplink mode.

Syntax **svc-ping** *ip-address* [**service** *service-id*] [**local-sdp**] [**remote-sdp**]

Context <GLOBAL>

Description Tests a service ID for correct and consistent provisioning between two service end points.

The **svc-ping** command accepts a far-end IP address and a *service-id* for local and remote service testing. The following information can be determined from **svc-ping**:

1. Local and remote service existence
2. Local and remote service state
3. Local and remote service type correlation
4. Local and remote customer association
5. Local and remote service-to-SDP bindings and state
6. Local and remote ingress and egress service label association

Unlike **sdp-ping**, only a single message will be sent per command; no count nor interval parameter is supported and round trip time is not calculated. A timeout value of 10 seconds is used before failing the request. The forwarding class is assumed to be Best-Effort Out-of-Profile

If no request is sent or a reply is not received, all remote information will be shown as N/A.

To terminate a **svc-ping** in progress, use the CLI break sequence <Ctrl-C>.

Upon request timeout, message response, request termination, or request error the following local and remote information will be displayed. Local and remote information will be dependent upon service existence and reception of reply.

| Field | Description | Values |
|-------|-------------|--------|
|-------|-------------|--------|

| Field | Description | Values (Continued) |
|----------------|--|--|
| Request Result | The result of the svc-ping request message. | Sent - Request Timeout
Sent - Request Terminated
Sent - Reply Received
Not Sent - Non-Existent Service-ID
Not Sent - Non-Existent SDP for Service
Not Sent - SDP For Service Down
Not Sent - Non-existent Service Egress Label |
| Service-ID | The ID of the service being tested. | <i>service-id</i> |

Service Diagnostics

| Field | Description | Values (Continued) |
|----------------------------|---|---|
| Local Service Type | The type of service being tested. If <i>service-id</i> does not exist locally, N/A is displayed. | Epip
TLS
IES
Mirror-Dest
N/A |
| Local Service Admin State | The local administrative state of <i>service-id</i> . If the service does not exist locally, the administrative state will be Non-Existent. | Admin-Up
Admin-Down
Non-Existent |
| Local Service Oper State | The local operational state of <i>service-id</i> . If the service does not exist locally, the state will be N/A. | Oper-Up
Oper-Down
N/A |
| Remote Service Type | The remote type of service being tested. If <i>service-id</i> does not exist remotely, N/A is displayed. | Epip, Ipipe
TLS
IES
Mirror-Dest
N/A |
| Remote Service Admin State | The remote administrative state of <i>service-id</i> . If the service does not exist remotely, the administrative state is Non-Existent. | Up
Down
Non-Existent |
| Local Service MTU | The local service-mtu for <i>service-id</i> . If the service does not exist, N/A is displayed. | <i>service-mtu</i>
N/A |
| Remote Service MTU | The remote service-mtu for <i>service-id</i> . If the service does not exist remotely, N/A is displayed. | <i>remote-service-mtu</i>
N/A |
| Local Customer ID | The local <i>customer-id</i> associated with <i>service-id</i> . If the service does not exist locally, N/A is displayed. | <i>customer-id</i>
N/A |
| Remote Customer ID | The remote <i>customer-id</i> associated with <i>service-id</i> . If the service does not exist remotely, N/A is displayed. | <i>customer-id</i>
N/A |
| Local Service IP Address | The local system IP address used to terminate remotely configured SDP-ID (as the far-end address). If an IP interface has not been configured to be the system IP address, N/A is displayed. | <i>system-ip-address</i>
N/A |

| Field | Description | Values (Continued) |
|---------------------------------|---|-------------------------------------|
| Local Service IP Interface Name | The name of the local system IP interface. If the local system IP interface has not been created, N/A is displayed. | <i>system-interface-name</i>
N/A |

Service Diagnostics

| Field | Description | Values (Continued) |
|---|--|--|
| Local Service IP Interface State | The state of the local system IP interface. If the local system IP interface has not been created, Non-Existent is displayed. | Up
Down
Non-Existent |
| Expected Far-end Address | The expected IP address for the remote system IP interface. This must be the far-end address entered for the svc-ping command. | <i>orig-sdp-far-end-addr</i>
<i>dest-ip-addr</i>
N/A |
| Actual Far-end Address | The returned remote IP address. If a response is not received, the displayed value is N/A. If the far-end service IP interface is down or non-existent, a message reply is not expected. sdp-ping should also fail. | <i>resp-ip-addr</i>
N/A |
| Responders Expected Far-end Address | The expected source of the originator's <i>sdp-id</i> from the perspective of the remote router terminating the <i>sdp-id</i> . If the far-end cannot detect the expected source of the ingress <i>sdp-id</i> or the request is transmitted outside the <i>sdp-id</i> , N/A is displayed. | <i>resp-rec-tunnel-far-end-address</i>
N/A |
| Originating SDP-ID | The <i>sdp-id</i> used to reach the far-end IP address if sdp-path is defined. The originating <i>sdp-id</i> must be bound to the <i>service-id</i> and terminate on the far-end IP address. If an appropriate originating <i>sdp-id</i> is not found, Non-Existent is displayed. | orig-sdp-id
Non-Existent |
| Originating SDP-ID Path Used | Whether the Originating router used the originating <i>sdp-id</i> to send the svc-ping request. If a valid originating <i>sdp-id</i> is found, operational and has a valid egress service label, the originating router should use the <i>sdp-id</i> as the requesting path if sdp-path has been defined. If the originating router uses the originating <i>sdp-id</i> as the request path, Yes is displayed. If the originating router does not use the originating <i>sdp-id</i> as the request path, No is displayed. If the originating <i>sdp-id</i> is non-existent, N/A is displayed. | Yes
No
N/A |
| Originating SDP-ID Administrative State | The local administrative state of the originating <i>sdp-id</i> . If the <i>sdp-id</i> has been shutdown, Admin-Down is displayed. If the originating <i>sdp-id</i> is in the no shutdown state, Admin-Up is displayed. If an originating <i>sdp-id</i> is not found, N/A is displayed. | Admin-Up
Admin-Up
N/A |
| Originating SDP-ID Operating State | The local operational state of the originating <i>sdp-id</i> . If an originating <i>sdp-id</i> is not found, N/A is displayed. | Oper-Up
Oper-Down
N/A |
| Originating SDP-ID Binding Admin State | The local administrative state of the originating <i>sdp-ids</i> binding to <i>service-id</i> . If an <i>sdp-id</i> is not bound to the service, N/A is displayed. | Admin-Up
Admin-Up
N/A |

| Field | Description | Values (Continued) |
|--|---|------------------------------------|
| Originating SDP-ID Binding Oper State | The local operational state of the originating <i>sdp-ids</i> binding to <i>service-id</i> . If an <i>sdp-id</i> is not bound to the service, N/A is displayed. | Oper-Up
Oper-Down
N/A |
| Responding SDP-ID | The <i>sdp-id</i> used by the far end to respond to the svc-ping request. If the request was received without the sdp-path parameter, the responding router will not use an <i>sdp-id</i> as the return path, but the appropriate responding <i>sdp-id</i> will be displayed. If a valid <i>sdp-id</i> return path is not found to the originating router that is bound to the <i>service-id</i> , Non-Existent is displayed. | <i>resp-sdp-id</i>
Non-Existent |
| Responding SDP-ID Path Used | Whether the responding router used the responding <i>sdp-id</i> to respond to the svc-ping request. If the request was received via the originating <i>sdp-id</i> and a valid return <i>sdp-id</i> is found, operational and has a valid egress service label, the far-end router should use the <i>sdp-id</i> as the return <i>sdp-id</i> . If the far end uses the responding <i>sdp-id</i> as the return path, Yes is displayed. If the far end does not use the responding <i>sdp-id</i> as the return path, No is displayed. If the responding <i>sdp-id</i> is non-existent, N/A is displayed. | Yes
No
N/A |
| Responding SDP-ID Administrative State | The administrative state of the far-end <i>sdp-id</i> associated with the return path for <i>service-id</i> . When a return path is administratively down, Admin-Down is displayed. If the return <i>sdp-id</i> is administratively up, Admin-Up is displayed. If the responding <i>sdp-id</i> is non-existent, N/A is displayed. | Admin-Up
Admin-Up
N/A |
| Responding SDP-ID Operational State | The operational state of the far-end <i>sdp-id</i> associated with the return path for <i>service-id</i> . When a return path is operationally down, Oper-Down is displayed. If the return <i>sdp-id</i> is operationally up, Oper-Up is displayed. If the responding <i>sdp-id</i> is non-existent, N/A is displayed. | Oper-Up
Oper-Down
N/A |
| Responding SDP-ID Binding Admin State | The local administrative state of the responder's <i>sdp-id</i> binding to <i>service-id</i> . If an <i>sdp-id</i> is not bound to the service, N/A is displayed. | Admin-Up
Admin-Down
N/A |
| Responding SDP-ID Binding Oper State | The local operational state of the responder's <i>sdp-id</i> binding to <i>service-id</i> . If an <i>sdp-id</i> is not bound to the service, N/A is displayed. | Oper-Up
Oper-Down
N/A |
| Originating VC-ID | The originator's VC-ID associated with the <i>sdp-id</i> to the far-end address that is bound to <i>service-id</i> . If the <i>sdp-id</i> signaling is off, <i>originator-vc-id</i> is 0. If the <i>originator-vc-id</i> does not exist, N/A is displayed. | <i>originator-vc-id</i>
N/A |

Service Diagnostics

| Field | Description | Values (Continued) |
|---|---|--|
| Responding VC-ID | The responder's VC-ID associated with the <i>sdp-id</i> to <i>originator-id</i> that is bound to <i>service-id</i> . If the <i>sdp-id</i> signaling is off or the service binding to <i>sdp-id</i> does not exist, <i>responder-vc-id</i> is 0. If a response is not received, N/A is displayed. | <i>responder-vc-id</i>
N/A |
| Originating Egress Service Label | The originating service label (VC-Label) associated with the <i>service-id</i> for the originating <i>sdp-id</i> . If <i>service-id</i> does not exist locally, N/A is displayed. If <i>service-id</i> exists, but the egress service label has not been assigned, Non-Existent is displayed. | <i>egress-vc-label</i>
N/A
Non-Existent |
| Originating Egress Service Label Source | The originating egress service label source. If the displayed egress service label is manually defined, Manual is displayed. If the egress service label is dynamically signaled, Signaled is displayed. If the <i>service-id</i> does not exist or the egress service label is non-existent, N/A is displayed. | Manual
Signaled
N/A |
| Originating Egress Service Label State | The originating egress service label state. If the originating router considers the displayed egress service label operational, Up is displayed. If the originating router considers the egress service label inoperative, Down is displayed. If the <i>service-id</i> does not exist or the egress service label is non-existent, N/A is displayed. | Up
Down
N/A |
| Responding Service Label | The actual responding service label in use by the far-end router for this <i>service-id</i> to the originating router. If <i>service-id</i> does not exist in the remote router, N/A is displayed. If <i>service-id</i> does exist remotely but the remote egress service label has not been assigned, Non-Existent is displayed. | <i>rec-vc-label</i>
N/A
Non-Existent |
| Responding Egress Service Label Source | The responder's egress service label source. If the responder's egress service label is manually defined, Manual is displayed. If the responder's egress service label is dynamically signaled, Signaled is displayed. If the <i>service-id</i> does not exist on the responder or the responder's egress service label is non-existent, N/A is displayed. | Manual
Signaled
N/A |
| Responding Service Label State | The responding egress service label state. If the responding router considers it is an egress service label operational, Up is displayed. If the responding router considers it is an egress service label inoperative, Down is displayed. If the <i>service-id</i> does not exist or the responder's egress service label is non-existent, N/A is displayed. | Up
Down
N/A |
| Expected Ingress Service Label | The locally assigned ingress service label. This is the service label that the far-end is expected to use for <i>service-id</i> when sending to the originating router. If <i>service-id</i> does not exist locally, N/A is displayed. If <i>service-id</i> exists but an ingress service label has not been assigned, Non-Existent is displayed. | <i>ingress-vc-label</i>
N/A
Non-Existent |

| Field | Description | Values (Continued) |
|--|--|---|
| Expected Ingress Label Source | The originator's ingress service label source. If the originator's ingress service label is manually defined, Manual is displayed. If the originator's ingress service label is dynamically signaled, Signaled is displayed. If the <i>service-id</i> does not exist on the originator or the originators ingress service label has not been assigned, N/A is displayed. | Manual
Signaled
N/A |
| Expected Ingress Service Label State | The originator's ingress service label state. If the originating router considers it as an ingress service label operational, Up is displayed. If the originating router considers it as an ingress service label inoperative, Down is displayed. If the <i>service-id</i> does not exist locally, N/A is displayed. | Up
Down
N/A |
| Responders Ingress Service Label | The assigned ingress service label on the remote router. This is the service label that the far end is expecting to receive for <i>service-id</i> when sending to the originating router. If <i>service-id</i> does not exist in the remote router, N/A is displayed. If <i>service-id</i> exists, but an ingress service label has not been assigned in the remote router, Non-Existent is displayed. | <i>resp-ingress-vc-label</i>
N/A
Non-Existent |
| Responders Ingress Label Source | The assigned ingress service label source on the remote router. If the ingress service label is manually defined on the remote router, Manual is displayed. If the ingress service label is dynamically signaled on the remote router, Signaled is displayed. If the <i>service-id</i> does not exist on the remote router, N/A is displayed. | Manual
Signaled
N/A |
| Responders Ingress Service Label State | The assigned ingress service label state on the remote router. If the remote router considers it as an ingress service label operational, Up is displayed. If the remote router considers it as an ingress service label inoperative, Down is displayed. If the <i>service-id</i> does not exist on the remote router or the ingress service label has not been assigned on the remote router, N/A is displayed. | Up
Down
N/A |
| Parameters | <p><i>ip-address</i> — The far-end IP address to which to send the svc-ping request message in dotted decimal notation.</p> <p>service <i>service-id</i> — The service ID of the service being tested must be indicated with this parameter. The service ID need not exist on the local to receive a reply message.</p> <p>Values 1 — 2147483647</p> <p>local-sdp — Specifies the svc-ping request message should be sent using the same service tunnel encapsulation labeling as service traffic. If local-sdp is specified, the command attempts to use an egress <i>sdp-id</i> bound to the service with the specified far-end IP address with the VC-Label for the service. The far-end address of the specified <i>sdp-id</i> is the expected <i>responder-id</i> within the reply received. The <i>sdp-id</i> defines the encapsulation of the SDP tunnel encapsulation used to reach the far end; this can be IP or MPLS. On originator egress, the service-ID must have an associated VC-Label to reach the far-end</p> | |

address of the *sdp-id* and the *sdp-id* must be operational for the message to be sent.

If **local-sdp** is not specified, the **svc-ping** request message is sent with encapsulation with the OAM label.

The following table indicates whether a message is sent and how the message is encapsulated based on the state of the service ID.

| Local Service State | local-sdp Not Specified | | local-sdp Specified | |
|---|-------------------------|-----------------------|---------------------|---|
| | Message Sent | Message Encapsulation | Message Sent | Message Encapsulation |
| Invalid Local Service | Yes | Generic IP/OAM (PLP) | No | None |
| No Valid SDP-ID Bound | Yes | Generic IP/OAM (PLP) | No | None |
| SDP-ID Valid But Down | Yes | Generic IP/OAM (PLP) | No | None |
| SDP-ID Valid and Up, But No Service Label | Yes | Generic IP/OAM (PLP) | No | None |
| SDP-ID Valid, Up and Egress Service Label | Yes | Generic IP/OAM (PLP) | Yes | SDP Encapsulation with Egress Service Label (SLP) |

remote-sdp — Specifies **svc-ping** reply message from the **far-end** should be sent using the same service tunnel encapsulation labeling as service traffic.

If **remote-sdp** is specified, the **far-end** responder attempts to use an egress *sdp-id* bound to the service with the message originator as the destination IP address with the VC-Label for the service. The *sdp-id* defines the encapsulation of the SDP tunnel encapsulation used to reply to the originator; this can be IP or MPLS. On responder egress, the service-ID must have an associated VC-Label to reach the originator address of the *sdp-id* and the *sdp-id* must be operational for the message to be sent.

The following table indicates how the message response is encapsulated based on the state of the remote service ID.

| Remote Service State | Message Encapsulation | |
|---|-----------------------------|---|
| | remote-sdp
Not Specified | remote-sdp
Specified |
| Invalid Ingress Service Label | Generic IP/GRE OAM (PLP) | Generic IP/GRE OAM (PLP) |
| Invalid Service-ID | Generic IP/GRE OAM (PLP) | Generic IP/GRE OAM (PLP) |
| No Valid SDP-ID Bound on Service-ID | Generic IP/GRE OAM (PLP) | Generic IP/GRE OAM (PLP) |
| SDP-ID Valid But Down | Generic IP/GRE OAM (PLP) | Generic IP/GRE OAM (PLP) |
| SDP-ID Valid and Up, but No Service Label | Generic IP/GRE OAM (PLP) | Generic IP/GRE OAM (PLP) |
| SDP-ID Valid and Up, Egress Service Label, but VC-ID Mismatch | Generic IP/GRE OAM (PLP) | Generic IP/GRE OAM (PLP) |
| SDP-ID Valid and Up, Egress Service Label, but VC-ID Match | Generic IP/GRE OAM (PLP) | SDP Encapsulation with Egress Service Label (SLP) |

Sample Output

```
A:ALU_G7X1>config# oam svc-ping 10.20.1.3 service 1
Service-ID: 1
```

```
Err Info          Local          Remote
-----
Type:             EPIPE         EPIPE
Admin State:      Up            Up
==> Oper State:   Down          Down
Service-MTU:      1514         1514
Customer ID:      1            1

IP Interface State: Up
Actual IP Addr:   10.20.1.1    10.20.1.3
Expected Peer IP: 10.20.1.3    10.20.1.1

SDP Path Used:    No            No
SDP-ID:           1            2
Admin State:      Up            Up
Operative State:  Up            Up
Binding Admin State: Up          Up
Binding Oper State: Up          Up
Binding VC ID:    10           10
Binding Type:     Spoke         Spoke
Binding Vc-type:  Ether         Ether
Binding Vlan-vc-tag: N/A        N/A

Egress Label:     131070       131068
Ingress Label:    131068       131070
Egress Label Type: Signaled     Signaled
Ingress Label Type: Signaled     Signaled
```

```
Request Result: Send - Reply Received: Responder Service ID Oper-Down
```

```
A:ALU_G7X1>config#
```

vprn-ping

Syntax **vprn-ping** *service-id* **source** *ip-address* **destination** *ip-address* [**fc** *fc-name*][**size** *size*] [**ttl** *vc-label-ttl*] [**return-control**] [**interval** *interval*] [**send-count** *send-count*] [**timeout** *timeout*]

Context <GLOBAL>
config>saa>test>type

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command performs a VPRN ping.

Parameters **service** *service-id* — The VPRN service ID to diagnose or manage.

Values *service-id:* 1 — 2147483647

source *ip-address* — The IP prefix for the source IP address in dotted decimal notation.

Values *ipv4-address:* 0.0.0.0 — 255.255.255.255

destination *ip-address* — The IP prefix for the destination IP address in dotted decimal notation.

Values 0.0.0.0 — 255.255.255.255

size *octets* — The OAM request packet size in octets, expressed as a decimal integer.

Values 1 — 9198

ttl *vc-label-ttl* — The TTL value in the VC label for the OAM request, expressed as a decimal integer.

Default 255

Values 1 — 255

return-control — Specifies the response to come on the control plane.

interval *interval* — The **interval** parameter in seconds, expressed as a decimal integer. This parameter is used to override the default request message send interval and defines the minimum amount of time that must expire before the next message request is sent.

If the **interval** is set to 1 second where the **timeout** value is set to 10 seconds, then the maximum time between message requests is 10 seconds and the minimum is 1 second. This depends upon the receipt of a message reply corresponding to the outstanding message request.

Default 1

Values 1 — 10 seconds

send-count *send-count* — The number of messages to send, expressed as a decimal integer. The **count** parameter is used to override the default number of message requests sent. Each message request must either timeout or receive a reply before the next message request is sent. The message **interval** value must be expired before the next message request is sent.

Default 1

Values 1 — 100

timeout *timeout* — The **timeout** parameter in seconds, expressed as a decimal integer. This value is used to override the default **timeout** value and is the amount of time that the router will wait for a message reply after sending the message request. Upon the expiration of message timeout, the requesting router assumes that the message response will not be received. Any response received after the request times out will be silently discarded.

Default 5

Values 1 — 100

fc-name — The forwarding class of the MPLS echo request encapsulation.

Default be

Values be, l2, af, l1, h2, ef, h1, nc

Sample Output

```
A:PE_1# oam vprn-ping 25 source 10.4.128.1 destination 10.16.128.0
Sequence Node-id                               Reply-Path Size      RTT
-----
[Send request Seq. 1.]
1         10.128.0.3:cpm                       In-Band   100       0ms
-----
...
A:PE_1#
-----
A:PE_1#
```

vprn-trace

Syntax **vprn-trace** *service-id* **source** *src-ip* **destination** *ip-address* [**fc** *fc-name*] [**size** *size*] [**min-ttl** *vc-label-ttl*] [**max-ttl** *vc-label-ttl*] [**return-control**] [**send-count** *send-countf*] [**interval** *seconds*] [**timeout** *timeoutf*]

Context <GLOBAL>
config>saa>test>type

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

Performs VPRN trace.

Parameters **service** *service-id* — The VPRN service ID to diagnose or manage.

Values *service-id:* 1 — 2147483647

source *src-ip* — The IP prefix for the source IP address in dotted decimal notation.

Values ipv4-address: 0.0.0.0 — 255.255.255.255

destination *dst-ip* — The IP prefix for the destination IP address in dotted decimal notation.

Values 0.0.0.0 — 255.255.255.255

size *octets* — The OAM request packet size in octets, expressed as a decimal integer.

Values [1..9198]

min-ttl *vc-label-ttl* — The minimum TTL value in the VC label for the trace test, expressed as a decimal integer.

Default 1

Values 1 — 255

max-ttl *vc-label-ttl* — The maximum TTL value in the VC label for the trace test, expressed as a decimal integer.

Default 4

Values 1 — 255

return-control — Specifies the OAM reply to a data plane OAM request be sent using the control plane instead of the data plane.

Default OAM reply sent using the data plane.

send-count *sendcount* — The number of OAM requests sent for a particular TTL value, expressed as a decimal integer.

Default 1

Values 1 — 10

interval *seconds* — The **interval** parameter in seconds, expressed as a decimal integer. This parameter is used to override the default request message send interval and defines the minimum amount of time that must expire before the next message request is sent.

If the **interval** is set to 1 second where the **timeout** value is set to 10 seconds, then the maximum time between message requests is 10 seconds and the minimum is 1 second. This depends upon the receipt of a message reply corresponding to the outstanding message request.

Default 1

Values 1 — 10 seconds

timeout *timeout* — The **timeout** parameter in seconds, expressed as a decimal integer. This value is used to override the default **timeout** value and is the amount of time that the router will wait for a message reply after sending the message request. Upon the expiration of message timeout, the requesting router assumes that the message response will not be received. Any response received after the request times out will be silently discarded.

Default 3

Values 1 — 60

fc-name — The forwarding class of the MPLS echo request encapsulation.

Default be

Values be, l2, af, ll, h2, ef, h1, nc

Sample Output

```
A:PE_1# oam vprn-trace 25 source 10.4.128.1 destination 10.16.128.0
```

```
TTL Seq Reply Node-id          Rcvd-on          Reply-Path RTT
```

```
-----
```

```
[Send request TTL: 1, Seq. 1.]
```

```
1 1 1 10.128.0.4 cpm In-Band 0ms
```

```
Requestor 10.128.0.1 Route: 0.0.0.0/0
```

```
Vpn Label: 131071 Metrics 0 Pref 170 Owner bgpVpn
```

```
Next Hops: [1] ldp tunnel
```

```
Route Targets: [1]: target:65100:1
```

```
Responder 10.128.0.4 Route: 10.16.128.0/24
```

```
Vpn Label: 131071 Metrics 0 Pref 170 Owner bgpVpn
```

```
Next Hops: [1] ldp tunnel
```

```
Route Targets: [1]: target:65001:100
```

```
[Send request TTL: 2, Seq. 1.]
```

```
2 1 1 10.128.0.3 cpm In-Band 0ms
```

```
Requestor 10.128.0.1 Route: 0.0.0.0/0
```

```
Vpn Label: 131071 Metrics 0 Pref 170 Owner bgpVpn
```

```
Next Hops: [1] ldp tunnel
```

```
Route Targets: [1]: target:65100:1
```

```
Responder 10.128.0.3 Route: 10.16.128.0/24
```

```
Vpn Label: 0 Metrics 0 Pref 0 Owner local
```

```
Next Hops: [1] ifIdx 2 nextHopIp 10.16.128.0
```

```
[Send request TTL: 3, Seq. 1.]
```

```
[Send request TTL: 4, Seq. 1.]
```

```
...
```

```
-----
```

```
A:PE_1#
```

VPLS MAC Diagnostics

Note: VPLS MAC diagnostics commands are not supported on 7210 SAS-M and 7210 SAS-T devices configured in Access uplink mode.

cpe-ping

Syntax **cpe-ping service** *service-id* **destination** *ip-address* **source** *ip-address* [**ttl** *vc-label-ttl*] [**return-control**] [**source-mac** *ieee-address*] [**fc** *fc-name*] [**interval** *interval*] [**count** *send-count*]

Context oam
config>saa>test>type

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This ping utility determines the IP connectivity to a CPE within a specified VPLS service.

Parameters **service** *service-id* — The service ID of the service to diagnose or manage.

Values *service-id:* 1 — 2147483647

destination *ip-address* — Specifies the IP address to be used as the destination for performing an OAM ping operations.

source *ip-address* — Specify an unused IP address in the same network that is associated with the VPLS.

ttl *vc-label-ttl* — The TTL value in the VC label for the OAM MAC request, expressed as a decimal integer.

Default 255

Values 1 — 255

return-control — Specifies the MAC OAM reply to a data plane MAC OAM request be sent using the control plane instead of the data plane.

Default MAC OAM reply sent using the data plane.

source-mac *ieee-address* — Specify the source MAC address that will be sent to the CPE. If not specified or set to 0, the MAC address configured for the CPM is used.

fc-name — The forwarding class of the MPLS echo request encapsulation.

Default be

Values be, l2, af, l1, h2, ef, h1, nc

interval *interval* — The **interval** parameter in seconds, expressed as a decimal integer. This parameter is used to override the default request message send interval and defines the minimum amount of time that must expire before the next message request is sent.

If the **interval** is set to 1 second where the **timeout** value is set to 10 seconds, then the maximum time between message requests is 10 seconds and the minimum is 1 second. This depends upon the receipt of

a message reply corresponding to the outstanding message request.

Default 1

Values 1 — 10

count *send-count* — The number of messages to send, expressed as a decimal integer. The **count** parameter is used to override the default number of message requests sent. Each message request must either time-out or receive a reply before the next message request is sent. The message **interval** value must be expired before the next message request is sent.

Default 1

Values 1 — 100

mac-populate

Syntax **mac-populate** *service-id mac ieee-address* [**flood**] [**age seconds**] [**force**] [*target-sap sap-id*]

Context oam

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command populates the FIB with an OAM-type MAC entry indicating the node is the egress node for the MAC address and optionally floods the OAM MAC association throughout the service. The **mac-populate** command installs an OAM MAC into the service FIB indicating the device is the egress node for a particular MAC address. The MAC address can be bound to a particular SAP (the **target-sap**) or can be associated with the control plane in that any data destined to the MAC address is forwarded to the control plane (cpm). As a result, if the service on the node has neither a FIB nor an egress SAP, then it is not allowed to initiate a **mac-populate**.

The MAC address that is populated in the FIBs in the provider network is given a type OAM, so that it can be treated distinctly from regular dynamically learned or statically configured MACs. Note that OAM MAC addresses are operational MAC addresses and are not saved in the device configuration. An exec file can be used to define OAM MACs after system initialization.

The **force** option in **mac-populate** forces the MAC in the table to be type OAM in the case it already exists as a dynamic, static or an OAM induced learned MAC with some other type binding.

An OAM-type MAC cannot be overwritten by dynamic learning and allows customer packets with the MAC to either ingress or egress the network while still using the OAM MAC entry.

The **flood** option causes each upstream node to learn the MAC (that is, populate the local FIB with an OAM MAC entry) and to flood the request along the data plane using the flooding domain. The flooded **mac-populate** request can be sent via the data plane or the control plane. The **send-control** option specifies the request be sent using the control plane. If **send-control** is not specified, the request is sent using the data plane.

An **age** can be provided to age a particular OAM MAC using a specific interval. By default, OAM MAC addresses are not aged and can be removed with a **mac-purge** or with an FDB clear operation.

When split horizon group (SHG) is configured, the flooding domain depends on which SHG the packet originates from. The **target-sap sap-id** value dictates the originating SHG information.

VPLS MAC Diagnostics

| | |
|-------------------|--|
| Parameters | service <i>service-id</i> — The Service ID of the service to diagnose or manage.
Values 1 — 2147483647 |
| | destination <i>ieee-address</i> — The MAC address to be populated. |
| | flood — Sends the OAM MAC populate to all upstream nodes.
Default MAC populate only the local FIB. |
| | age <i>seconds</i> — The age for the OAM MAC, expressed as a decimal integer.
Default The OAM MAC does not age.
Values 1 — 65535 |
| | force — Converts the MAC to an OAM MAC even if it currently another type of MAC.
Default Do not overwrite type. |
| | target-sap <i>sap-id</i> — The local target SAP bound to a service on which to associate the OAM MAC. By default, the OAM MAC is associated with the control plane, that is, it is associated with the CPU on the router.
When the target-sap <i>sap-id</i> value is not specified the MAC is bound to the CPM. The originating SHG is 0 (zero). When the target-sap <i>sap-id</i> value is specified, the originating SHG is the SHG of the target-sap.
Default Associate OAM MAC with the control plane (CPU). |

mac-purge

| | |
|--------------------|---|
| Syntax | mac-purge <i>service-id</i> target <i>ieee-address</i> [flood] [register] |
| Context | oam |
| Description | <p>Platforms Supported: 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.</p> <p>This command removes an OAM-type MAC entry from the FIB and optionally floods the OAM MAC removal throughout the service. A mac-purge can be sent via the forwarding path or via the control plane. When sending the MAC purge using the data plane, the TTL in the VC label is set to 1. When sending the MAC purge using the control plane, the packet is sent directly to the system IP address of the next hop.</p> <p>A MAC address is purged only if it is marked as OAM. A mac-purge request is an HVPLS OAM packet, with the following fields. The Reply Flags is set to 0 (since no reply is expected), the Reply Mode and Reserved fields are set to 0. The Ethernet header has source set to the (system) MAC address, the destination set to the broadcast MAC address. There is a VPN TLV in the FEC Stack TLV to identify the service domain.</p> <p>If the register option is provided, the R bit in the Address Delete flags is turned on.</p> <p>The flood option causes each upstream node to be sent the OAM MAC delete request and to flood the request along the data plane using the flooding domain. The flooded mac-purge request can be sent via the</p> |

data plane or the control plane. The **send-control** option specifies the request be sent using the control plane. If **send-control** is not specified, the request is sent using the data plane.

The **register** option reserves the MAC for OAM testing where it is no longer an active MAC in the FIB for forwarding, but it is retained in the FIB as a registered OAM MAC. Registering an OAM MAC prevents relearns for the MAC based on customer packets. Relearning a registered MAC can only be done through a **mac-populate** request. The originating SHG is always 0 (zero).

Parameters **service** *service-id* — The service ID of the service to diagnose or manage.

Values 1 — 2147483647

target *ieee-address* — The MAC address to be purged.

flood — Sends the OAM MAC purge to all upstream nodes.

Default MAC purge only the local FIB.

register — Reserve the MAC for OAM testing.

Default Do not register OAM MAC.

mac-ping

Syntax **mac-ping service** *service-id* **destination** *dst-ieee-address* [**source** *src-ieee-address*] [**fc** *fc-name*] [**size** *octets*] [**tll** *vc-label-ttl*] [**count** *send-count*] [**return-control**] [**interval** *interval*] [**timeout** *timeout*]

Context oam
config>saa>test>type

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

The **mac-ping** utility is used to determine the existence of an egress SAP binding of a given MAC within a VPLS service.

A **mac-ping** packet can be sent through the control plane or the data plane. The **send-control** option specifies the request be sent using the control plane. If **send-control** is not specified, the request is sent using the data plane.

A **mac-ping** is forwarded along the flooding domain if no MAC address bindings exist. If MAC address bindings exist, then the packet is forwarded along those paths, provided they are active. A response is generated only when there is an egress SAP binding for that MAC address or if the MAC address is a “local” OAM MAC address associated with the device’s control plan.

A **mac-ping** reply can be sent using the data plane or the control plane. The **return-control** option specifies the reply be sent using the control plane. If **return-control** is not specified, the request is sent using the data plane.

A **mac-ping** with data plane reply can only be initiated on nodes that can have an egress MAC address binding. A node without a FIB and without any SAPs cannot have an egress MAC address binding, so it is not a node where replies in the data plane will be trapped and sent up to the control plane.

A control plane request is responded to through a control plane reply only.

By default, MAC OAM requests are sent with the system or chassis MAC address as the source MAC. The **source** option allows overriding of the default source MAC for the request with a specific MAC address.

When a **source** *ieee-address* value is specified and the source MAC address is locally registered within a split horizon group (SHG), then this SHG membership will be used as if the packet originated from this SHG. In all other cases, SHG 0 (zero) will be used. Note that if the **mac-trace** is originated from a non-zero SHG, such packets will not go out to the same SHG.

If EMG is enabled, mac-ping will return only the first SAP in each chain.

Parameters

service *service-id* — The service ID of the service to diagnose or manage.

Values 1 — 2147483647

destination *ieee-address* — The destination MAC address for the OAM MAC request.

size *octets* — The MAC OAM request packet size in octets, expressed as a decimal integer. The request payload is padded to the specified size with a 6 byte PAD header and a byte payload of 0xAA as necessary. If the octet size specified is less than the minimum packet, the minimum sized packet necessary to send the request is used.

Default No OAM packet padding.

Values 1 — 9198

ttl *vc-label-ttl* — The TTL value in the VC label for the OAM MAC request, expressed as a decimal integer.

Default 255

Values 1 — 255

return-control — Specifies the MAC OAM reply to a data plane MAC OAM request be sent using the control plane instead of the data plane.

Default MAC OAM reply sent using the data plane.

source *src-ieee-address* — The source MAC address from which the OAM MAC request originates. By default, the system MAC address for the chassis is used.

Default The system MAC address.

Values Any unicast MAC value.

fc *fc-name* — The **fc** parameter is used to test the forwarding class of the MPLS echo request packets. The actual forwarding class encoding is controlled by the network egress LSP-EXP mappings.

Values be, l2, af, l1, h2, ef, h1, nc

interval *interval* — The **interval** parameter in seconds, expressed as a decimal integer. This parameter is used to override the default request message send interval and defines the minimum amount of time that must expire before the next message request is sent.

If the **interval** is set to 1 second where the **timeout** value is set to 10 seconds, then the maximum time between message requests is 10 seconds and the minimum is 1 second. This depends upon the receipt of a message reply corresponding to the outstanding message request.

Default 1

Values 1 — 10

count *send-count* — The number of messages to send, expressed as a decimal integer. The **count** parameter is used to override the default number of message requests sent. Each message request must either time-out or receive a reply before the next message request is sent. The message **interval** value must be expired before the next message request is sent.

Default 1

Values 1 — 100

timeout *timeout* — The **timeout** parameter in seconds, expressed as a decimal integer. This value is used to override the default **timeout** value and is the amount of time that the router will wait for a message reply after sending the message request. Upon the expiration of message timeout, the requesting router assumes that the message response will not be received. Any response received after the request times out will be silently discarded.

Default 5

Values 1 — 10

mac-trace

Syntax **mac-trace service** *service-id* **destination** *ieee-address* [**fc** *fc-name*] [**size** *octets*] [**min-ttl** *vc-label-ttl*] [**max-ttl** *vc-label-ttl*] [**return-control**] [**source** *ieee-address*] [**send-count** *send-count*] [**interval** *interval*] [**timeout** *timeout*]

Context oam
config>saa>test>type

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command displays the hop-by-hop path for a destination MAC address within a VPLS.

The MAC traceroute operation is modeled after the IP traceroute utility which uses ICMP echo request and reply packets with increasing TTL values to determine the hop-by-hop route to a destination IP. The MAC traceroute command uses Alcatel-Lucent OAM packets with increasing TTL values to determine the hop-by-hop route to a destination MAC.

In a MAC traceroute, the originating device creates a MAC ping echo request packet for the MAC to be tested with increasing values of the TTL. The echo request packet is sent through the control plane or data plane and awaits a TTL exceeded response or the echo reply packet from the device with the destination MAC. The devices that reply to the echo request packets with the TTL exceeded and the echo reply are displayed.

When a **source** *ieee-address* value is specified and the source MAC address is locally registered within a split horizon group (SHG), then this SHG membership will be used as if the packet originated from this SHG. In all other cases, SHG 0 (zero) will be used. Note that if the **mac-ping** is originated from a non-zero SHG, such packets will not go out to the same SHG.

If EMG is enabled, mac-trace will return only the first SAP in each chain.

Parameters **service** *service-id* — The Service ID of the service to diagnose or manage.

Values 1 — 2147483647

destination *ieee-address* — The destination MAC address to be traced.

size *octets* — The MAC OAM request packet size in octets, expressed as a decimal integer. The request payload is padded to the specified size with a 6 byte PAD header and a byte payload of 0xAA as necessary. If the octet size specified is less than the minimum packet, the minimum sized packet necessary to send the request is used.

Default No OAM packet padding.

Values 1 — 9198

fc *fc-name* — The **fc** parameter is used to test the forwarding class of the MPLS echo request packets. The actual forwarding class encoding is controlled by the network egress LSP-EXP mappings.

Values be, l2, af, l1, h2, ef, h1, nc

min-ttl *vc-label-ttl* — The minimum TTL value in the VC label for the MAC trace test, expressed as a decimal integer.

Default 1

Values 1 — 255

max-ttl *vc-label-ttl* — The maximum TTL value in the VC label for the MAC trace test, expressed as a decimal integer.

Default 4

Values 1 — 255

return-control — Specifies the MAC OAM reply to a data plane MAC OAM request be sent using the control plane instead of the data plane.

Default MAC OAM reply sent using the data plane.

source *ieee-address* — The source MAC address from which the OAM MAC request originates. By default, the system MAC address for the chassis is used.

Default The system MAC address.

Values Any unicast MAC value.

send-count *send-count* — The number of MAC OAM requests sent for a particular TTL value, expressed as a decimal integer.

Default 1

Values 1 — 10

interval *interval* — The **interval** parameter in seconds, expressed as a decimal integer. This parameter is used to override the default request message send interval and defines the minimum amount of time that must expire before the next message request is sent.

If the **interval** is set to 1 second, and the **timeout** value is set to 10 seconds, then the maximum time between message requests is 10 seconds and the minimum is 1 second. This depends upon the receipt of a message reply corresponding to the outstanding message request.

Default 1

Values 1 — 10

timeout *timeout* — The **timeout** parameter in seconds, expressed as a decimal integer. This value is used to override the default **timeout** value and is the amount of time that the router will wait for a message reply after sending the message request. Upon the expiration of message timeout, the requesting router assumes that the message response will not be received. Any response received after the request times out will be silently discarded.

Default 5

Values 1 — 60

EFM Commands

efm

Syntax *port-id*

Context oam>efm

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command enables Ethernet in the First Mile (EFM) OAM tests loopback tests on the specified port. The EFM OAM remote loopback OAMPDU will be sent to the peering device to trigger remote loopback.

Parameters *port-id* — Specify the port ID in the slot/mda/port format.

local-loopback

Syntax **local-loopback {start | stop}**

Context oam>efm

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command enables local loopback tests on the specified port.

remote-loopback

Syntax **remote-loopback {start | stop}**

Context oam>efm

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command enables remote Ethernet in the First Mile (EFM) OAM loopback tests on the specified port. The EFM OAM remote loopback OAMPDU will be sent to the peering device to trigger remote loopback.

ETH-CFM OAM Commands

linktrace

Syntax `linktrace mac-address mep mep-id domain md-index association ma-index [ttl ttl-value]`

Context `oam>eth-cfm`

Default **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

The command specifies to initiate a linktrace test.

Parameters *mac-address* — Specifies a unicast destination MAC address.

mep *mep-id* — Specifies the target MAC address.

Values 1 — 8191

domain *md-index* — Specifies the MD index.

Values 1 — 4294967295

association *ma-index* — Specifies the MA index.

Values 1 — 4294967295

ttl *ttl-value* — Specifies the TTL for a returned linktrace.

Values 0 — 255

Default 64

loopback

Syntax `loopback mac-address mep mep-id domain md-index association ma-index [send-count send-count] [size data-size] [priority priority]`

Context `oam>eth-cfm`

Default **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

The command specifies to initiate a loopback test.

Parameters *mac-address* — Specifies a unicast MAC address.

mep *mep-id* — Specifies target MAC address.

Values 1 — 8191

ETH-CFM OAM Commands

domain *md-index* — Specifies the MD index.

Values 1 — 4294967295

association *ma-index* — Specifies the MA index.

Values 1 — 4294967295

send-count *send-count* — Specifies the number of messages to send, expressed as a decimal integer. Loop-back messages are sent back to back, with no delay between the transmissions.

Default 1

Values 1 — 5

size *data-size* — This is the size of the data portion of the data TLV. If 0 is specified no data TLV is added to the packet.

Values 0 — 1500

priority *priority* — Specifies a 3-bit value to be used in the VLAN tag, if present, in the transmitted frame.

Values 0 — 7

eth-test

Syntax *mac-address mep mep-id domain md-index association ma-index* [**priority** *priority*] [**data-length** *data-length*]

Context oam>eth-cfm

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command issues an ETH-CFM test.

Parameters *mac-address* — Specifies a unicast MAC address.

mep *mep-id* — Specifies target MAC address.

Values 1 — 8191

domain *md-index* — Specifies the MD index.

Values 1 — 4294967295

association *ma-index* — Specifies the MA index.

Values 1 — 4294967295

data-length *data-length* — Indicates the UDP data length of the echo reply, the length starting after the IP header of the echo reply.

Values 64 — 1500

Default 64

priority *priority* — Specifies the priority.

Values 0—7

Default The CCM and LTM priority of the MEP

one-way-delay-test

Syntax **one-way-delay-test** *mac-address* **mep** *mep-id* **domain** *md-index* **association** *ma-index* [**priority** *priority*]

Context oam>eth-cfm

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command issues an ETH-CFM one-way delay test.

Parameters *mac-address* — Specifies a unicast MAC address.

mep *mep-id* — Specifies target MAC address.

Values 1 — 8191

domain *md-index* — Specifies the MD index.

Values 1 — 4294967295

association *ma-index* — Specifies the MA index.

Values 1 — 4294967295

priority *priority* — Specifies the priority.

Values 0 — 7

Default The CCM and LTM priority of the MEP.

two-way-delay-test

Syntax **two-way-delay-test** *mac-address* **mep** *mep-id* **domain** *md-index* **association** *ma-index* [**priority** *priority*]

Context oam>eth-cfm

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command issues an ETH-CFM two-way delay test.

Parameters *mac-address* — Specifies a unicast MAC address.

mep *mep-id* — Specifies target MAC address.

Values 1 — 8191

domain *md-index* — Specifies the MD index.

Values 1 — 4294967295

association *ma-index* — Specifies the MA index.

Values 1 — 4294967295

priority *priority* — Specifies the priority.

Values 0 — 7

Default The CCM and LTM priority of the MEP.

two-way-slm-test

Syntax **two-way-slm-test** *mac-address* **mep** *mep-id* **domain** *md-index* **association** *ma-index* [**priority** *priority*] [**send-count** *send-count*] [**size** *data-size*] [**timeout** *timeout*] [**interval** *interval*]

Context oam>eth-cfm

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command configures an Ethernet CFM two-way SLM test in SAA.

mac-address — Specifies a unicast destination MAC address.

mep *mep-id* — Specifies the target MAC address.

Values 1 — 8191

domain *md-index* — Specifies the MD index.

Values 1 — 4294967295

association *ma-index* — Specifies the MA index.

Values 1 — 4294967295

priority *priority* — Specifies a 3-bit value to be used in the VLAN tag, if present, in the transmitted frame.

send-count *send-count* — The number of messages to send, expressed as a decimal integer. The **count** parameter is used to override the default number of message requests sent. Each message request must either timeout or receive a reply before the next message request is sent. The message **interval** value must be expired before the next message request is sent.

Default 1

Values 1 — 1000

size *data-size* — This is the size of the data portion of the data TLV. If 0 is specified no data TLV is added to the packet.

Default 0

Values 0 — 1500

timeout *timeout* — The **timeout** parameter in seconds, expressed as a decimal integer. This value is used to override the default **timeout** value and is the amount of time that the router waits for a reply message after sending the message request. Upon the expiration of message timeout, the requesting router assumes that the message response is not received. Any response received after the request times out is silently discarded.

Default 5

Values 1 — 10

interval *interval* — The **interval** parameter in seconds, expressed as a decimal integer. This parameter is used to override the default request message send interval and defines the minimum amount of time that must expire before the next message request is sent.

If the **interval** is set to 1 second, and the **timeout** value is set to 10 seconds, then the maximum time between message requests is 10 seconds and the minimum is 1 second. This depends upon the receipt of a message reply corresponding to the outstanding message request.

Default 5

Values [0.1, 0.2, .. 0.9] | [1, 2, .. 10]

Testhead Commands

test-oam

Syntax `test-oam`

Context `config`

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command enables the context to configure Operations, Administration, and Maintenance test parameters

testhead-profile

Syntax `testhead-profile profile-id create`

Context `config> test-oam`

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

Provides the context to the create service testhead profiles which is used by the Y.1564/RFC 2544 testhead (also known as, traffic generator) OAM tool. A service testhead profile allows user to configure the parameters such as contents of the frame payload that is generated by traffic generator, the size of the frame, test duration, test acceptance criteria, and other criteria to be used by the testhead tool.

The profile is used the testhead OAM tool to generate the appropriate frame at the configured rate and measure the performance parameters (FD, FDV, and loss). At the end of the test run, the tool compares the measured values against the test acceptance criteria that is configured in the profile to determine if the service is within bounds of the acceptance criteria or not.

The no form the command removes user created profile from the system.

Default `none`

Parameters `profile-id` — Identifies the profile.

Values `1-10`

description

| | |
|--------------------|---|
| Syntax | description <i>profile-description</i> |
| Context | config> test-oam>testhead-profile |
| Description | <p>Platforms Supported: 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.</p> <p>Allows user to associate a description with profile.</p> <p>The no form the command removes description.</p> |
| Default | none |
| Parameters | <p><i>profile-description</i> — Provides a way to add a description to the profile based on it use or as per user choice.</p> <p>Values ASCII string</p> |

rate

| | |
|--------------------|---|
| Syntax | [no] rate cir-rate-in-kbps [cir-adaptation-rule adaptation-rule] [pir cir-rate-in-kbps] |
| Context | config> test-oam>testhead-profile |
| Description | <p>Platforms Supported: 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.</p> <p>The testhead tool generates traffic up to the configured CIR rate, if the PIR rate is not specified. In other words, CIR rate specifies the bandwidth or throughput the user needs to validate. User can specify the optional PIR rate. If the PIR rate is specified, it needs to be greater than or equal to the CIR rate and the test-head generates traffic up to the configured PIR rate.</p> <p>The <i>cir-adaptation-rule</i> parameter rule can be specified to let the system derive the operational hardware rate for both the CIR and PIR rate. This allows the software to find the best operational rate based on the user specified constraint and the hardware based rate steps supported on the platform. For more information about the hardware rate steps supported for meters on different platforms, see the “7210 SAS QoS User Guide”.</p> <p>The no form of the command sets the CIR to default and PIR is not set. If the test is run after executing no rate command, the test generates traffic up to CIR rate. If PIR rate is specified, it must be greater than or equal to the CIR rate.</p> |
| Default | rate cir 1000kbps adaptation-rule closest |
| Parameters | <p><i>cir-rate</i> — The cir parameter overrides the default administrative CIR to use. When the rate command has not been executed or the cir parameter is not explicitly specified, the default CIR is assumed. Fractional values are not allowed and must be given as a positive integer. The actual CIR rate is dependent on the meter’s adaptation-rule parameters and the hardware. It is specified in kilo-bits per second (kbps).</p> <p>Values 0 — 10000000, max</p> |

adaptation-rule — Defines the constraints enforced when adapting the CIR and PIR rate defined with the rate command to the hardware rates supported by the platform. This parameter requires a qualifier that defines the constraint used when deriving the operational CIR and PIR value. If this parameter is not specified then the default adaptation-rule closest is applied.

Values [closest|max|min]

max - The max (maximum) option is mutually exclusive with the min and closest options. When max is defined, the operational CIR will be the next multiple of the hardware step-size that is equal to or lesser than the specified rate. The hardware step-size is determined by the configured administrative CIR and varies based on the platform.

min - The min (minimum) option is mutually exclusive with the max and closest options. When min is defined, the operational PIR will be the next multiple of the hardware step-size that is equal to or lesser than the specified rate. The hardware step-size is determined by the configured administrative CIR and varies based on the platform.

closest - The closest parameter is mutually exclusive with the min and max parameter. When closest is defined, the operational CIR will be the next multiple of the hardware step-size that is equal to or lesser than the specified rate. The hardware step-size is determined by the configured administrative CIR and varies based on the platform.

pir-rate — The pir parameter overrides the default administrative PIR to use. When the rate command has not been executed or the PIR parameter is not explicitly specified, the default PIR is assumed. Fractional values are not allowed and must be given as a positive integer. The actual PIR rate is dependent on the meter's adaptation-rule parameters and the hardware. The value is specified in kilo-bits per second (kbps).

Values 0 — 10000000, max

test-duration

Syntax test-duration {[hours hours] [minutes minutes] [seconds seconds]}

Context config> test-oam>testhead-profile

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command allows the user to specify the total test duration to be used for throughput measurement. The CLI parameters, hours, minutes, and seconds, allows the user to specify the number of hours, number of minutes and number of seconds to used for throughput measurement. User can specify all the parameters together. If all the parameters are specified together then the total test duration is set to the sum of the values specified for hours, minutes and seconds.

The no form of the command sets the value to the default value

Default no test-duration (sets the test duration for 3 minutes).

Parameters *hours* — The total number of hours to run the test. The total test duration is determined by the sum of the hours, minutes and seconds specified by the user.

Values 0 - 24

minutes — The total number of minutes to run the test. The total test duration is determined by the sum of the hours, minutes and seconds specified by the user.

Values 0 — 60

seconds — The total number of seconds to run the test. The total test duration is determined by the sum of the hours, minutes and seconds specified by the user.

Values 0 — 60

frame-size

Syntax [no] frame-size [64..9212]

Context config> test-oam>testhead-profile

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command allows the user to specify the frame size of the packets generated by the testhead tool. Any frame size in the given range can be specified.

The no form of the command sets the value to the default value

Default no frame-size - set to a default value of 1514 bytes.

Parameters *frame-size* — The size of the frame generated by the testhead tool. Choose from among the value allowed in the available range.

Values 64 ... 9212

acceptance-criteria

Syntax [no] acceptance-criteria *acceptance-criteria-id* create

Context configure> test-oam> testhead-profile

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

Testhead Commands

This command provides the context to specify the test acceptance criteria to be used by the testhead OAM tool to declare the PASS/FAIL result at the completion of the test.

User can create upto 4 different acceptance criteria per profile to measure different SLA needs. User has an option to specify only one of the acceptance criteria to be specified with the testhead OAM tool during the invocation of the test.

The no form of the command removes the test acceptance criteria.

Default no defaults

Parameters *acceptance-criteria-id* — A number to identify the test acceptance criteria. It is a decimal number used to identify the test acceptance criteria and to use when starting the throughput test.

Values 1- 4

cir-threshold

Syntax **[no]** *cir-threshold* *cir-threshold*

Context configure> test-oam> testhead-profile> acceptance-criteria

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

The specified value for the CIR rate is compared with the measured CIR rate at the end of the test to declare the test result. If the measured value is greater than the specified value the test is declared as 'PASS', else it is considered to be 'FAIL'.

The no form of the command disables the comparison of the parameter with the measured value at the end of the test. Basically, the threshold value is ignored and not considered for declaring the test result.

Default no cir-threshold

Parameters *threshold* — Specifies the value for comparison with measured value

Values 0 – 1000000kbps

pir-threshold

Syntax **[no]** *pir-threshold* *pir-threshold*

Context configure> test-oam> testhead-profile> acceptance-criteria

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

The specified value for the PIR rate is compared with the measured PIR rate at the end of the test to declare the test result. If the measured value is greater than the specified value the test is declared as 'PASS', else it is considered to be 'FAIL'.

The no form of the command disables the comparison of the parameter with the measured value at the end of the test. Basically, the threshold value is ignored and not considered for declaring the test result.

Default no pir-threshold

Parameters *threshold* — Specifies the value for comparison with measured value

Values 0 – 1000000kbps

latency-rising-threshold

Syntax **[no] latency-rising-threshold** *threshold*

Context configure> test-oam> testhead-profile> acceptance-criteria

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

The specified value for the latency is compared with the measured latency at the end of the test to declare the test result. If the measured value is greater than the specified value the test is declared as 'FAIL', else it is considered to be 'PASS'.

The no form of the command disables the comparison of the parameter with the measured value at the end of the test. Basically, the threshold value is ignored and not considered for declaring the test result.

Default no latency-rising-threshold

Parameters *threshold* — Specifies the value for comparison with measured value.

Values [0..2147483000], Specified in microseconds.

latency-rising-threshold-in

Syntax **[no] latency-rising-threshold-in** *in-profile-threshold*

Context configure> test-oam> testhead-profile> acceptance-criteria

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

The specified value for the latency is compared with the measured latency for green/in-profile packets at the end of the test to declare the test result. If the measured value is greater than the specified value the test is declared as 'FAIL', else it is considered to be 'PASS'.

The no form of the command disables the comparison of the parameter with the measured value at the end of the test. Basically, the threshold value is ignored and not considered for declaring the test result.

Default no latency-rising-threshold-in

Parameters *In-profile-threshold* — Specifies the value for comparison with measured value

Values [0..2147483000], Specified in microseconds.

latency-rising-threshold-out

| | |
|--------------------|--|
| Syntax | [no] latency-rising-threshold out-profile-threshold |
| Context | configure> test-oam> testhead-profile> acceptance-criteria |
| Description | <p>Platforms Supported: 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.</p> <p>The specified value for the latency is compared with the measured latency of yellow or out-of-profile packets at the end of the test to declare the test result. If the measured value is greater than the specified value the test is declared as 'FAIL', else it is considered to be 'PASS'.</p> <p>The no form of the command disables the comparison of the parameter with the measured value at the end of the test. Basically, the threshold value is ignored and not considered for declaring the test result.</p> |
| Default | no latency-rising-threshold-out |
| Parameters | <p><i>out-profile-threshold</i> — Specifies the value for comparison with measured value</p> <p>Values [0..2147483000], Specified in microseconds.</p> |

jitter-rising-threshold

| | |
|--------------------|--|
| Syntax | [no] jitter-rising-threshold <i>threshold</i> |
| Context | configure> test-oam> testhead-profile> acceptance-criteria |
| Description | <p>Platforms Supported: 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.</p> <p>The specified value for the jitter is compared with the measured jitter at the end of the test to declare the test result. If the measured value is greater than the specified value the test is declared as 'FAIL', else it is considered to be 'PASS'.</p> <p>The no form of the command disables the comparison of the parameter with the measured value at the end of the test. Basically, the threshold value is ignored and not considered for declaring the test result.</p> |
| Default | no jitter-rising-threshold |
| Parameters | <p><i>threshold</i> — Specifies the value for comparison with measured value.</p> <p>Values [0..2147483000], Specified in microseconds.</p> |

jitter-rising-threshold-in

| | |
|--------------------|---|
| Syntax | [no] jitter-rising-threshold-in <i>in-profile-threshold</i> |
| Context | configure> test-oam> testhead-profile> acceptance-criteria |
| Description | <p>Platforms Supported: 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.</p> <p>The specified value for the jitter is compared with the measured jitter for green/in-profile packets at the end of the test to declare the test result. If the measured value is greater than the specified value the test is declared as 'FAIL', else it is considered to be 'PASS'.</p> <p>The no form of the command disables the comparison of the parameter with the measured value at the end of the test. Basically, the threshold value is ignored and not considered for declaring the test result.</p> |
| Default | no jitter-rising-threshold-in |
| Parameters | <p><i>In-profile-threshold</i> — Specifies the value for comparison with measured value</p> <p>Values [0..2147483000], Specified in microseconds.</p> |

jitter-rising-threshold-out

| | |
|--------------------|--|
| Syntax | [no] jitter-rising-threshold-out <i>out-profile-threshold</i> |
| Context | configure> test-oam> testhead-profile> acceptance-criteria |
| Description | <p>Platforms Supported: 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.</p> <p>The specified value for the jitter is compared with the measured jitter for yellow/out-of-profile packets at the end of the test to declare the test result. If the measured value is greater than the specified value the test is declared as 'FAIL', else it is considered to be 'PASS'.</p> <p>The no form of the command disables the comparison of the parameter with the measured value at the end of the test. Basically, the threshold value is ignored and not considered for declaring the test result.</p> |
| Default | no jitter-rising-threshold-out |
| Parameters | <p><i>out-profile-threshold</i> — Specifies the value for comparison with measured value</p> <p>Values [0..2147483000], Specified in microseconds.</p> |

loss-rising-threshold

| | |
|--------------------|--|
| Syntax | [no] loss-rising-threshold <i>threshold</i> |
| Context | configure> test-oam> testhead-profile> acceptance-criteria |
| Description | <p>Platforms Supported: 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.</p> <p>The specified value for the Frame Loss Ratio (FLR) is compared with the measured FLR at the end of the test to declare the test result. If the measured value is greater than the specified value the test is declared as 'FAIL', else it is considered to be 'PASS'.</p> <p>Frame Loss Ratio is computed as a ratio of the difference of number of received frames to number of injected or sent frames divided by the number of sent frames.</p> <p>The no form of the command disables the comparison of the parameter with the measured value at the end of the test. Basically, the threshold value is ignored and not considered for declaring the test result.</p> |
| Default | no loss-rising-threshold |
| Parameters | <p><i>threshold</i> — Specifies the value for comparison with measured value.</p> <p>Values 1 – 1000000, Loss-rising-threshold is specified as a number which denotes one ten-thousandth (1/10000) of a percent. For example, specifying a value of 1 is equivalent to 0.0001%, and specifying a value of 10000 is equivalent to 1%.</p> |

loss-rising-threshold-in

| | |
|--------------------|---|
| Syntax | [no] loss-rising-threshold-in <i>in-profile-threshold</i> |
| Context | configure> test-oam> testhead-profile> acceptance-criteria |
| Description | <p>Platforms Supported: 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.</p> <p>The specified value for the frame loss ratio (FLR) is compared with the measured FLR for green or in-profile packets at the end of the test to declare the test result. If the measured value is greater than the specified value the test is declared as 'FAIL', else it is considered to be 'PASS'.</p> <p>Frame Loss Ratio for green/in-profile packets is computed as a ratio of the difference of number of received green or in-profile frames to number of injected/sent green/in-profile frames divided by the number of sent green frames.</p> <p>The no form of the command disables the comparison of the parameter with the measured value at the end of the test. Basically, the threshold value is ignored and not considered for declaring the test result.</p> |
| Default | no loss-rising-threshold-in |

| | |
|-------------------|--|
| Parameters | <i>in-profile-threshold</i> — Specifies the value for comparison with measured value |
| Values | 1 – 1000000, Loss-rising-threshold is specified as a number which denotes one ten-thousandth (1/10000) of a percent. For example, specifying a value of 1 is equivalent to 0.0001%, and specifying a value of 10000 is equivalent to 1%. |

loss-rising-threshold-out

| | |
|--------------------|--|
| Syntax | [no] loss-rising-threshold-out <i>out-profile-threshold</i> |
| Context | configure> test-oam> testhead-profile> acceptance-criteria |
| Description | <p>Platforms Supported: 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.</p> <p>The specified value for the frame loss ratio (FLR) is compared with the measured FLR for yellow/out-of-profile packets at the end of the test to declare the test result. If the measured value is greater than the specified value the test is declared as 'FAIL', else it is considered to be 'PASS'.</p> <p>Frame Loss ratio for yellow/out-of-profile packets is computed as a ratio of the difference of number of received yellow frames to number of injected/sent yellow frames divided by the number of sent yellow frames.</p> <p>The no form of the command disables the comparison of the parameter with the measured value at the end of the test. Basically, the threshold value is ignored and not considered for declaring the test result.</p> |
| Default | no loss-rising-threshold |
| Parameters | <i>out-profile-threshold</i> — Specifies the value for comparison with measured value |
| Values | 1 – 1000000, Loss-rising-threshold is specified as a number which denotes one ten-thousandth (1/10000) of a percent. For example, specifying a value of 1 is equivalent to 0.0001%, and specifying a value of 10000 is equivalent to 1%. |

test-completion-trap-enable

| | |
|--------------------|--|
| Syntax | [no] test-completion-trap-enable |
| Context | configure> test-oam> testhead-profile |
| Description | <p>Platforms Supported: 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.</p> <p>Executing this command allows the user to specify that the test completion trap needs to be generated after the completion of the test or if the test is stopped. The trap contains the details of test configuration, the measured values, test completion status and PASS/FAIL result.</p> <p>The no form of the command disables the generation of the event/log/trap after test completion.</p> |
| Default | no test-completion-trap-enable – that is, trap is not generated on completion of the test. |

dot1p

Syntax [no] dot1p in-profile *dot1p-value* out-of-profile *dot1p-value*

Context configure> test-oam> testhead-profile

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command allows the user to configure the Dot1p values to identify the in-profile or green packets and out-of-profile or yellow packets. The values configured using this command are used by the testhead tool on the local end (that is, the node on which the testhead tool is executed) to match the dot1p values received in the packet header and identify green and yellow packets and appropriately account the packets. These values are used only when the testhead tool is invoked with the parameter *color-aware* is set to 'enable'.

The dot1p in-profile value (that is, packets with dot1p values in the L2 header equal to the dot1p-in-profile value configured is considered to be in-profile or green packet) is used to count the number of in-profile packets and measure the latency, jitter, and FLR for in-profile packets. Similarly, the dot1p out-profile is used to count the total out-of-profile or yellow packets and measure latency, jitter, and FLR for out-of-profile or yellow packets.

While the testhead tool is initiated, if color-aware is set to enable and no values are specified (that is, the no form of the command is used in the profile), the CLI gives an error. If values are specified, then the configured values are used to match and identify in-profile and out-of-profile packets.

The no form of the command disables the use of dot1p to identify a green or yellow packet.

Note: Testhead OAM tool does not mark the packets below CIR as in-profile packets and packets above CIR and below PIR as out-of-profile packets using the Dot1p or DSCP or other packet header bits to indicate the color of the packet (for example: DEI bit), as the 7210 SAS access SAP ingress does not support color-aware metering. It is used to only identify green and yellow packets and maintain a count of received green and yellow packets when the tests are run in color-aware mode.

Default The no form of this command is the default. There are no defaults for the dot1p values.

Parameters *in-profile dot1p-value* — Specifies the dot1p value used to identify green or in-profile packets. It must be different than the value configured for yellow or out-of-profile packets.

Values 0-7

out-profile dot1p-value — Specifies the dot1p value used to identify green or out-of-profile packets. It must be different than the value configured for green or in-profile packets.

Values 0-7

frame-payload

Syntax [no] frame-payload *frame-payload-id* [payload-type [l2|tcp-ipv4|udp-ipv4|ipv4] create

Context configure> test-oam> testhead-profile

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command provides the context to specify the packet header values to be used in frames generated by testhead tool.

User can create up to 4 different types of frame payload representing different kinds of traffic, within a profile. User chooses one among these when starting the throughput test.

The parameter payload-type determines the packet header fields that are used to populate the frame generated by the testhead OAM tool. The packet header fields use the value from the parameters configured under the frame-payload. For example, when the payload-type is configured as "l2", software uses the parameters src-mac, dst-mac, vlan-tag-1 (if configured), vlan-tag-2 (if configured), ethertype, and data-pattern. See below for parameters used when other values are specified with payload-type.

The no form of the command removes the frame payload context.

Default no defaults – no frame payload is created by default.

Parameters *frame-payload-id* — A number to identify the frame-payload. it is an integer used to identify the frame type to use when starting the throughput test.

Values 1-4

frame-payload-type — Identifies whether the frame payload is L2 traffic, IP traffic, TCP/IP traffic or UDP/IP traffic and uses appropriate parameters to build the frame to be generated by the testhead OAM tool. It defaults to tcp-ipv4, if the user does not specify the value during creation of the new frame-payload.

Values l2|tcp-ipv4|udp-ipv4|ipv4
 If l2 is specified, use src-mac+dst-mac+vlan-tag-1(if available)+vlan-tag-2 (if available)+ethertype+data-pattern.
 If tcp-ipv4 or udp-ipv4 is specified, use src-mac+dst-mac+vlan-tag-1(if available)+vlan-tag-2 (if available)+ethertype=0x0800+src-ipv4+dst-ipv4+ip-ttl+ip-dscp or ip-tos+TCP/UDP Protocol Number+src-port+dst-port+data-pattern.
 If ipv4 is specified, use src-mac+dst-mac+vlan-tag-1(if available)+vlan-tag-2 (if available)+ethertype=0x0800+src-ipv4+dst-ipv4+ip-ttl+ip-dscp or ip-tos+ip-proto+data-pattern.

description

Syntax [no] *description frame-description*

Context configure> test-oam> testhead-profile> frame-payload

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command allows user to add some description to the frame type created to describe the purpose or identify the usage or any other such purpose.

The no form of the command removes the description.

Default no description

Parameters *frame-description* — it is an ASCII string used to describe the frame.

Values ASCII string

src-mac

Syntax [no] *src-mac mac-address*

Context configure> test-oam> testhead-profile> frame-payload

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

Specifies the value of source MAC address to use in the frame generated by the testhead OAM tool. Only unicast MAC address must be specified.

This value must be specified for all possible values of payload-type.

The no form of the command indicates that the field is not to be used in the frame generated by the tool.

Default no src-mac

Parameters *mac-address* — Specify the unicast source MAC address.

Values It is specified as a hexadecimal string using the notation xx:xx:xx:xx:xx:xx. The values for xx can be in the range 0-9 and a-f.

dst-mac

Syntax [no] *dst-mac mac-address*

Context configure> test-oam> testhead-profile> frame-payload

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

Specifies the value of source MAC address to use in the frame generated by the testhead OAM tool. Only unicast MAC address must be specified.

This value must be specified for all possible values of payload-type.

The no form of the command indicates that the field is not to be used in the frame generated by the tool.

Default no dst-mac

Parameters *mac-address* — Specify the unicast source MAC address.

Values It is specified as a hexadecimal string using the notation xx:xx:xx:xx:xx:xx. The values for xx can be in the range 0-9 and a-f.

vlan-tag-1

Syntax **[no] vlan-tag-1 vlan-id *vlan-id-value* [tpid *tpid value*] [dot1p *dot1p-value*]**

Context configure> test-oam> testhead-profile> frame-payload

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command allows the user to specify the values to be used for the outermost vlan-tag (often called the outer vlan) in the frame generated by the testhead OAM tool. The tool uses the values specified for VLAN ID, dot1p bits and TPID in populating the outermost VLAN tag in the frame generated.

Configuration of this parameter is optional and it is used for all possible values of payload-type, if configured.

The no form of the command indicates that the field is not to be used in the frame generated by the tool.

NOTES:

- User must ensure that TPID/ethertype configured with this command matches the QinQ ethertype value in use on the port on which the test SAP is configured or must match 0x8100 if the test SAP is configured on a Dot1q encapsulation port, for the frame generated by the tool to be processed successfully on SAP ingress. If this value does not match the one configured under the port, frames generated by the testhead will be dropped by the node on SAP ingress due to ethertype mismatch.
- User must ensure that VLAN ID configured with this command matches the outermost VLAN tag of the QinQ SAP or the Dot1q SAP used for the test SAP for the frame generated by the tool to be processed successfully on SAP ingress. If this value does not match the one configured for the SAP, frames generated by the testhead will be dropped by the node on SAP ingress due to VLAN ID mismatch.
- The Dot1p bits specified for the outermost tag can be used for SAP ingress QoS classification.

Default no vlan-tag-1

Testhead Commands

- Parameters** *vlan-id-value* — Specify the VLAN ID to use to populate the VLAN ID value of the VLAN tag. No defaults are chosen and user has to specify a value to use, if they configure this command.
- Values** Values can be in the range 0-4094.
- tpid-value* — Specify the TPID (also known as, ethertype) to use for the VLAN tag addition. It defaults to 0x8100 if user does not specify it.
- Values** Values can be any of the valid ethertype values allowed for use with VLAN tags in the range 0x0600..0xffff.
- Dot1p-value* — Specify the Dot1p value to use to populate the Dot1p bits in the VLAN tag. It defaults to 0, if the user does not specify it.
- Values** Values can be in the range of 0 – 7.

vlan-tag-2

- Syntax** **[no] vlan-tag-2 vlan-id *vlan-id-value* [tpid *tpid value*] [dot1p *dot1p-value*]**
- Context** configure> test-oam> testhead-profile> frame-payload
- Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.
- This command allows the user to specify the values to be used for the second vlan-tag (often called the inner vlan or the C-vlan) in the frame generated by the testhead OAM tool. The tool uses the values specified for VLAN ID, dot1p bits and TPID in populating the second VLAN tag in the frame generated.
- Configuration of this parameter is optional and it is used for all possible values of payload-type, if configured.
- The no form of the command indicates that the field is not to be used in the frame generated by the tool.
- NOTES:**
- User must ensure that TPID/ethertype configured with this command is 0x8100 for the frame generated by the tool to be processed successfully on SAP ingress. If this value does not match 0x8100, frames generated by the testhead will be dropped by the node on SAP ingress due to ethertype mismatch (7210 supports only 0x8100 as the ethertype value for the inner vlan tag).
 - User must ensure that VLAN ID configured with this command matches the outermost VLAN tag of the QinQ SAP or the Dot1q SAP used for the test SAP for the frame generated by the tool to be processed successfully on SAP ingress. If this value does not match the one configured for the SAP, frames generated by the testhead will be dropped by the node on SAP ingress due to VLAN ID mismatch.
 - The Dot1p bits specified for the outermost tag can be used for SAP ingress QoS classification.
- Default** no vlan-tag-2
- Parameters** *vlan-id-value* — Specify the VLAN ID to use to populate the VLAN ID value of the VLAN tag. No defaults are chosen and user has to specify a value to use, if they configure this command.
- Values** Values can be in the range 0-4094.

tpid-value — Specify the TPID (also known as, ethertype) to use for the VLAN tag addition. It defaults to 0x8100 if user does not specify it.

Values Values can be any of the valid ethertype values allowed for use with VLAN tags in the range 0x0600..0xffff.

Dot1p-value — Specify the Dot1p value to use to populate the Dot1p bits in the VLAN tag. It defaults to 0, if the user does not specify it.

Values Values can be in the range of 0 – 7

ethertype

Syntax **[no] ethertype** *ethertype-value*

Context configure> test-oam> testhead-profile> frame-payload

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command allows the user to specify the ethertype of the frame generated by the testhead tool.

This value must be specified if the payload-type is “l2”. The testhead tool uses the value specified with this command only if the payload-type is “l2”. For all other values of payload-type, the ethertype value used in the frame generated by the testhead tool uses specific value based on the payload-type. See the frame-payload CLI description for more information.

The no form of the command indicates that the field is not to be used in the frame generated by the tool.

Default no ethertype, if the payload-type is set to l2, else the values used depends on the payload-type specified.

Parameters *ethertype-value* — Specify the frame payload ethertype value.

Values Valid ethertype values specified in the range 0x0600..0xffff, as hexadecimal string.

src-ip

Syntax **[no] src-ip ipv4** *ipv4-address*

Context configure> test-oam> testhead-profile> frame-payload

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command allows the user to specify the source IPv4 address to use in the IP header for the frame generated by the testhead tool.

This value must be specified if the payload-type is configured as ipv4 or tcp-ipv4 or udp-ipv4. The testhead tool does not use the value specified with this command if the payload-type is “l2”.

Testhead Commands

The no form of the command indicates that the field is not to be used in the frame generated by the tool.

Default no src-ip, if the payload-type is set to ipv4, tcp-ipv4, udp-ipv4.

Parameters *ipv4-address* — Specify the IPv4 source IP address to use in the IP header

Values Valid IPv4 address specified in dotted decimal format (that is, a.b.c.d) where a, b, c, d are decimal values in the range1-255

dst-ip

Syntax **[no] dst-ip ipv4 *ipv4-address***

Context configure> test-oam> testhead-profile> frame-payload

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command allows the user to specify the destination IPv4 address to use in the IP header for the frame generated by the testhead tool.

This value must be specified if the payload-type is configured as ipv4 or tcp-ipv4 or udp-ipv4. The testhead tool does not use the value specified with this command if the payload-type is “12”.

The no form of the command indicates that the field is not to be used in the frame generated by the tool.

Default no dst-ip, if the payload-type is set to ipv4, tcp-ipv4, udp-ipv4.

Parameters *ipv4-address* — Specify the IPv4 destination IP address to use in the IP header

Values Valid IPv4 address specified in dotted decimal format (that is, a.b.c.d) where a, b, c, d are decimal values in the range1-255.

ip-proto

Syntax **[no] ip-proto *ip-protocol-number***

Context configure> test-oam> testhead-profile> frame-payload

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command allows the user to specify the IP protocol value to use in the IP header for the frame payload generated by the testhead tool.

This value must be specified if the payload-type is configured as ipv4. If the payload-type is specified as tcp-ipv4 or udp-ipv4, the appropriate standard defined values are used. The testhead tool does not use the value specified with this command if the payload-type is “12”.

The no form of the command indicates that the field is not to be used in the frame generated by the tool.

Default no ip-PROTO

Parameters *ip-protocol-number* — Specify the IP-protocol number to use in the IP header.

Values Valid IP protocol number specified as a decimal number in the range 0-255.

dscp

Syntax [**no**] **dscp** *dscp-name*

Context configure> test-oam> testhead-profile> frame-payload

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command allows the user to specify the IP DSCP value to use in the IP header for the frame generated by the testhead tool.

This value can be specified if the payload-type is configured as ipv4 or tcp-ipv4 or udp-ipv4 and if configured is used by the testhead tool to populate the IP DSCP field of the IP header. If it is not specified it defaults to 0 when the payload type is ipv4, tcp-ipv4, and udp-ipv4. The testhead tool does not use the value specified with this command if the payload-type is "12".

NOTES:

- If both IP DSCP and IP ToS is configured them IP DSCP take precedence.
- If IP DSCP is not configured but IP ToS is configured, then IP ToS value is used.

The no form of the command indicates that the field is not to be used in the frame generated by the tool.

Default no dscp

Parameters *dscp-name* — Specify the IPv4 DSCP value to use in the IP header.

Values Valid values from the list of DSCP names.
 be|ef|cp1|cp2|cp3|cp4|cp5|cp6|cp7|cp9|cs1|cs2|cs3|cs4|cs5|nc1|nc2|af11|af12|af13|af21|af22|af23|af31|af32|af33|af41|af42|af43|cp11|cp13|cp15|cp17|cp19|cp21|cp23|cp25|cp27|cp29|cp31|cp33|cp35|cp37|cp39|cp41|cp42|cp43|cp44|cp45|cp47|cp49|cp50|cp51|cp52|cp53|cp54|cp55|cp57|cp58|cp59|cp60|cp61|cp62|cp63

ip-ttl

Syntax [no] ip-ttl *ttl-value*

Context configure> test-oam> testhead-profile> frame-payload

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command allows the user to specify the IP TTL (Time-to-Live) value to use in the IP header for the frame generated by the testhead tool.

This value can be specified if the payload-type is configured as ipv4 or tcp-ipv4 or udp-ipv4 and if configured is used by the testhead tool to populate the IP TTL field of the IP header. If it is not specified it defaults to 1 when the payload type is ipv4, tcp-ipv4, and udp-ipv4. The testhead tool does not use the value specified with this command if the payload-type is “12”.

The no form of the command indicates that the field is not to be used in the frame generated by the tool.

Default no ip-ttl

Parameters *ttl-value* — Specify the IP TTL value to use in the IP header.

Values Specified as a decimal number in the range 1-255.

ip-tos

Syntax [no] ip-tos *type-of-service*

Context configure> test-oam> testhead-profile> frame-payload

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command allows the user to specify the IP TOS (Type of Service) value to use in the IP header for the frame generated by the testhead tool.

This value can be specified if the payload-type is configured as ipv4 or tcp-ipv4 or udp-ipv4 and if configured is used by the testhead tool to populate the IP TOS field of the IP header. If it is not specified it defaults to 0 when the payload type is ipv4, tcp-ipv4, and udp-ipv4. The testhead tool does not use the value specified with this command if the payload-type is “12”.

NOTES:

- If both IP DSCP and IP ToS is configured them IP DSCP take precedence.
- If IP DSCP is not configured but IP ToS is configured, then IP ToS value is used.

The no form of the command indicates that the field is not to be used in the frame generated by the tool.

Default no ip-tos

Parameters *type-of-service* — Specify the value of ToS bits to use in the IP header.

Values Valid number in the range 0-8.

src-port

Syntax [no] src-port *src-port-number*

Context configure> test-oam> testhead-profile> frame-payload

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command allows the user to specify the source port to use in the TCP header for the frame generated by the testhead tool.

This value must be specified if the payload-type is configured as tcp-ipv4 or udp-ipv4. The testhead tool does not use the value specified with this command if the payload-type is l2 or ipv4.

The no form of the command indicates that the field is not to be used in the frame generated by the tool.

Default no src-port, if the payload-type is set to tcp-ipv4 or udp-ipv4

Parameters *src-port-number* — Specify the source TCP/UDP port number to use in the frame's TCP/UDP header.

Values Valid TCP/UDP port number specified in decimal or hexadecimal in the range 0-65535.

dst-port

Syntax [no] dst-port

Context configure> test-oam> testhead-profile> frame-payload

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command allows the user to specify the destination port to use in the TCP header for the frame generated by the testhead tool.

This value must be specified if the payload-type is configured as tcp-ipv4 or udp-ipv4. The testhead tool does not use the value specified with this command if the payload-type is l2 or ipv4.

The no form of the command indicates that the field is not to be used in the frame generated by the tool.

Default no dst-port, if the payload-type is set to tcp-ipv4 or udp-ipv4

Parameters *dst-port-number* — Specify the destination TCP/UDP port number to use in the frame's TCP/UDP header.

Values Valid TCP/UDP port number specified in decimal or hexadecimal in the range 0-65535.

data-pattern

Syntax [no] **data-pattern** *data-pattern*

Context configure> test-oam> testhead-profile> frame-payload

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command allows the user to specify the data pattern to populate the payload portion of the frame generated by the testhead tool.

This value can be specified if the payload-type is configured as l2 or ipv4 or tcp-ipv4 or udp-ipv4. For all these payload types, the frame with the appropriate headers is created and the payload portion of the frame, is filled up with the data-pattern-value specified with this command, repeating it as many times as required to fill up the remaining length of the payload.

The no form of the command uses the default data-pattern value of 0xa1b2c3d4e5f6.

Default no data-pattern

Parameters *data-pattern* — Used to specify the data-pattern to fill the payload data.

Values A string of decimal or hexadecimal numbers of length in the range 1-64.

OAM testhead commands

testhead

Syntax `testhead test-name owner owner-name testhead-profile profile-id [frame-payload frame-payload-id] [acceptance-criteria acceptance-criteria-id] [color-aware enable|disable] sap sap-id [fc fc-name]`

Context oam

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command allows the user to execute the throughput test by generating the traffic up to the configured rate and measures the delay, delay-variation and frame-loss ratio. At the end of the test run the testhead command compares the measured values against the test acceptance criteria that is specified to determine if the service is within bounds of the acceptance criteria or not. It declares the test to have PASSED if the configured rate thresholds are achieved and the measured performance parameters (that is, latency, jitter, and FLR) values are lesser than the thresholds configured in the acceptance criteria. It reports a FAILURE, if the configured rate thresholds are not achieved or if any of the measured values for the performance parameters exceeds the thresholds configured in the acceptance criteria.

The user must specify the testhead-profile parameter to use. This profile parameter determines the rate at which traffic is generated and the content of the frames used for traffic generation. If both CIR and PIR is specified or if only PIR is specified (by setting CIR to zero), the tool generates traffic up to the configured PIR rate. If only CIR is specified the tool generates traffic up to the configured CIR rate.

If the acceptance-criteria parameter is not specified and color-aware is set to disable, then by default software will display the test result as “PASS”, if the frame loss is zero and desired rate is achieved. For comparison with measured rate, the test uses the configured CIR rate, if only CIR is configured or it uses the PIR rate, if either only PIR rate is specified or if both CIR and PIR rates are set to non-zero values. Measured value of latency, jitter and delay variation is not compared.

If the acceptance-criteria parameter is not specified and color-aware is set to enable, then the test is declared to be pass, if the measured CIR and PIR rates matches the configured CIR and PIR values and frame loss is zero OR if one of the following is true:

- If the measured throughput rate (CIR + PIR) is equal to the configured CIR rate and if no PIR rate is configured.
- If the measured throughput rate (CIR + PIR) is equal to the configured PIR rate and if either no CIR rate is configured or if CIR rate is configured.

The test is declared to ‘FAIL’ otherwise. Measured value of latency, jitter, and delay variation is not compared.

If acceptance-criteria is specified and color-aware is set to enable, the test will use the configured packet header marking values (that is, dot1p) to identify the color of the packet and classify it as green (in-profile) or yellow (out-of-profile). It measures the green packet (that is, CIR rate) and the green/in-profile packet performance parameter values and the yellow packet rate (that is, PIR rate) and the yellow/out-of-profile

Testhead Commands

packet performance parameter values individually based on the packet markings. In addition to comparing the measured performance parameter values against the normal performance parameter threshold values (if enabled), if user has enabled in/out thresholds for performance parameters in the acceptance-criteria, the tool will use these values to compare against the measured values and declare a pass/fail result. The tool uses the *cir-threshold* and *pir-threshold* to compare against the measured CIR and PIR throughput rates and declare PASS/FAIL, if the thresholds specified by the *cir-threshold* and *pir-threshold* are achieved.

NOTE: When color-aware mode is set to enable, the marking values used to identify both in-profile/green packet and out-of-profile/yellow packet must be configured. If either of the packet header marking values (For example: dot1p) are not configured by the user, then the CLI displays an error.

If acceptance-criteria is specified and color-aware is set to disable, the tests are color blind (not color-aware). The tool does not use the configured packet header marking values to identify the color of the packet and treats all packets the same. The tool uses the normal thresholds configured in the acceptance-criteria (i.e. the threshold values other than the in/out profile thresholds) to compare the measured values and declare a pass/fail result. The tool will not make any attempt to compare the in/out thresholds against measured values. The tool uses the *cir-threshold* and *pir-threshold* as follows:

- If no PIR rate is configured and if the measured throughput rate is equal to the configured cir-threshold rate, the desired rate is said to have been achieved and the tests continue to compare the measured performance parameter thresholds with the configured performance parameter thresholds (if any).
- If PIR rate is configured and no CIR rate is configured and if the measured throughput rate is equal to the configured pir-threshold rate, the desired rate is said to have been achieved and the tests continue to compare the measured performance parameter thresholds with the configured performance parameter thresholds (if any).
- If PIR rate is configured and CIR rate is configured and if the measured throughput rate is equal to the configured pir-threshold rate, the desired rate is said to have been achieved and the tests continue to compare the measured performance parameter thresholds with the configured performance parameter thresholds (if any).

The test-name and owner-name together identify a particular testhead invocation/session uniquely. The results of the testhead session are associated with the test-name and owner-name. These parameters must be used if the user needs to display the results of the testhead tool and to clear the results of a completed run. Multiple invocations of the testhead tool with the same test-name and owner-name is not allowed if the results of the old run using the same pair of test-name and owner-name are present. In other words, the results are not overwritten when the testhead is invoked again with the same values for test-name and owner-name. The results needs to be cleared explicitly using the clear command before invoking the testhead tool with the same test-name and owner-name. Results for up to 100 unique sessions each using a different test-name and owner-name is saved in memory (in other words, the results are not available for use after a reboot).

NOTE: This command is not saved in the configuration file across a reboot.

For more information, See the, “Pre-requisites for using the Testhead Tool on page 183”

Default no defaults

Parameters *test-name* — Name of the test

Values ASCII string upto 32 characters in length

owner test-owner — Specifies the owner of an testhead operation.

Values ASCII string upto 32 characters in length

testhead-profile profile-id — Specifies the testhead profile ID to use with this run/session of testhead invocation. Testhead profile must be configured beforehand using the commands under config> test-oam> test-head-profile>.

Values 1- 10

frame-payload frame-payload-id — Optional parameter used to specify the frame payload ID to use for this run. It identifies the parameters used to construct the frame generated by the testhead tool.

Values 1 – 4, if this parameter is not specified, then by default parameters configured under frame-payload-id 1 is used by this run.

acceptance-criteria acceptance-criteria-id — Optional parameter used to specify the test acceptance criteria parameters to use. for this run. It identifies the parameters used to compare the measured performance values against the configured thresholds configured in the acceptance criteria.

Values 1 – 4. If this parameter is not specified then the run is declared pass if the throughput configured in the testhead-profile is achieved without any loss.

color-aware — Optional parameter that specifies if color aware tests need to be executed. If set to enable, then color-aware test enabled. If set to disable, which is the default, then non-color-ware test is enabled.

sap sap-id — Identifies the test SAP. Must be specified by the user.

Values null - <port-id|lag-id>
 dot1q - <port-id|lag-id>:qtag1
 qinq - <port-id|lag-id>:qtag1.qtag2
 port-id - slot/mda/port
 lag-id - lag-<id>
 lag - keyword
 id - [1..200]
 qtag1 - [0..4094]
 qtag2 - [*|1..4094]

For more information, see “SAP configuration guidelines”.

fc fc-name — Optional parameter that specifies the forwarding class (FC) to use to send the frames generated by the testhead tool.

Values be, l2, af, l1, h2, ef, h1, nc

testhead

Syntax `testhead test-name owner owner-name stop`

Context oam

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

The currently running test, if any will be stopped. All performance results based on the data available upto the time the test is stopped is used determine the pass/fail criteria. Additionally, the test-status will display “Stopped” and Test completion status will be marked “Incomplete or No”.

Parameters *test-name* — Name of the test

Values ASCII string upto 32 characters in length

owner test-owner — Specifies the owner of an testhead operation.

Values ASCII string upto 32 characters in length

OAM Performance Monitoring and Binning Commands

oam-pm

Syntax oam-pm

Context config

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command enables the context to configure storage parameters (including binning structures), availability/resiliency, and the individual proactive and on-demand tests used to gather performance and statistical data.

bin-group

Syntax **bin-group** *bin-group-number* [**fd-bin-count** *fd-bin-count* **fdr-bin-count** *fdr-bin-count* **ifdv-bin-count** *ifdv-bin-count* **create**]
no bin-group *bin-group-number*

Context config>oam-pm

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command configures the parameters for a specific bin group. Bin-group 1 is a default bin group and cannot be modified. If no bin group is assigned to an OAM-PM session, bin-group 1 will be assigned by default. The default values for bin-group 1 are **fd-bin-count 3 bin 1 lower-bound 5000 bin 2 lower-bound 10000**, **fdr-bin-count 2 bin 1 lower-bound 5000**, and **ifdv-bin-count 2 bin 1 lower-bound 5000**.

The **no** form of the command removes the specified bin group.

Parameters *bin-group-number* — Specifies the numerical identifier for a bin group. A bin group can only shut down and modified when all of the PM sessions referencing the bin group have been shut down. The bin group **description** may still be modified for active bin groups.

Values 1 to 255

fd-bin-count — Specifies the number of frame delay bins that will be created.

Values 2 to 10

fdr-bin-count — Specifies the number of frame delay range bins that will be created.

Values 2 to 10

ifdv-bin-count — Specifies the number of inter-frame delay variation bins that will be created.

Values 2 to 10

OAM Performance Monitoring and Binning Commands

create — Keyword to create the specified bin group.

bin-type

Syntax **bin-type** {**fd** | **fdr** | **ifdv**}
no bin-type

Context config>oam-pm>bin-group

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command enables the specified delay metric configuration context.

The **no** form of the command restores the default value.

Default bin-type fd

Parameters **fd** — Keyword to enter the frame delay bin threshold configuration context.

fdr — Keyword to enter the frame delay range bin threshold configuration context.

ifdv — Keyword to enter the inter-frame delay variation bin threshold configuration context.

bin

Syntax **bin** *bin-number*

Context config>oam-pm>bin-group>bin

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command enables the context to configure the floor threshold for an individual bin.

Parameters *bin-number* — Specifies the bin to configure.

Values 1 to 9

lower-bound

Syntax **lower-bound** *microseconds*
no lower-bound

Context config>oam-pm>bin-group>bin-type>bin

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command configures the lower threshold for an individual bin. The operator does not have to specify a lower threshold for every bin that was previously defined by the bin-count for the specific type. By default, the lower threshold for each bin is the bin-number * 5000 microseconds. Lower thresholds in the previous adjacent bin must be lower than the threshold of the next higher bin threshold; otherwise, an error will prevent the bin from entering the active state when the **no shutdown** command is issued for the bin group. Bin 0 is the result of the difference between 0 and the configured **lower-bound** of bin 1. The highest bin in the bin-count will capture every result above the threshold. Any negative delay metric result will be treated as zero and placed in bin 0.

The **no** form of the command restores the default threshold for the bin.

Parameters *microseconds* — Specifies the lower threshold for the bin, in microseconds.

Values 1 to 4294967295

Default bin-number * 5000

description

Syntax **description** *description-string*
no description

Context config>oam-pm>bin-group
config>oam-pm>session

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

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.

Parameters *description-string* — Specifies the description character string. Allowed values are any characters 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 in double quotes.

shutdown

Syntax [**no**] shutdown

Context config>oam-pm>bin-group
config>oam-pm>session>ethernet>dmm
config>oam-pm>session>ethernet>slm

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command activates and deactivates the bin group or test.

OAM Performance Monitoring and Binning Commands

When a bin group is active, only the **description** of the bin group can be modified. The bin group can only be shut down and modified when all references in the various PM sessions or individual tests have been shut down. If an active PM session is referencing the bin group, it will generate an error indicating there are a number of active tests referencing the bin group, and it cannot be shut down.

When a test is shut down, no active measurements are made and any outstanding requests are ignored. If the test is started or stopped during a measurement interval, the suspect flag will be set to “yes” to indicate that the data for the specific data set is in questionable.

The **no** form of the command activates the bin group or test.

session

Syntax **session** *session-name* [**test-family** {**ethernet** | **ip**} [**session-type** {**proactive** | **on-demand**}]
create]
no session *session-name*

Context config>oam-pm

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command configures the individual session containers that will house the test-specific configuration parameters. Since this session context provides only a container abstract to house the individual test functions, it cannot be shut down. Only individual tests sessions within the container may be shut down. No values, parameters, or configuration within this context may be changed if any individual test is active. Changes may only be made when all tests within the context are shut down, with the exception of the **description**.

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

Parameters *session-name* — Specifies the name of the session container. 32 characters maximum.

ethernet — Specifies that the test will be based on the Ethernet layer.

ip — Specifies that the test will be based on the IP layer.

proactive — Specifies that the test will always be on, with no stop. Tests are proactive by default.

on-demand — Specifies that the test will run on demand, with an immediate start and no stop, or a stop based on offset.

create — Keyword to create the session container.

bin-group

| | |
|--------------------|---|
| Syntax | bin-group <i>bin-group-number</i>
no bin-group |
| Context | config>oam-pm>session |
| Description | Platforms Supported: 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command links the individual test to the group of bins that map the probe responses.
The no form of this command installs the default bin-group 1 as the bin group for the session. |
| Default | bin-group 1 |
| Parameters | <i>bin-group-number</i> — Specifies the number of the bin-group that will be referenced during this session.

Values 1 to 255 |

ethernet

| | |
|--------------------|--|
| Syntax | ethernet |
| Context | config>oam-pm>session |
| Description | Platforms Supported: 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command enables the context to configure Ethernet-specific source and destination information, priority, and Ethernet test tools on the launch point. |

dest-mac

| | |
|--------------------|---|
| Syntax | dest-mac <i>ieee-address</i>
no dest-mac |
| Context | config>oam-pm>session>ethernet |
| Description | Platforms Supported: 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command defines the destination MAC address of the peer MEP and sets the destination MAC address in the Layer 2 header to match. This must be a unicast address.
The no form of the command removes the session parameter. |
| Parameters | <i>ieee-address</i> — Specifies the Layer 2 unicast MAC address of the destination MEP.

Values 6-byte unicast mac-address (xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx) |

dmm

| | |
|--------------------|---|
| Syntax | dmm [test-id <i>test-id</i>] [create]
no dmm |
| Context | config>oam-pm>session>ethernet |
| Description | Platforms Supported: 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command configures the test ID to be assigned to the delay test, and creates the container to allow the individual test parameters to be configured.

The no form of the command removes the DMM test function from the PM session. |
| Parameters | <i>test-id</i> — Specifies the value to be placed in the 4-byte test ID field of the ETH-DMM PDU.

Values 0 to 2147483647

create — Keyword to create the test. |

data-tlv-size

| | |
|--------------------|---|
| Syntax | data-tlv-size <i>octets</i>
no data-tlv-size |
| Context | config>oam-pm>session>ethernet>dmm
config>oam-pm>session>ethernet>slm |
| Description | Platforms Supported: 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command is used to add an optional Data TLV to the PDU and increase the frame on the wire by the specified amount. This value is not the total size of the frame on the wire, but rather the size of the additional padding added to the PDU.

The no form of the command removes the optional TLV. |
| Default | data-tlv-size 0 |
| Parameters | <i>octets</i> — Specifies the size of the optional Data TLV, in octets.

Values 0, 3 to 2000 |

interval

| | |
|--------------------|---|
| Syntax | interval <i>milliseconds</i>
no interval |
| Context | config>oam-pm>session>ethernet>dmm |
| Description | Platforms Supported: 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command configures the message period, or probe spacing, for the transmission of DMM frames. The no form of the command restores the default value. |
| Default | interval 1000 |
| Parameters | <i>milliseconds</i> — Specifies the number of milliseconds between the transmission of DMM frames. The default value for the DMM interval is intentionally different from the default value for the SLM interval.

Values 100, 1000, 10000 |

test-duration

| | |
|--------------------|---|
| Syntax | test-duration <i>seconds</i>
no test-duration |
| Context | config>oam-pm>session>ethernet>dmm
config>oam-pm>session>ethernet>slm |
| Description | Platforms Supported: 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This optional command defines the length of time the test will run before stopping automatically. This command is only a valid option when a session has been configured with a session-type of on-demand . This is not an option when the session-type is configured as proactive . All tests start immediately following the execution of a no shutdown command.

The test duration value, remaining time, or completed state, is not synchronized with the backup CPM. This means that a failover will re-launch any active test without regard to the test-duration timer on the previously active CPM. When the test starts on the newly active CPM, the test-duration will be reset to the beginning.

The no form of the command will remove a previously configured test-duration and allow the test to execute until manually stopped. |
| Parameters | <i>seconds</i> — Specifies the interval, in seconds, during which the test will continue to execute after the start time.

Values 1 to 86400 |

priority

| | |
|--------------------|--|
| Syntax | priority <i>priority</i>
no priority |
| Context | config>oam-pm>session>ethernet |
| Description | <p>Platforms Supported: 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.</p> <p>This command configures the CoS priority across all tests configured under this session. This CoS value is exposed to the various QoS policies the frame will pass through and does not necessarily map directly to the CoS value on the wire.</p> <p>The no form of the command restores the default value.</p> |
| Default | priority 0 |
| Parameters | <p><i>priority</i> — Specifies the CoS value.</p> <p>Values 0 to 7</p> |

slm

| | |
|--------------------|---|
| Syntax | slm [<i>test-id test-id</i>] [create]
no slm |
| Context | config>oam-pm>session>ethernet |
| Description | <p>Platforms Supported: 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.</p> <p>This command configures the test ID to be assigned to the synthetic loss test, and creates the container to allow the individual test parameters to be configured.</p> <p>The no form of the command removes the SLM test function from the PM session.</p> |
| Parameters | <p><i>test-id</i> — Specifies the value to be placed in the 4-byte test ID field of the ETH-SLM PDU.</p> <p>Values 0 to 2147483647</p> <p>create — Keyword to create the test.</p> |

flr-threshold

| | |
|--------------------|--|
| Syntax | flr-threshold <i>percentage</i>
no flr-threshold |
| Context | config>oam-pm>session>ethernet>slm |
| Description | <p>Platforms Supported: 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.</p> <p>This command configures the frame loss threshold that is used to determine if the delta-t is available or unavailable. An individual delta-t with a frame loss threshold equal to or higher than the configured threshold will be marked unavailable. An individual delta-t with a frame loss threshold lower than the configured threshold will be marked as available.</p> <p>The no form of the command restores the default value.</p> |
| Default | flr-threshold 50 |
| Parameters | <p><i>percentage</i> — Specifies the percentage of the threshold.</p> <p>Values 50</p> |

timing

| | |
|--------------------|--|
| Syntax | timing frames-per-delta-t <i>frames</i> consec-delta-t <i>deltas</i> interval <i>milliseconds</i> chli-threshold
<i>threshold</i>
no timing |
| Context | config>oam-pm>session>ethernet>slm |
| Description | <p>Platforms Supported: 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.</p> <p>This command configures various availability parameters and the probe spacing (interval) for the SLM frames. The maximum size of the availability window cannot exceed 10 s (10000 ms).</p> <p>The no form of the command will install the default values for all timing parameters and use those values to compute availability and set the SLM frequency. If an SLM test is active, it will always have timing parameters, whether default or operator-configured.</p> |
| Default | timing frames-per-delta-t 10 consec-delta-t 10 interval 100 chli-threshold 5 |
| Parameters | <p><i>frames</i> — Specifies the number of frames that define the size of the delta-t. Each delta-t will be marked as available or unavailable based on the flr-threshold. The size of the delta-t measurement is the product of the number of frames and the interval</p> <p>Values 1 to 50</p> <p><i>deltas</i> — Specifies the number of consecutive delta-ts that make up the sliding window over which availability and unavailability will be determined. Transitions from one state to another will occur when the con-</p> |

OAM Performance Monitoring and Binning Commands

secutive delta-ts are in a new state.

Values 2 to 10

milliseconds — Specifies the number of milliseconds between the transmission of the SLM frames. The default value for the SLM interval is intentionally different from the default interval for DMM.

Values 100, 1000

threshold — Specifies the number of consecutive unavailable delta-ts that will cause the CHLI counter to be incremented. A CHLI counter is an indication that the sliding window is available but has crossed a threshold of consecutive unavailable delta-t intervals. A CHLI can only be incremented once during a sliding window and will only be incremented during times of availability.

Values 1 to 9

SOURCE

Syntax **source mep** *mep-id* **domain** *md-index* **association** *ma-index*
no source

Context config>oam-pm>session>ethernet

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command configures the source launch point for Y.1731 parameters that will be used by the individual tests within the session. If an MEP matching the configuration does not exist, the session will be allowed to become active; however, the frames sent and received as seen under the **show>oam-pm>statistics>session** *session-name* command will be zero.

The **no** form of the command removes this session parameter.

Parameters *mep-id* — Specifies the maintenance association end point identifier of the launch point.

Values 1 to 8191

md-index — Specifies the maintenance domain (MD) index value of the launch point.

Values 1 to 4294967295

ma-index — Specifies the maintenance association (MA) index value of the launch point.

Values 1 to 4294967295

meas-interval

Syntax **meas-interval** {5-mins | 15-mins | 1-hour | 1-day} [create]
no meas-interval

Context config>oam-pm>session

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command establishes the parameters of the individual measurement intervals utilized by the session. A maximum of three different measurement intervals may be configured under each session.

The **no** form of the command deletes the specified measurement interval.

Parameters **5-mins** — Specifies a 5 minute measurement interval duration.

15-mins — Specifies a 15 minute measurement interval duration.

1-hour — Specifies a 1 hour measurement interval duration.

1-day — Specifies a 1 day measurement interval duration.

create — Keyword to create the measurement interval.

accounting-policy

Syntax **accounting-policy** *acct-policy-id*
no accounting-policy

Context config>oam-pm>session>meas-interval

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This optional command assigns a record-type of complete-pm to the specified accounting policy (configured using the **config>log>accounting-policy** command). This runs the data collection process for completed measurement intervals in memory, file storage, and maintenance functions, moving data from memory to flash. A single accounting policy can be applied to a measurement interval.

The **no** form of the command removes the accounting policy.

Parameters *acct-policy-id* — Specifies the accounting policy to be applied to the measurement interval.

Values 1 to 99

boundary-type

| | |
|--------------------|--|
| Syntax | boundary-type { clock-aligned test-relative }
no boundary-type |
| Context | config>oam-pm>session>meas-interval |
| Description | <p>Platforms Supported: 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.</p> <p>This command establishes the alignment of the start of the measurement interval with either the time of day clock or the start of the test.</p> <p>Test-relative start times will launch the measurement interval when the individual test enters the active (no shutdown) state.</p> <p>Alignment with the time of day clock always defaults to the representative top of the hour. Clocks aligned at 15-minute measurement intervals will divide the hour into four equal sections at 00, 15, 30, and 45. Clocks aligned at 1-hour measurement intervals will start at 00. Clocks aligned at 1-day measurement intervals will start at midnight. It is typical for the first measurement interval of a clock-aligned test to have the suspect flag set to yes because it is unlikely that the no shutdown command will exactly correspond to the clock-based measurement interval start time. Clock-aligned measurement intervals can include an additional offset. See the clock-offset command option under this context.</p> <p>The no form of the command restores the default value.</p> |
| Default | boundary-type clock-aligned |
| Parameters | <p>clock-aligned — Keyword that aligns the start of the measurement interval with the time of day clock.</p> <p>test-relative — Keyword that aligns the start of the measurement interval with the start of the test.</p> |

clock-offset

| | |
|--------------------|---|
| Syntax | clock-offset <i>seconds</i>
no clock-offset |
| Context | config>oam-pm>session>meas-interval |
| Description | <p>Platforms Supported: 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.</p> <p>This command configures an offset between measurement intervals with a boundary-type of clock-aligned and the default time of day clock. The configured offset must be smaller than the size of the measurement interval. As an example, an offset of 300 seconds will shift the start times of the measurement intervals by 5 minutes from their default alignments with respect to the time of day clock.</p> <p>The no form of the command restores the default value.</p> |
| Default | clock-offset 0 |
| Parameters | <i>seconds</i> — Specifies the number of seconds to offset a clock-alignment measurement interval from its |

default.

Values 0 to 86399

intervals-stored

Syntax **intervals-stored** *intervals*
no intervals-stored

Context config>oam-pm>session>meas-interval

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command configures the number of completed measurement intervals per session to be stored in volatile system memory. The entire block of memory is allocated for the measurement interval when the test is active (**no shutdown**) to ensure that memory is available. The numbers increase from 1 to the configured value + 1. The active PM data will be stored in interval number 1, and older runs are stored, in order, to the upper most number, with the oldest run being deleted when the number of completed measurement intervals exceeds the configured value + 1. As new test measurement intervals complete for the session, the stored intervals will be renumbered to maintain the described order. Care must be taken when setting this value. There must be a balance between completed runs stored in volatile memory and the use of the write-to-flash function of the accounting policy.

The **5-mins** and **15-mins** measurement intervals share the same (1 to 96) retention pool. In the unlikely event that both intervals are required, the sum total of both cannot exceed 96. The **1-hour** and **1-day** measurement intervals utilize their own ranges. If this command is omitted when configuring the measurement interval, the default values will be used.

Parameters *intervals* — Specifies the number of stored intervals.

Values **5-mins** — 1 to 96 (default 32)
15-mins — 1 to 96 (default 32)
1-hour — 1 to 24 (default 8)
1-day — 1 (default 1)

Service Assurance Agent (SAA) Commands

saa

Syntax **saa**

Context config

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command creates the context to configure the Service Assurance Agent (SAA) tests.

test

Syntax **test** *name* [**owner** *test-owner*]
no test *name*

Context config>saa

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command identifies a test and create/modify the context to provide the test parameters for the named test. Subsequent to the creation of the test instance the test can be started in the OAM context.

A test can only be modified while it is shut down.

The **no** form of this command removes the test from the configuration. In order to remove a test it can not be active at the time.

Parameters *name* — Identify the saa test name to be created or edited.

owner *test-owner* — Specifies the owner of an SAA operation upto 32 characters in length.

Values If a *test-owner* value is not specified, tests created by the CLI have a default owner “TiMOS CLI”.

accounting-policy

Syntax **accounting-policy** *acct-policy-id*
no accounting-policy

Context config>saa>test

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command associates an accounting policy to the SAA test. The accounting policy must already be defined before it can be associated else an error message is generated.

A notification (trap) when a test is completed is issued whenever a test terminates.

The **no** form of this command removes the accounting policy association.

Default none

Parameters *acct-policy-id* — Enter the accounting *policy-id* as configured in the **config>log>accounting-policy** context.

Values 1 — 99

description

Syntax **description** *description-string*
no description

Context config>saa>test

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

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 this command removes the string from the configuration.

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.

continuous

Syntax [**no**] **continuous**

Context config>saa>test

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command specifies whether the SAA test is continuous. Once the test is configured as continuous, it cannot be started or stopped by using the **saa** command.

The **no** form of the command disables the continuous running of the test. Use the **shutdown** command to disable the test.

jitter-event

Syntax `jitter-event rising-threshold threshold [falling-threshold threshold] [direction]`
`no jitter-event`

Context config>saa>test

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

Specifies that at the termination of an SAA test probe, the calculated jitter value is evaluated against the configured rising and falling jitter thresholds. SAA threshold events are generated as required.

Once the threshold (rising/falling) is crossed, it is disabled from generating additional events until the opposite threshold is crossed. If a falling-threshold is not supplied, the rising threshold will be re-enabled when it falls below the threshold after the initial crossing that generate the event.

The configuration of jitter event thresholds is optional.

Parameters **rising-threshold *threshold*** — Specifies a rising threshold jitter value. When the test run is completed, the calculated jitter value is compared to the configured jitter rising threshold. If the test run jitter value is greater than the configured rising threshold value then an SAA threshold event is generated. The SAA threshold event is tmnxOamSaaThreshold, logger application OAM, event #2101.

Default 0

Values 0 — 2147483 milliseconds

falling-threshold *threshold* — Specifies a falling threshold jitter value. When the test run is completed, the calculated jitter value is compared to the configured jitter falling threshold. If the test run jitter value is greater than the configured falling threshold value then an SAA threshold event is generated. The SAA threshold event is tmnxOamSaaThreshold, logger application OAM, event #2101.

Default 0

Values 0 — 2147483 milliseconds

direction — Specifies the direction for OAM ping responses received for an OAM ping test run.

Values **inbound** — Monitor the value of jitter calculated for the inbound, one-way, OAM ping responses received for an OAM ping test run.

outbound — Monitor the value of jitter calculated for the outbound, one-way, OAM ping requests sent for an OAM ping test run.

roundtrip — Monitor the value of jitter calculated for the round trip, two-way, OAM ping requests and replies for an OAM ping test run.

Default roundtrip

latency-event

Syntax `latency-event rising-threshold threshold [falling-threshold threshold] [direction]`
`no latency-event`

Context `config>saa>test`

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

Specifies that at the termination of an SAA test probe, the calculated latency event value is evaluated against the configured rising and falling latency event thresholds. SAA threshold events are generated as required.

Once the threshold (rising/falling) is crossed, it is disabled from generating additional events until the opposite threshold is crossed. If a falling-threshold is not supplied, the rising threshold will be re-enabled when it falls below the threshold after the initial crossing that generate the event.

The configuration of latency event thresholds is optional.

Parameters **rising-threshold *threshold*** — Specifies a rising threshold latency value. When the test run is completed, the calculated latency value is compared to the configured latency rising threshold. If the test run latency value is greater than the configured rising threshold value then an SAA threshold event is generated. The SAA threshold event is `tmnxOamSaaThreshold`, logger application OAM, event #2101.

Default 0

Values 0 — 2147483 milliseconds

falling-threshold *threshold* — Specifies a falling threshold latency value. When the test run is completed, the calculated latency value is compared to the configured latency falling threshold. If the test run latency value is greater than the configured falling threshold value then an SAA threshold event is generated. The SAA threshold event is `tmnxOamSaaThreshold`, logger application OAM, event #2101.

Default 0

Values 0 — 2147483 milliseconds

direction — Specifies the direction for OAM ping responses received for an OAM ping test run.

Values **inbound** — Monitor the value of jitter calculated for the inbound, one-way, OAM ping responses received for an OAM ping test run.

outbound — Monitor the value of jitter calculated for the outbound, one-way, OAM ping requests sent for an OAM ping test run.

roundtrip — Monitor the value of jitter calculated for the round trip, two-way, OAM ping requests and replies for an OAM ping test run.

Default roundtrip

loss-event

Syntax **loss-event rising-threshold** *threshold* [**falling-threshold** *threshold*] [**direction**]
no loss-event

Context config>saa>test

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

Specifies that at the termination of an SAA testrun, the calculated loss event value is evaluated against the configured rising and falling loss event thresholds. SAA threshold events are generated as required.

The configuration of loss event thresholds is optional.

Parameters **rising-threshold** *threshold* — Specifies a rising threshold loss event value. When the test run is completed, the calculated loss event value is compared to the configured loss event rising threshold. If the test run loss event value is greater than the configured rising threshold value then an SAA threshold event is generated. The SAA threshold event is tmnxOamSaaThreshold, logger application OAM, event #2101.

Default 0

Values 0 — 2147483647 packets

falling-threshold *threshold* — Specifies a falling threshold loss event value. When the test run is completed, the calculated loss event value is compared to the configured loss event falling threshold. If the test run loss event value is greater than the configured falling threshold value then an SAA threshold event is generated. The SAA threshold event is tmnxOamSaaThreshold, logger application OAM, event #2101.

Default 0

Values 0 — 2147483647 packets

direction — Specifies the direction for OAM ping responses received for an OAM ping test run.

Values **inbound** — Monitor the value of jitter calculated for the inbound, one-way, OAM ping responses received for an OAM ping test run.

outbound — Monitor the value of jitter calculated for the outbound, one-way, OAM ping requests sent for an OAM ping test run.

roundtrip — Monitor the value of jitter calculated for the round trip, two-way, OAM ping requests and replies for an OAM ping test run.

Default roundtrip

probe-history

Syntax **probe-history** [**auto** | **drop** | **keep**]

Context config>saa>test

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

Defines history probe behavior. Defaults are associated with various configured parameters within the SAA test. Auto (keep) is used for test with probe counts of 100 or less, and intervals of 1 second and above. Auto (drop) will only maintain summary information for tests marked as continuous with file functions, probe counts in excess of 100 and intervals of less than 1 second. SAA tests that are not continuous with a write to file will default to Auto (keep). The operator is free to change the default behaviors for each type. Each test that maintains per probe history will consume more system memory. When per probe entries are required the probe history is available at the completion of the test.

Default auto

Parameters *auto* — An auto selector that determines the storage of the history information.

drop — Store summarized min/max/ave data not per probe information for test runs. This may be configured for all tests in an effort to conserve memory.

keep — Store per probe information for tests. This consumes significantly more memory than summary information and should only be used if necessary.

trap-gen

Syntax trap-gen

Context config>saa>test

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command enables the context to configure trap generation for the SAA test.

probe-fail-enable

Syntax [no] probe-fail-enable

Context config>saa>test>trap-gen

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command enables the generation of an SNMP trap when probe-fail-threshold consecutive probes fail during the execution of the SAA ping test. This command is not applicable to SAA trace route tests.

The **no** form of the command disables the generation of an SNMP trap.

probe-fail-threshold

Syntax [no] **probe-fail-threshold** *threshold*

Context config>saa>test>trap-gen

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command has no effect when probe-fail-enable is disabled. This command is not applicable to SAA trace route tests.

The **probe-fail-enable** command enables the generation of an SNMP trap when the probe-fail-threshold consecutive probes fail during the execution of the SAA ping test. This command is not applicable to SAA trace route tests.

The **no** form of the command returns the threshold value to the default.

Default 1

test-completion-enable

Syntax [no] **test-completion-enable**

Context config>saa>test>trap-gen

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command enables the generation of a trap when an SAA test completes.

The **no** form of the command disables the trap generation.

test-fail-enable

Syntax [no] **test-fail-enable**

Context config>saa>test>trap-gen

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command enables the generation of a trap when a test fails. In the case of a ping test, the test is considered failed (for the purpose of trap generation) if the number of failed probes is at least the value of the **test-fail-threshold** parameter.

The **no** form of the command disables the trap generation.

test-fail-threshold

Syntax [no] test-fail-threshold *threshold*

Context config>saa>test>trap-gen

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command configures the threshold for trap generation on test failure.

This command has no effect when test-fail-enable is disabled. This command is not applicable to SAA trace route tests.

The **no** form of the command returns the threshold value to the default.

Default 1

type

Syntax type
no type

Context config>saa>test

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command creates the context to provide the test type for the named test. Only a single test type can be configured.

A test can only be modified while the test is in shut down mode.

Once a test type has been configured the command can be modified by re-entering the command, the test type must be the same as the previously entered test type.

To change the test type, the old command must be removed using the **config>saa>test>no type** command.

cpe-ping

Note: This command is not supported on 7210 SAS-M and 7210 SAS-T devices configured in Access uplink mode. Only on the devices configured in Network mode.

Service Assurance Agent (SAA) Commands

| | |
|-------------|--|
| Syntax | cpe-ping service <i>service-id</i> destination <i>ip-address</i> source <i>ip-address</i> [ttl <i>vc-label-ttl</i>] [return-control] [source-mac <i>ieee-address</i>] [fc <i>fc-name</i>] [interval <i>interval</i>] [send-count <i>send-count</i>] [send-control] |
| Context | oam
config>saa>test>type |
| Description | Platforms Supported: 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This ping utility determines the IP connectivity to a CPE within a specified VPLS service. |
| Parameters | service <i>service-id</i> — The service ID of the service to diagnose or manage.

Values
<i>service-id:</i> 1 — 2147483647
<i>svc-name:</i> 64 characters maximum

destination <i>ip-address</i> — Specifies the IP address to be used as the destination for performing an OAM ping operations.

source <i>ip-address</i> — Specify an unused IP address in the same network that is associated with the VPLS.

ttl <i>vc-label-ttl</i> — The TTL value in the VC label for the OAM MAC request, expressed as a decimal integer.

Default 255
Values 1 — 255

return-control — Specifies the MAC OAM reply to a data plane MAC OAM request be sent using the control plane instead of the data plane.

Default MAC OAM reply sent using the data plane.

source-mac <i>ieee-address</i> — Specify the source MAC address that will be sent to the CPE. If not specified or set to 0, the MAC address configured for the CPMCFM is used.

fc-name — The forwarding class of the MPLS echo request encapsulation.

Default be
Values be, l2, af, ll, h2, ef, h1, nc

interval <i>interval</i> — The interval parameter in seconds, expressed as a decimal integer. This parameter is used to override the default request message send interval and defines the minimum amount of time that must expire before the next message request is sent.

If the interval is set to 1 second where the timeout value is set to 10 seconds, then the maximum time between message requests is 10 seconds and the minimum is 1 second. This depends upon the receipt of a message reply corresponding to the outstanding message request.

Default 1
Values 1 — 10

send-count <i>send-count</i> — The number of messages to send, expressed as a decimal integer. The count parameter is used to override the default number of message requests sent. Each message request must either timeout or receive a reply before the next message request is sent. The message interval value must be expired before the next message request is sent. |

Default 1

Values 1 — 255

send-control — Specifies the MAC OAM request be sent using the control plane instead of the data plane.

Default MAC OAM request sent using the data plane.

dns

Syntax **dns target-addr** *dns-name* **name-server** *ip-address* [**source** *ip-address*] [**send-count** *send-count*] [**timeout** *timeout*] [**interval** *interval*] [**record-type** {**ipv4-a-record** | **ipv6-aaaa-record**}]

Context <GLOBAL>
config>saa>test>type

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command configures a DNS name resolution test.

Parameters **target-addr** — The IP host address to be used as the destination for performing an OAM ping operation. *dns-name* — The DNS name to be resolved to an IP address.

name-server *ip-address* — Specifies the server connected to a network that resolves network names into network addresses.

Values **ipv4-address** - a.b.c.d
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

source *ip-address* — Specifies the IP address to be used as the source for performing an OAM ping operation.

Values **ipv4-address** - a.b.c.d
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

send-count *send-count* — The number of messages to send, expressed as a decimal integer. The **send-count** parameter is used to override the default number of message requests sent. Each message request must either timeout or receive a reply before the next message request is sent. The message **interval** value must be expired before the next message request is sent.

Default 1

Values 1 — 255

timeout *timeout* — The **timeout** parameter in seconds, expressed as a decimal integer. This value is used to override the default **timeout** value and is the amount of time that the router will wait for a message reply

Service Assurance Agent (SAA) Commands

after sending the message request. Upon the expiration of message timeout, the requesting router assumes that the message response will not be received. Any response received after the request times out will be silently discarded.

Default 5

Values 1 — 120

interval *interval* — The **interval** parameter in seconds, expressed as a decimal integer. This parameter is used to override the default request message send interval and defines the minimum amount of time that must expire before the next message request is sent.

If the **interval** is set to 1 second, and the **timeout** value is set to 10 seconds, then the maximum time between message requests is 10 seconds and the minimum is 1 second. This depends upon the receipt of a message reply corresponding to the outstanding message request.

Default 1

Values 1 — 10

record-type — Specifies a record type.

Values **ipv4-a-record** - A record specific mapping a host name to an IPv4 address.

ipv6-aaaa-record - A record specific to the Internet class that stores a single IPv6 address.

eth-cfm-linktrace

Syntax **eth-cfm-linktrace** *mac-address* **mep** *mep-id* **domain** *md-index* **association** *ma-index* [**ttl** *ttl-value*] [**fc** {*fc-name*}] [**send-count** *send-count*] [**timeout** *timeout*] [**interval** *interval*] [**record-type** {*ipv4-a-record*|*ipv6-aaaa-record*}]

Context config>saa>test>type

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command configures a CFM linktrace test in SAA.

Parameters- *mac-address* — Specifies a unicast destination MAC address.

mep *mep-id* — Specifies the target MAC address.

Values 1 — 8191

domain *md-index* — Specifies the MD index.

Values 1 — 4294967295

association *ma-index* — Specifies the MA index.

Values 1 — 4294967295

ttl *ttl-value* — Specifies the maximum number of hops traversed in the linktrace.

Default 64

Values 1—255

fc *fc-name* — The **fc** parameter is used to indicate the forwarding class of the CFM Linktrace request messages.

The actual forwarding class encoding is controlled by the network egress mappings.

Default nc

Values be, l2, af, l1, h2, ef, h1, nc

send-count *send-count* — The number of messages to send, expressed as a decimal integer. The **count** parameter

is used to override the default number of message requests sent. Each message request must either timeout or receive a reply before the next message request is sent. The message **interval** value must be expired before the next message request is sent.

Default 1

Values 1 — 10

timeout *timeout* — The **timeout** parameter in seconds, expressed as a decimal integer. This value is used to override the default **timeout** value and is the amount of time that the router will wait for a message reply after sending the message request. Upon the expiration of message timeout, the requesting router assumes that the message response will not be received. Any response received after the request times out will be silently discarded. The **timeout** value must be less than the **interval**.

Default 5

Values 1 — 10

interval *interval* — The **interval** parameter in seconds, expressed as a decimal integer. This parameter is used to override the default request message send interval and defines the minimum amount of time that must expire before the next message request is sent.

If the **interval** is set to “1” second, and the **timeout** value is set to “10” seconds, then the maximum time between message requests is “10” seconds and the minimum is “1” second. This depends upon the receipt of a message reply corresponding to the outstanding message request. The **timeout** value must be less than the **interval**.

Default 5

Values 1 — 10

record-type — Specifies a record type.

Values
 ipv4-a-record — A record specific mapping a host name to an IPv4 address.
 ipv6-aaaa-record — A record specific to the Internet class that stores a single IPv6 address.

eth-cfm-loopback

Syntax `eth-cfm-loopback mac-address mep mep-id domain md-index association ma-index [size datasize] [fc {fc-name}] [count send-count][timeout timeout] [interval interval]`

Context config>saa>test>type

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command configures an Ethernet CFM loopback test in SAA.

mac-address — Specifies a unicast destination MAC address.

mep *mep-id* — Specifies the target MAC address.

Values 1 — 8191

domain *md-index* — Specifies the MD index.

Values 1 — 4294967295

association *ma-index* — Specifies the MA index.

Values 1 — 4294967295

size *data-size* — The packet size in bytes, expressed as a decimal integer.

Default 0

Values 0 — 1500

fc *fc-name* — The **fc** parameter is used to indicate the forwarding class of the CFM Loopback request messages.

The actual forwarding class encoding is controlled by the network egress mappings.

Default nc

Values be, l2, af, l1, h2, ef, h1, nc

count *send-count* — The number of messages to send, expressed as a decimal integer. The **count** parameter is used to override the default number of message requests sent. Each message request must either timeout or receive a reply before the next message request is sent. The message **interval** value must be expired before the next message request is sent.

Default 1

Values 1 — 100

timeout *timeout* — The **timeout** parameter in seconds, expressed as a decimal integer. This value is used to override the default **timeout** value and is the amount of time that the router will wait for a message reply after sending the message request. Upon the expiration of message timeout, the requesting router assumes that the message response will not be received. Any response received after the request times out will be silently discarded. The **timeout** value must be less than the **interval**.

Default 5

Values 1 — 10

interval *interval* — The **interval** parameter in seconds, expressed as a decimal integer. This parameter is used to override the default request message send interval and defines the minimum amount of time that must expire before the next message request is sent.

If the **interval** is set to “1” second, and the **timeout** value is set to “10” seconds, then the maximum time between message requests is “10” seconds and the minimum is “1” second. This depends upon the receipt of a message reply corresponding to the outstanding message request. The **timeout** value must be less than the **interval**.

Default 5

Values 1 — 10

eth-cfm-two-way-delay

Syntax **eth-cfm-two-way-delay** *mac-address* **mep** *mep-id* **domain** *md-index* **association** *ma-index* [**fc** *{fc-name}*] [**send-count** *send-count*] [**timeout** *timeout*] [**interval** *interval*]

Context config>saa>test>type

Description- **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command configures an Ethernet CFM two-way delay test in SAA.

mac-address — Specifies a unicast destination MAC address.

mep *mep-id* — Specifies the target MAC address.

Values 1 — 8191

domain *md-index* — Specifies the MD index.

Values 1 — 4294967295

association *ma-index* — Specifies the MA index.

Values 1 — 4294967295

tll ttl-value — Specifies the maximum number of hops traversed in the linktrace.

Default 64

Values 1 — 255

fc *fc-name* — The **fc** parameter is used to indicate the forwarding class of the CFM two-delay request messages.

The actual forwarding class encoding is controlled by the network egress mappings.

Default nc

Values be, l2, af, l1, h2, ef, h1, nc

send-count *send-count* — The number of messages to send, expressed as a decimal integer. The **count** parameter

is used to override the default number of message requests sent. Each message request must either timeout

Service Assurance Agent (SAA) Commands

or receive a reply before the next message request is sent. The message **interval** value must be expired before the next message request is sent.

Default 1

Values 1 — 100

timeout *timeout* — The **timeout** parameter in seconds, expressed as a decimal integer. This value is used to override the default **timeout** value and is the amount of time that the router will wait for a message reply after sending the message request. Upon the expiration of message timeout, the requesting router assumes that the message response will not be received. Any response received after the request times out will be silently discarded. The **timeout** value must be less than the **interval**.

Default 5

Values 1 — 10

interval *interval* — The **interval** parameter in seconds, expressed as a decimal integer. This parameter is used to override the default request message send interval and defines the minimum amount of time that must expire before the next message request is sent.

If the **interval** is set to “1” second, and the **timeout** value is set to “10” seconds, then the maximum time between message requests is “10” seconds and the minimum is “1” second. This depends upon the receipt of a message reply corresponding to the outstanding message request. The **timeout** value must be less than the **interval**.

Default 5

Values [0.1, 0.2, .. 0.9] | [1, 2, .. 10]

eth-cfm-two-way-slm

Syntax **eth-cfm-two-way-delay** *mac-address* **mep** *mep-id* **domain** *md-index* **association** *ma-index* [**fc** *fc-name*] [**send-count** *send-count*] [**size** *data-size*] [**timeout** *timeout*] [**interval** *interval*]

Context config>saa>test>type

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command configures an Ethernet CFM two-way SLM test in SAA.

mac-address — Specifies a unicast destination MAC address.

mep *mep-id* — Specifies the target MAC address.

Values 1 — 8191

domain *md-index* — Specifies the MD index.

Values 1 — 4294967295

association *ma-index* — Specifies the MA index.

Values 1 — 4294967295

fc *fc-name* — The *fc* parameter is used to indicate the forwarding class of the CFM SLM request messages. The actual forwarding class encoding is controlled by the network egress mappings.

Default nc

Values be, l2, af, l1, h2, ef, h1, nc

send-count *send-count* — The number of messages to send, expressed as a decimal integer. The count parameter is used to override the default number of message requests sent. Each message request must either timeout or receive a reply before the next message request is sent. The message interval value must be expired before the next message request is sent.

Default 1

Values 1 — 100

size *data-size* — This is the size of the data portion of the data TLV. If 0 is specified no data TLV is added to the packet.

Default 0

Values 0 — 1500

timeout *timeout* — The timeout parameter in seconds, expressed as a decimal integer. This value is used to override the default timeout value and is the amount of time that the router waits for a message reply after sending the message request. Upon the expiration of message timeout, the requesting router assumes that the message response is not received. Any response received after the request times out is silently discarded. The timeout value must be less than the interval.

Default 5

Values 1 — 10

interval *interval* — The interval parameter in seconds, expressed as a decimal integer. This parameter is used to override the default request message send interval and defines the minimum amount of time that must expire before the next message request is sent. If the interval is set to 1 second, and the timeout value is set to 10 seconds, then the maximum time between message requests is 10 seconds and the minimum is 1 second. This depends upon the receipt of a message reply corresponding to the outstanding message request. The timeout value must be less than the interval.

Default 5

Values [0.1, 0.2, .. 0.9] | [1, 2, .. 10]

icmp-ping

Syntax **icmp-ping** [*ip-address* | *dns-name*] [**rapid** | **detail**] [**ttl** *time-to-live*] [**tos** *type-of-service*] [**size** *bytes*] [**pattern** *pattern*] [**source** *ip-address* | *dns-name*] [**interval** *seconds*] [{**next-hop** *ip-address*}]

Service Assurance Agent (SAA) Commands

| {**interface** *interface-name*} | **bypass-routing**] [**count** *requests*] [**do-not-fragment**] [**router** *router-instance* | **service-name** *service-name*] [**timeout** *timeout*]

Context config>saa>test>type

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command configures an ICMP ping test.

Parameters *ip-address* — The far-end IP address to which to send the **svc-ping** request message in dotted decimal notation.

Values

| | |
|---------------|---|
| ipv4-address: | a.b.c.d |
| 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 |

dns-name — The DNS name of the far-end device to which to send the **svc-ping** request message, expressed as a character string up to 63 characters maximum.

rapid — Packets will be generated as fast as possible instead of the default 1 per second.

detail — Displays detailed information.

tll *time-to-live* — The TTL value for the IP packet, expressed as a decimal integer.

Values 1 — 128

tos *type-of-service* — Specifies the service type.

Values 0 — 255

size *bytes* — The request packet size in bytes, expressed as a decimal integer.

Values 0 — 16384

pattern *pattern* — The data portion in a ping packet will be filled with the pattern value specified. If not specified, position info will be filled instead.

Values 0 — 65535

source *ip-address|dns-name* — Specifies the IP address to be used.

Values ipv4-address: a.b.c.ddns-name: 128 characters max

interval *seconds* — This parameter is used to override the default request message send interval and defines the minimum amount of time that must expire before the next message request is sent.

If the **interval** is set to 1 second, and the **timeout** value is set to 10 seconds, then the maximum time between message requests is 10 seconds and the minimum is 1 second. This depends upon the receipt of a message reply corresponding to the outstanding message request.

Default 1

Values 1 — 10

next-hop *ip-address* — Only displays static routes with the specified next hop IP address.

Values ipv4-address: a.b.c.d (host bits must be 0)

interface *interface-name* — The name used to refer to the interface. The name must already exist in the **config>router>interface** context.

bypass-routing — Specifies whether to send the ping request to a host on a directly attached network bypassing the routing table.

count *requests* — Specifies the number of times to perform an OAM ping probe operation. Each OAM echo message request must either timeout or receive a reply before the next message request is sent.

Values 1 — 100000

Default 5

do-not-fragment — Sets the DF (Do Not Fragment) bit in the ICMP ping packet.

router *router-instance* — Specifies the router name or service ID.

Values *router-name:* Base , management
 service-id: 1 — 2147483647

Default Base

service-name *service-name* — Specifies the service name as an integer.

Values *service-id:* 1 — 2147483647

timeout *timeout* — Overrides the default **timeout** value and is the amount of time that the router will wait for a message reply after sending the message request. Upon the expiration of message timeout, the requesting router assumes that the message response will not be received. Any response received after the request times out will be silently discarded.

Default 5

Values 1 — 10

icmp-trace

Syntax **icmp-trace** [*ip-address* | *dns-name*] [**ttl** *time-to-live*] [**wait** *milli-seconds*] [**tos** *type-of-service*] [**source** *ip-address*] [**tos** *type-of-service*] [**router** *router-instance* | **service-name** *service-name*]

Context config>saa>test>type

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command configures an ICMP traceroute test.

Parameters *ip-address* — The far-end IP address to which to send the **svc-ping** request message in dotted decimal notation.

Values ipv4-address: a.b.c.d
 ipv6-address: x:x:x:x:x:x:x (eight 16-bit pieces)

Service Assurance Agent (SAA) Commands

x:x:x:x:x:d.d.d.d
x: [0 .. FFFF]H
d: [0 .. 255]D

dns-name — The DNS name of the far-end device to which to send the **svc-ping** request message, expressed as a character string to 63 characters maximum.

ttl *time-to-live* — The TTL value for the MPLS label, expressed as a decimal integer.

Values 1 — 255

wait *milliseconds* — The time in milliseconds to wait for a response to a probe, expressed as a decimal integer.

Default 5000

Values 10 — 60000

tos *type-of-service* — Specifies the service type.

Values 0 — 255

source *ip-address* — Specifies the IP address to be used.

Values ipv4-address: a.b.c.d
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

router *router-instance* — Specifies the router name or service ID.

Values *router-name:* Base, management
service-id: 1 — 2147483647

Default Base

Values

mac-ping

Note: This command is not supported on 7210 SAS-M and 7210 SAS-T devices configured in Access uplink mode.

Syntax **mac-ping service** *service-id destination dst-ieee-address* [**source** *src-ieee-address*] [**fc** *fc-name*] [**size** *octets*] [**ttl** *vc-label-ttl*] [**send-count** *send-count*] [**send-control**] [**return-control**] [**interval** *interval*] [**timeout** *timeout*]

Context oam
config>saa>test>type

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

The `mac-ping` utility is used to determine the existence of an egress SAP binding of a given MAC within a VPLS service.

A **mac-ping** packet can be sent via the control plane or the data plane. The **send-control** option specifies the request be sent using the control plane. If **send-control** is not specified, the request is sent using the data plane.

A **mac-ping** is forwarded along the flooding domain if no MAC address bindings exist. If MAC address bindings exist, then the packet is forwarded along those paths, provided they are active. A response is generated only when there is an egress SAP binding for that MAC address or if the MAC address is a “local” OAM MAC address associated with the device’s control plan.

A **mac-ping** reply can be sent using the data plane or the control plane. The **return-control** option specifies the reply be sent using the control plane. If **return-control** is not specified, the request is sent using the data plane.

A **mac-ping** with data plane reply can only be initiated on nodes that can have an egress MAC address binding. A node without a FIB and without any SAPs cannot have an egress MAC address binding, so it is not a node where replies in the data plane will be trapped and sent up to the control plane.

A control plane request is responded to via a control plane reply only.

By default, MAC OAM requests are sent with the system or chassis MAC address as the source MAC. The **source** option allows overriding of the default source MAC for the request with a specific MAC address.

When a **source** *ieee-address* value is specified and the source MAC address is locally registered within a split horizon group (SHG), then this SHG membership will be used as if the packet originated from this SHG. In all other cases, SHG 0 (zero) will be used. Note that if the **mac-trace** is originated from a non-zero SHG, such packets will not go out to the same SHG.

If EMG is enabled, `mac-ping` will return only the first SAP in each chain.

Parameters

service *service-id* — The service ID of the service to diagnose or manage.

Values *service-id*: 1 — 2147483647

destination *ieee-address* — The destination MAC address for the OAM MAC request.

size *octets* — The MAC OAM request packet size in octets, expressed as a decimal integer. The request payload is padded to the specified size with a 6 byte PAD header and a byte payload of 0xAA as necessary. If the octet size specified is less than the minimum packet, the minimum sized packet necessary to send the request is used.

Default No OAM packet padding.

Values 1 — 65535

ttl *vc-label-ttl* — The TTL value in the VC label for the OAM MAC request, expressed as a decimal integer.

Default 255

Values 1 — 255

send-control — Specifies the MAC OAM request be sent using the control plane instead of the data plane.

Default MAC OAM request sent using the data plane.

return-control — Specifies the MAC OAM reply to a data plane MAC OAM request be sent using the con-

Service Assurance Agent (SAA) Commands

trol plane instead of the data plane.

Default MAC OAM reply sent using the data plane.

source *src-ieee-address* — The source MAC address from which the OAM MAC request originates. By default, the system MAC address for the chassis is used.

Default The system MAC address.

Values Any unicast MAC value.

fc *fc-name* — The **fc** parameter is used to test the forwarding class of the MPLS echo request packets. The actual forwarding class encoding is controlled by the network egress LSP-EXP mappings.

Values be, l2, af, ll, h2, ef, h1, nc

interval *interval* — The **interval** parameter in seconds, expressed as a decimal integer. This parameter is used to override the default request message send interval and defines the minimum amount of time that must expire before the next message request is sent.

If the **interval** is set to 1 second where the **timeout** value is set to 10 seconds, then the maximum time between message requests is 10 seconds and the minimum is 1 second. This depends upon the receipt of a message reply corresponding to the outstanding message request.

Default 1

Values 1 — 10

send-count *send-count* — The number of messages to send, expressed as a decimal integer. The **count** parameter is used to override the default number of message requests sent. Each message request must either timeout or receive a reply before the next message request is sent. The message **interval** value must be expired before the next message request is sent.

Default 1

Values 1 — 100

timeout *timeout* — The **timeout** parameter in seconds, expressed as a decimal integer. This value is used to override the default **timeout** value and is the amount of time that the router will wait for a message reply after sending the message request. Upon the expiration of message timeout, the requesting router assumes that the message response will not be received. Any response received after the request times out will be silently discarded.

Default 5

Values 1 — 10

Sample Output

```
oam mac-ping service 1 destination 00:bb:bb:bb:bb:bb
Seq Node-id Path RTT
-----
[Send request Seq. 1, Size 126]
1 2.2.2.2:sap1/1/1:1 In-Band 960ms
-----
```

sdp-ping

Note: This command is not supported on 7210 SAS-M and 7210 SAS-T devices configured in Access uplink mode.

Syntax **sdp-ping** *orig-sdp-id* [**resp-sdp** *resp-sdp-id*] [**fc** *fc-name* [**profile** {**in** | **out**}]] [**timeout** *seconds*] [**interval** *seconds*] [**size** *octets*] [**send-count** *send-count*] [**interval** *<interval>*]

Context oam
config>saa>test>type

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command tests SDPs for uni-directional or round trip connectivity and performs SDP MTU Path tests.

The **sdp-ping** command accepts an originating SDP-ID and an optional responding SDP-ID. The size, number of requests sent, message time-out and message send interval can be specified. All **sdp-ping** requests and replies are sent with PLP OAM-Label encapsulation, as a *service-id* is not specified.

For round trip connectivity testing, the **resp-sdp** keyword must be specified. If **resp-sdp** is not specified, a uni-directional SDP test is performed.

To terminate an **sdp-ping** in progress, use the CLI break sequence <Ctrl-C>.

An **sdp-ping** response message indicates the result of the **sdp-ping** message request. When multiple response messages apply to a single SDP echo request/reply sequence, the response message with the highest precedence will be displayed. The following table displays the response messages sorted by precedence.

| Result of Request | Displayed Response Message | Precedence |
|--|--------------------------------|------------|
| Request timeout without reply | Request Timeout | 1 |
| Request not sent due to non-existent <i>orig-sdp-id</i> | Orig-SDP Non-Existent | 2 |
| Request not sent due to administratively down <i>orig-sdp-id</i> | Orig-SDP Admin-Down | 3 |
| Request not sent due to operationally down <i>orig-sdp-id</i> | Orig-SDP Oper-Down | 4 |
| Request terminated by user before reply or timeout | Request Terminated | 5 |
| Reply received, invalid <i>origination-id</i> | Far End: Originator-ID Invalid | 6 |
| Reply received, invalid <i>responder-id</i> | Far End: Responder-ID Error | 7 |
| Reply received, non-existent <i>resp-sdp-id</i> | Far End: Resp-SDP Non-Existent | 8 |
| Reply received, invalid <i>resp-sdp-id</i> | Far End: Resp-SDP Invalid | 9 |
| Reply received, <i>resp-sdp-id</i> down (admin or oper) | Far-end: Resp-SDP Down | 10 |
| Reply received, No Error | Success | 11 |

Service Assurance Agent (SAA) Commands

Parameters *orig-sdp-id* — The SDP-ID to be used by **sdp-ping**, expressed as a decimal integer. The far-end address of the specified SDP-ID is the expected *responder-id* within each reply received. The specified SDP-ID defines the encapsulation of the SDP tunnel encapsulation used to reach the far end. This can be IP or MPLS. If *orig-sdp-id* is invalid or administratively down or unavailable for some reason, the SDP Echo Request message is not sent and an appropriate error message is displayed (once the **interval** timer expires, sdp-ping will attempt to send the next request if required).

Values 1 — 17407

resp-sdp *resp-sdp-id* — Optional parameter is used to specify the return SDP-ID to be used by the far-end 7210 SAS M for the message reply for round trip SDP connectivity testing. If *resp-sdp-id* does not exist on the far-end 7210 SAS M, terminates on another 7210 SAS M different than the originating 7210 SAS M, or another issue prevents the far-end router from using *resp-sdp-id*, the SDP echo reply will be sent using generic IP/OAM encapsulation. The received forwarding class (as mapped on the ingress network interface for the far end) defines the forwarding class encapsulation for the reply message.

Default null. Use the non-SDP return path for message reply.

Values 1 — 17407

fc *fc-name* — The **fc** parameter is used to indicate the forwarding class of the SDP encapsulation. The actual forwarding class encoding is controlled by the network egress DSCP or LSP-EXP mappings.

The DSCP or LSP-EXP mappings on the receive network interface controls the mapping back to the internal forwarding class used by the far-end 7210 SAS M that receives the message request. The egress mappings of the egress network interface on the far-end 7210 SAS M controls the forwarding class markings on the return reply message.

The DSCP or LSP-EXP mappings on the receive network interface controls the mapping of the message reply back at the originating 7210 SAS M. This is displayed in the response message output upon receipt of the message reply.

Default be

Values be, l2, af, ll, h2, ef, h1, nc

timeout *seconds* — The **timeout** parameter in seconds, expressed as a decimal integer. This value is used to override the default **timeout** value and is the amount of time that the router will wait for a message reply after sending the message request. Upon the expiration of message timeout, the requesting router assumes that the message response will not be received. A 'request timeout' message is displayed by the CLI for each message request sent that expires. Any response received after the request times out will be silently discarded.

Default 5

Values 1 — 10

interval *seconds* — The **interval** parameter in seconds, expressed as a decimal integer. This parameter is used to override the default request message send interval and defines the minimum amount of time that must expire before the next message request is sent.

If the **interval** is set to 1 second, and the **timeout** value is set to 10 seconds, then the maximum time between message requests is 10 seconds and the minimum is 1 second. This depends upon the receipt of a message reply corresponding to the outstanding message request.

Default 1

Values 1 — 10

size *octets* — The **size** parameter in octets, expressed as a decimal integer. This parameter is used to override the default message size for the **sdp-ping** request. Changing the message size is a method of checking the ability of an SDP to support a **path-mtu**. The size of the message does not include the SDP encapsulation, VC-Label (if applied) or any DLC headers or trailers.

When the OAM message request is encapsulated in an IP/ SDP, the IP ‘DF’ (Do Not Fragment) bit is set. If any segment of the path between the sender and receiver cannot handle the message size, the message is discarded. MPLS LSPs are not expected to fragment the message either, as the message contained in the LSP is not an IP packet.

Default 72

Values 72 — 9198

send-count *send-count* — The number of messages to send, expressed as a decimal integer. The **count** parameter is used to override the default number of message requests sent. Each message request must either timeout or receive a reply before the next message request is sent. The message **interval** value must be expired before the next message request is sent.

Default 1

Values 1 — 100

Special Cases Single Response Connectivity Tests — A single response sdp-ping test provides detailed test results.

Upon request timeout, message response, request termination, or request error the following local and remote information will be displayed. Local and remote information will be dependent upon SDP-ID existence and reception of reply.

| Field | Description | Values |
|---|--|--|
| Request Result | The result of the sdp-ping request message. | Sent - Request Timeout
Sent - Request Terminated
Sent - Reply Received
Not Sent - Non-Existent Local SDP-ID
Not Sent - Local SDP-ID Down |
| Originating SDP-ID | The originating SDP-ID specified by orig-sdp . | <i>orig-sdp-id</i> |
| Originating SDP-ID Administrative State | The local administrative state of the originating SDP-ID. If the SDP-ID has been shutdown, Admin-Down is displayed. If the originating SDP-ID is in the no shutdown state, Admin-Up is displayed. If the <i>orig-sdp-id</i> does not exist, Non-Existent is displayed. | Admin-Up
Admin-Down
Non-Existent |

Service Assurance Agent (SAA) Commands

| Field | Description | Values |
|---|--|--|
| Originating SDP-ID
Operating State | The local operational state of the originating SDP-ID. If <i>orig-sdp-id</i> does not exist, N/A will be displayed. | Oper-Up
Oper-Down
N/A |
| Originating SDP-ID
Path MTU | The local path-mtu for <i>orig-sdp-id</i> . If <i>orig-sdp-id</i> does not exist locally, N/A is displayed. | <i>orig-path-mtu</i>
N/A |
| Responding SDP-ID | The SDP-ID requested as the far-end path to respond to the sdp-ping request. If resp-sdp is not specified, the responding router will not use an SDP-ID as the return path and N/A will be displayed. | <i>resp-sdp-id</i>
N/A |
| Responding SDP-ID
Path Used | Displays whether the responding 7210 SAS M used the responding <i>sdp-id</i> to respond to the sdp-ping request. If <i>resp-sdp-id</i> is a valid, operational SDP-ID, it must be used for the SDP echo reply message. If the far-end uses the responding <i>sdp-id</i> as the return path, Yes will be displayed. If the far-end does not use the responding <i>sdp-id</i> as the return path, No will be displayed. If resp-sdp is not specified, N/A will be displayed. | Yes
No
N/A |
| Responding SDP-ID
Administrative State | The administrative state of the responding <i>sdp-id</i> . When <i>resp-sdp-id</i> is administratively down, Admin-Down will be displayed. When <i>resp-sdp-id</i> is administratively up, Admin-Up will be displayed. When <i>resp-sdp-id</i> exists on the far-end 7210 SAS M but is not valid for the originating router, Invalid is displayed. When <i>resp-sdp-id</i> does not exist on the far-end router, Non-Existent is displayed. When resp-sdp is not specified, N/A is displayed. | Admin-Down
Admin-Up
Invalid
Non-Existent
N/A |

| Field | Description | Values |
|-------------------------------------|--|--|
| Responding SDP-ID Operational State | The operational state of the far-end <i>sdp-id</i> associated with the return path for <i>service-id</i> . When a return path is operationally down, Oper-Down is displayed. If the return <i>sdp-id</i> is operationally up, Oper-Up is displayed. If the responding <i>sdp-id</i> is non-existent, N/A is displayed. | Oper-Up
Oper-Down
N/A |
| Responding SDP-ID Path MTU | The remote path-mtu for <i>resp-sdp-id</i> . If <i>resp-sdp-id</i> does not exist remotely, N/A is displayed | <i>resp-path-mtu</i>
N/A |
| Local Service IP Address | The local system IP address used to terminate remotely configured <i>sdp-ids</i> (as the <i>sdp-id</i> far-end address). If an IP address has not been configured to be the system IP address, N/A is displayed. | <i>system-ip-addr</i>
N/A |
| Local Service IP Interface Name | The name of the local system IP interface. If the local system IP interface has not been created, N/A is displayed. | <i>system-interface-name</i>
N/A |
| Local Service IP Interface State | The state of the local system IP interface. If the local system IP interface has not been created, Non-Existent is displayed. | Up
Down
Non-Existent |
| Expected Far End Address | The expected IP address for the remote system IP interface. This must be the far-end address configured for the <i>orig-sdp-id</i> . | <i>orig-sdp-far-end-addr</i>
<i>dest-ip-addr</i>
N/A |
| Actual Far End Address | The returned remote IP address. If a response is not received, the displayed value is N/A. If the far-end service IP interface is down or non-existent, a message reply is not expected. | <i>resp-ip-addr</i>
N/A |
| Responders Expected Far End Address | The expected source of the originators <i>sdp-id</i> from the perspective of the remote terminating the <i>sdp-id</i> . If the far-end cannot detect the expected source of the ingress <i>sdp-id</i> , N/A is displayed. | <i>resp-rec-tunnel-far-end-addr</i>
N/A |
| Round Trip Time | The round trip time between SDP echo request and the SDP echo reply. If the request is not sent, times out or is terminated, N/A is displayed. | <i>delta-request-reply</i>
N/A |

Multiple Response Connectivity Tests — When the connectivity test count is greater than one (1), a single line is displayed per SDP echo request send attempt.

The request number is a sequential number starting with 1 and ending with the last request sent, incrementing by one (1) for each request. This should not be confused with the *message-id* contained in each request and reply message.

A response message indicates the result of the message request. Following the response message is the round trip time value. If any reply is received, the round trip time is displayed.

Service Assurance Agent (SAA) Commands

After the last reply has been received or response timed out, a total is displayed for all messages sent and all replies received. A maximum, minimum and average round trip time is also displayed. Error response and timed out requests do not apply towards the average round trip time.

Multiple Response Round Trip Connectivity Test Sample Output

```
*A:DUT-A# oam sdp-ping 101 resp-sdp 102
Err SDP-ID Info                Local                Remote
-----
SDP-ID:                        101                102
Administrative State:          Up                  Up
Operative State:               Up                  Up
Path MTU:                      9186               N/A
Response SDP Used:             Yes
IP Interface State:            Up
Actual IP Address:             10.20.1.1          10.20.1.2
Expected Peer IP:              10.20.1.2          10.20.1.1
Forwarding Class               be                  be
Profile                         Out                 Out

Request Result: Sent - Reply Received
RTT: 10(ms)

*A:DUT-A# oam sdp-ping 101 resp-sdp 102 count 10
Request      Response      RTT
-----
1            Success       10ms
2            Success       0ms
3            Success       0ms
4            Success       0ms
5            Success       0ms
6            Success       0ms
7            Success       0ms
8            Success       0ms
9            Success       0ms
10           Success       0ms

Sent: 10      Received: 10
Min: 0ms     Max: 10ms    Avg: 1ms
*A:DUT-A#
```

vccv-ping

Note: This command is not supported on 7210 SAS devices configured in Access uplink mode.

Syntax **vccv-ping** *sdp-id:vc-id* [*src-ip-address ip-addr dst-ip-address ip-addr pw-id pw-id*][*reply-mode {ip-routed|control-channel}*] [*fc fc-name*] [*size octets*] [*send-count send-count*] [*timeout timeout*] [*interval interval*] [*tll vc-label-ttl*]

vccv-ping saii-type2 *global-id:prefix:ac-id* **taii-type2** *global-id:prefix:ac-id* [**reply-mode** *ip-routed|control-channel*] [**src-ip-address** *ip-addr dst-ip-address ip-addr*]
vccv-ping spoke-sdp-fec *spoke-sdp-fec-id* [**reply-mode** *ip-routed|control-channel*] [**saii-type2** *global-id:prefix:ac-id* **taii-type2** *global-id:prefix:ac-id*] [**src-ip-address** *ip-addr dst-ip-address ip-addr*]
options common to all vccv-ping cases: [**count** *send-count*] [**fc** *fc-name* [**profile** *in|out*]]
[**interval** *interval*] [**size** *octets*] [**timeout** *timeout*] [**ttl** *vc-label-ttl*]

Context oam
config>saa>test

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command configures a Virtual Circuit Connectivity Verification (VCCV) ping test. A vccv-ping test checks connectivity of a VLL inband. It checks to verify that the destination (target) PE is the egress for the Layer 2 FEC. It provides for a cross-check between the dataplane and the control plane. It is inband which means that the vccv-ping message is sent using the same encapsulation and along the same path as user packets in that VLL. The vccv-ping test is the equivalent of the lsp-ping test for a VLL service. The vccv-ping reuses an lsp-ping message format and can be used to test a VLL configured over an MPLS.

Note that VCCV ping can be initiated on TPE or SPE. If initiated on the SPE, the **reply-mode** parameter must be used with the ip-routed value. The ping from the TPE can have either values or can be omitted, in which case the default value is used.

If a VCCV ping is initiated from TPE to neighboring a SPE (one segment only) it is sufficient to only use the *sdpid:vcid* parameter. However, if the ping is across two or more segments, at least the *sdpid:vcid*, **src-ip-address** *ip-addr*, **dst-ip-address** *ip-addr*, **ttl** *vc-label-ttl* and **pw-id** *pw-id* parameters are used where:

- The *src-ip-address* is system IP address of the router preceding the destination router.
- The *pwid* is actually the VC ID of the last pseudowire segment.
- The *vc-label-ttl* must have a value equal or higher than the number of pseudowire segments.

Note that VCCV ping is a multi-segment pseudowire. For a single-hop pseudowire, only the peer VCCV CC bit of the control word is advertised when the control word is enabled on the pseudowire. VCCV ping on multi-segment pseudowires require that the control word be enabled in all segments of the VLL.

If the control word is not enabled on spoke SDP it will not be signaled peer VCCV CC bits to the far end, consequently VCCV ping cannot be successfully initiated on that specific spoke SDP.

Parameters *sdp-id:vc-id* — The VC ID of the pseudowire being tested must be indicated with this parameter. The VC ID needs to exist on the local router and the far-end peer needs to indicate that it supports VCCV to allow the user to send vccv-ping message.

Values 1 — 17407:1 — 4294967295

spoke-sdp-fec *spoke-sdp-fec-id* — If a FEC 129 PW is being tested, then its *spoke-sdp-fec-id* must be indicated with this parameter. The *spoke-sdp-fec-id* must exist on the local router and the far-end peer needs to indicate that it supports VCCV to allow the user to send vccv-ping messages.

spoke-sdp-fec is mutually exclusive with the *sdp-id:vc-id* parameter.

Values 1 — 4294967295

Service Assurance Agent (SAA) Commands

src-ip-address *ip-addr* — Specifies the source IP address.

ipv4-address: a.b.c.d**control-channel** {**ipv4** | **non-ip**} — Specifies the encapsulation format to use for the VCCV ping echo request and echo reply packet.

Values **ipv4** — IPv4 encapsulation in an IPv4 pseudowire associated channel (channel type 0x0021)

Values **non-ip** — MPLS-TP encapsulation without UDP/IP headers, in pseudowire associated channel using channel type 0x025.

Default non-ip

saii-type2 *global-id:prefix:ac-id* — If a FEC129 AII Type 2 pseudowire is being tested, then the source attachment individual identifier (SAII) must be indicated.

The **saii-type2** parameter is mutually exclusive with the *sdp-id:vc-id* parameter.

Syntax: *global-id* — The global ID of this T-PE node.

Values: 1 — 4294967295

prefix — The prefix on this T-PE node that the spoke-SDP is associated with.

ac-id — An unsigned integer representing a locally unique identifier for the spoke-SDP.

Values: 1 — 4294967295

taii-type2 *global-id:prefix:ac-id* — If a FEC129 AII Type 2 pseudowire is being tested, then the target attachment individual identifier (TAII) must be indicated. The **taii-type2** parameter is mutually exclusive with *sdp-id:vc-id*.

Syntax: *global-id* — The global ID of the far end T-PE node of the FEC129 pseudowire.

Values: 1 — 4294967295

prefix — The prefix on far end T-PE node that the pseudowire being tested is associated with.

Values: ipv4-formatted address: a.b.c.d

ac-id — An unsigned integer representing a locally unique identifier for the pseudowire being tested at the far end T-PE.

Values: 1 — 4294967295

dst-ip-address *ip-address* — Specifies the destination IP address.

Values **ipv4-address:** a.b.c.d

ipv6-address - x:x:x:x:x:x:x (eight 16-bit pieces)

pw-id *pw-id* — Specifies the pseudowire ID to be used for performing a **vccv-ping** operation. The pseudowire ID is a non-zero 32-bit connection ID required by the FEC 128, as defined in RFC 4379, *Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures*.

reply-mode {**ip-routed** | **control-channel**} — The reply-mode parameter indicates to the far-end how to send the reply message. The option **control-channel** indicates a reply mode in-band using vccv control channel.

Default control-channel

fc *fc-name* — The **fc** parameter is used to indicate the forwarding class of the MPLS echo request packets. The actual forwarding class encoding is controlled by the network egress LSP-EXP mappings.

The LSP-EXP mappings on the receive network interface controls the mapping back to the internal for-

warding class used by the far-end 7210 SAS M that receives the message request. The egress mappings of the egress network interface on the far-end router controls the forwarding class markings on the return reply message. The LSP-EXP mappings on the receive network interface controls the mapping of the message reply back at the originating SR.

Default be

Values be, l2, af, l1, h2, ef, h1, nc

timeout *seconds* — The timeout parameter, in seconds, expressed as a decimal integer. This value is used to override the default timeout value and is the amount of time that the router will wait for a message reply after sending the message request. Upon the expiration of message timeout, the requesting router assumes that the message response will not be received. A 'request timeout' message is displayed by the CLI for each message request sent that expires. Any response received after the request times out will be silently discarded.

Default 5

Values 1 — 10

interval *seconds* — The interval parameter in seconds, expressed as a decimal integer. This parameter is used to override the default request message send interval and defines the minimum amount of time that must expire before the next message request is sent.

If the interval is set to 1 second, and the timeout value is set to 10 seconds, then the maximum time between message requests is 10 seconds and the minimum is 1 second. This depends upon the receipt of a message reply corresponding to the outstanding message request.

Default 1

Values 1 — 10

size *octets* — The VCCV ping echo request packet size in octets, expressed as a decimal integer. The request payload is padded with zeroes to the specified size.

Default 88

Values 88 — 9198

send-count *send-count* — The number of messages to send, expressed as a decimal integer. The count parameter is used to override the default number of message requests sent. Each message request must either timeout or receive a reply before the next message request is sent. The message interval value must be expired before the next message request is sent.

Default 1

Values 1 — 100

ttl *vc-label-ttl* — Specifies the time-to-live value for the vc-label of the echo request message. The outer label TTL is still set to the default of 255 regardless of this value.

vccv-trace

Syntax `vccv-trace sdp-id:vc-id [fc fc-name [profile {in | out}]] [size octets] [reply-mode ip-routed|control-channel] [probe-count probe-count] [timeout timeout] [interval interval] [min-ttl min-vc-label-ttl] [max-ttl max-vc-label-ttl] [max-fail no-response-count] [detail]`
`vccv-trace saii-type2 global-id:prefix:ac-id taii-type2 global-id:prefix:ac-id [reply-mode ip-routed|control-channel]`
`vccv-trace spoke-sdp-fec spoke-sdp-fec-id [reply-mode ip-routed| control-channel] [saii-type2 global-id:prefix:ac-id taii-type2 global-id:prefix:ac-id]`
options common to all vccv-trace cases: `[detail] [fc fc-name [profile in|out]] [interval interval-value] [max-fail no-response-count] [max-ttl max-vc-label-ttl] [min-ttl min-vc-label-ttl] [probe-count probe-count] [size octets] [timeout timeout-value]`

Context oam
 config>saa>test>type

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command configures a Virtual Circuit Connectivity Verification (VCCV) automated trace test. The automated VCCV-trace can trace the entire path of a PW with a single command issued at the T-PE or at an S-PE. This is equivalent to LSP-Trace and is an iterative process by which the source T-PE or S-PE node sends successive VCCV-Ping messages with incrementing the TTL value, starting from TTL=1.

In each iteration, the T-PE builds the MPLS echo request message in a way similar to vccv-ping. The first message with TTL=1 will have the next-hop S-PE T-LDP session source address in the Remote PE Address field in the PW FEC TLV. Each S-PE which terminates and processes the message will include in the MPLS echo reply message the FEC 128 TLV corresponding the PW segment to its downstream node. The source T-PE or S-PE node can then build the next echo reply message with TTL=2 to test the next-next hop for the MS-PW. It will copy the FEC TLV it received in the echo reply message into the new echo request message. The process is terminated when the reply is from the egress T-PE or when a timeout occurs.

The user can specify to display the result of the VCCV-trace for a fewer number of PW segments of the end-to-end MS-PW path. In this case, the min-ttl and max-ttl parameters are configured accordingly. However, the T-PE/S-PE node will still probe all hops up to min-ttl in order to correctly build the FEC of the desired subset of segments.

Parameters `sdpid:vcid` — The VC ID of the pseudowire being tested must be indicated with this parameter. The VC ID needs to exist on the local 7210 SAS M and the far-end peer needs to indicate that it supports VCCV to allow the user to send vccv-ping message.

Values 1-17407:1 — 4294967295

reply-mode {*ip-routed* | *control-channel*} — The reply-mode parameter indicates to the far-end how to send the reply message. The option control-channel indicates a reply mode in-band using vccv control channel.

Note that when a VCCV trace message is originated from an S-PE node, the user should use the IPv4 reply mode as the replying node does not know how to set the TTL to reach the sending S-PE node. If the user attempts this, a warning is issued to use the ipv4 reply mode.

Default control-channel

control-channel {**none** | **non-ip**} — The encapsulation format to use for the VCCV ping echo request and echo reply packet.

Values none — IPv4 encapsulation in an IPv4 pseudowire associated channel (channel type 0x0021)

Values non-ip — MPLS-TP encapsulation without UDP/IP headers, in pseudowire associated channel using channel type 0x025.

Default non-ip

spoke-sdp-fec *spoke-sdp-fec-id* — If a FEC 129 PW is being tested, then its *spoke-sdp-fec-id* must be indicated with this parameter. The *spoke-sdp-fec-id* must exist on the local router and the far-end peer needs to indicate that it supports VCCV to allow the user to send vccv-ping messages.

spoke-sdp-fec is mutually exclusive with the *sdp-id:vc-id* parameter.

Values 1 — 4294967295

saii-type2 *global-id:prefix:ac-id* — If a FEC129 AII Type 2 pseudowire is being tested, then the source attachment individual identifier (SAII) must be indicated.

The **saii-type2** parameter is mutually exclusive with the *sdp-id:vc-id* parameter.

Syntax: *global-id* — The global ID of this 7210 T-PE node.

Values: 1 — 4294967295

prefix — The prefix on this 7210 T-PE node that the spoke-SDP is associated with.

ac-id — An unsigned integer representing a locally unique identifier for the spoke-SDP.

Values: 1 — 4294967295

taii-type2 *global-id:prefix:ac-id* — If a FEC129 AII Type 2 pseudowire is being tested, then the target attachment individual identifier (TAII) must be indicated. The **taii-type2** parameter is mutually exclusive with *sdp-id:vc-id*.

Syntax: *global-id* — The global ID of the far end T-PE of the FEC129 pseudowire.

Values: 1 — 4294967295

prefix — The prefix on far end T-PE that the pseudowire being tested is associated with.

Values: ipv4-formatted address: a.b.c.d

ac-id — An unsigned integer representing a locally unique identifier for the pseudowire being tested at the far end T-PE.

Values: 1 — 4294967295

fc *fc-name* [**profile** {**in** | **out**}] — The **fc** and **profile** parameters are used to indicate the forwarding class of the VCCV trace echo request packets. The actual forwarding class encoding is controlled by the network egress LSP-EXP mappings.

The LSP-EXP mappings on the receive network interface controls the mapping back to the internal forwarding class used by the far-end router that receives the message request. The egress mappings of the egress network interface on the far-end router controls the forwarding class markings on the return reply message. The LSP-EXP mappings on the receive network interface controls the mapping of the message reply back at the originating router.

fc-name — The forwarding class of the VCCV trace echo request encapsulation.

Service Assurance Agent (SAA) Commands

Default be

Values be, l2, af, ll, h2, ef, h1, nc

profile {in | out} — The profile state of the VCCV trace echo request encapsulation.

Default out

size *octets* — The VCCV ping echo request packet size in octets, expressed as a decimal integer. The request payload is padded with zeroes to the specified size.

Default 88

Values 88 — 9198

probe-count *probe-count* — The number of VCCV trace echo request messages to send per TTL value.

Default 1

Values 1 — 10

timeout *timeout* — The timeout parameter in seconds, expressed as a decimal integer. This value is used to override the default timeout value and is the amount of time that the router will wait for a message reply after sending the message request. Upon the expiration of message timeout, the requesting router assumes that the message response will not be received. A request timeout message is displayed by the CLI for each message request sent that expires. Any response received after the request times out will be silently discarded.

Default 3

Values 1 — 60

interval *interval* — The interval parameter in seconds, expressed as a decimal integer. This parameter is used to override the default request message send interval and defines the minimum amount of time that must expire before the next message request is sent.

If the interval is set to 1 second, and the timeout value is set to 10 seconds, then the maximum time between message requests is 10 seconds and the minimum is 1 second. This depends upon the receipt of a message reply corresponding to the outstanding message request.

Default 1

Values 1 — 255

min-ttl *min-vc-label-ttl* — The TTL value for the VC label of the echo request message for the first hop of the MS-PW for which the results are to be displayed. This is expressed as a decimal integer. Note that the outer label TTL is still set to the default of 255 regardless of the value of the VC label.

Default 1

Values 1 — 255

max-ttl *max-vc-label-ttl* — The TTL value for the VC label of the echo request message for the last hop of the MS-PW for which the results are to be displayed. This is expressed as a decimal integer. Note that the outer label TTL is still set to the default of 255 regardless of the value of the VC label.

Default 8

Values 1 — 255

max-fail *no-response-count* — The maximum number of consecutive VCCV trace echo requests, expressed as a decimal integer that do not receive a reply before the trace operation fails for a given TTL value.

Default 5

Values 1 — 255

OAM SAA Commands

saa

Syntax `saa test-name [owner test-owner] {start | stop} [no-accounting]`

Context oam

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

Use this command to start or stop an SAA test.

test-name — Name of the SAA test. The test name must already be configured in the **config>saa>test** context.

owner *test-owner* — Specifies the owner of an SAA operation up to 32 characters in length.

Values If a *test-owner* value is not specified, tests created by the CLI have a default owner “TiMOS CLI”.

start — This keyword starts the test. A test cannot be started if the same test is still running.

A test cannot be started if it is in a shut-down state. An error message and log event will be generated to indicate a failed attempt to start an SAA test run. A test cannot be started if it is in a continuous state.

stop — This keyword stops a test in progress. A test cannot be stopped if it is not in progress. A log message will be generated to indicate that an SAA test run has been aborted. A test cannot be stopped if it is in a continuous state.

no-accounting — This parameter disables the recording results in the accounting policy. If **no-accounting** is specified, the MIB record produced at the end of the test will not be added to the accounting file. It will however use up one of the three MIB rows available for the accounting module to be collected.

LDP Treetrace Commands

ldp-treetrace

Syntax `ldp-treetrace {prefix ip-prefix/mask} [max-ttl tll-value] [max-path max-paths] [timeout timeout] [retry-count retry-count] [fc fc-name [profile profile]] [downstream-map-tlv {dsmap|ddmap}]`

Context oam

Description **Platforms Supported:** 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-R6, and 7210 SAS-R12.

This command allows the user to perform a single run of the LDP ECMP OAM tree trace to discover all ECMP paths of an LDP FEC.

Parameters **prefix** *ip-prefix/mask* — Specifies the address prefix and subnet mask of the target BGP IPv4 label route.
max-ttl *max-label-ttl* — The maximum TTL value in the MPLS label for the LSP trace test, expressed as a decimal integer.

Values 1 — 255

Default 30

max-paths *max-paths* — The maximum number of paths for a ldp-treetrace test, expressed as a decimal integer.

Values 1 — 255

Default 128

timeout *timeout* — The **timeout** parameter in seconds, expressed as a decimal integer. This value is used to override the default timeout value and is the amount of time that the router will wait for a message reply after sending the message request. Upon the expiration of message timeout, the requesting router assumes that the message response will not be received. Any response received after the request times out will be silently discarded.

Values 1 — 60

Default 3

fc *fc-name* — The fc parameters are used to indicate the forwarding class of the MPLS echo request packet.

When an MPLS echo request packet is generated in CPM and is forwarded to the outgoing interface, the packet is queued in the egress network queue corresponding to the specified fc and profile parameter values. The marking of the packet's EXP is dictated by the LSP-EXP mappings on the outgoing interface.

When the MPLS echo request packet is received on the responding node, The fc parameter values are dictated by the LSP-EXP mappings of the incoming interface.

When an MPLS echo reply packet is generated in CPM and is forwarded to the outgoing interface, the packet is queued in the egress network queue corresponding to the fc parameter values determined by the classification of the echo request packet, which is being replied to, at the incoming interface. The marking of the packet's EXP is dictated by the LSP-EXP mappings on the outgoing interface. The TOS byte is not modified. [Table 20](#) summarizes this behavior:

Table 20: Request Packet and Behavior

| | |
|-------------------------------------|---|
| cpm (sender node) | <p>echo request packet:</p> <ul style="list-style-type: none"> • packet{tos=1, fc1} • fc1 and profile1 are as entered by user in OAM command or default values • tos1 as per mapping of {fc1} to IP precedence in network egress QoS policy of outgoing interface |
| outgoing interface (sender node) | <p>echo request packet:</p> <ul style="list-style-type: none"> • pkt queued as {fc1} • ToS field=tos1 not remarked • EXP=exp1, as per mapping of {fc1} to EXP in network egress QoS policy of outgoing interface |
| Incoming interface (responder node) | <p>echo request packet:</p> <ul style="list-style-type: none"> • packet{tos1, exp1} • exp1 mapped to {fc2} as per classification in network QoS policy of incoming interface |
| cpm (responder node) | <p>echo reply packet:</p> <ul style="list-style-type: none"> • packet{tos=1, fc2} |
| outgoing interface (responder node) | <p>echo reply packet:</p> <ul style="list-style-type: none"> • pkt queued as {fc2} • ToS field= tos1 not remarked (reply inband or out-of-band) • EXP=exp2, if reply is inband, remarked as per mapping of {fc2, profile2} to EXP in network egress QoS policy of outgoing interface |
| Incoming interface (sender node) | <p>echo reply packet:</p> <ul style="list-style-type: none"> • packet{tos1, exp2} • exp2 mapped to {fc1, profile1} as per classification in network QoS policy of incoming interface |

Values be, l2, af, ll, h2, ef, h1, nc

Default be

profile *profile* — The profile state of the MPLS echo request packet.

Values in, out

Default out

retry-count *retry-count* — Specifies the maximum number of consecutive MPLS echo requests, expressed as a decimal integer that do not receive a reply before the trace operation fails for a given TTL.

Values 1 — 255

Default 5

downstream-map-tlv {**dsmap** | **ddmap**} — Specifies which format of the downstream mapping TLV to use in the LSP trace packet. The DSMAP TLV is the original format in RFC 4379. The DDMAP is the new enhanced format specified in RFC 6424.

Default Inherited from global configuration of downstream mapping TLV in option **mpls-echo-request-downstream-map** {**dsmap** | **ddmap**}.

Sample Output

```
*A:Dut-A# oam ldp-treetrace prefix 10.20.1.6/32
```

```
ldp-treetrace for Prefix 10.20.1.6/32:
```

```
      127.0.0.1, ttl = 3 dst =      127.1.0.255 rc = EgressRtr status = Done
Hops:      127.0.0.1      127.0.0.1
```

```
      127.0.0.1, ttl = 3 dst =      127.2.0.255 rc = EgressRtr status = Done
Hops:      127.0.0.1      127.0.0.1
```

```
ldp-treetrace discovery state: Done
ldp-treetrace discovery status: ' OK '
Total number of discovered paths: 2
Total number of failed traces: 0
```

test-oam

Syntax test-oam

Context config

Description **Platforms Supported:** 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-R6, and 7210 SAS-R12.

This command enables the context to configure Operations, Administration, and Maintenance test parameters.

ldp-tree trace

Syntax [no] ldp-tree trace

Context config>test-oam

Description **Platforms Supported:** 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-R6, and 7210 SAS-R12.

This command creates the context to configure the LDP ECMP OAM tree trace which consists of an LDP ECMP path discovery and an LDP ECMP path probing features.

The **no** option deletes the configuration for the LDP ECMP OAM tree discovery and path probing under this context.

Sample Output

```
*A:7210SAS# show test-oam ldp-tree trace prefix 7.7.7.7/32 detail
Discovery State : Done                Last Discovered : 01/01/2000 23:38:37
Discovery Status : ' OK '
Discovered Paths : 4
Probe State      : OK                 Failed Hops      : 0
                                           Failed Probes    : 0
```

```
=====
Discovered Paths
=====
```

| PathDest | DiscoveryTtl | Egr-NextHop | ProbeState | Remote-RtrAddr | ProbeTmOutCnt | Discovery-time | RtnCode |
|-------------|--------------|-------------|------------|----------------|---------------|---------------------|-----------|
| 127.1.0.255 | 1 | 37.37.37.7 | OK | 37.37.37.7 | 0 | 01/01/2000 23:38:36 | EgressRtr |
| 127.2.0.255 | 1 | 38.38.38.7 | OK | 38.38.38.7 | 0 | 01/01/2000 23:38:36 | EgressRtr |
| 127.3.0.255 | 1 | 39.39.39.7 | OK | 39.39.39.7 | 0 | 01/01/2000 23:38:36 | EgressRtr |
| 127.4.0.255 | 1 | 40.40.40.7 | OK | 40.40.40.7 | 0 | 01/01/2000 23:38:36 | EgressRtr |

```
=====
```

```
*A:7210SAS# show test-oam ldp-tree trace prefix 7.7.7.7/32
```

Sample output over a numbered IP interface

```
*A:Dut-B# oam ldp-tree trace prefix 10.20.1.5/32
```

```
ldp-tree trace for Prefix 10.20.1.5/32:
```

```
10.10.131.2, ttl = 2 dst = 127.1.0.253 rc = EgressRtr status = Done
Hops: 11.1.0.2
```

```
10.10.132.2, ttl = 2 dst = 127.1.0.255 rc = EgressRtr status = Done
Hops: 11.1.0.2
```

```
10.10.131.2, ttl = 2 dst = 127.2.0.255 rc = EgressRtr status = Done
Hops: 11.2.0.2
```

```

10.10.132.2, ttl = 2 dst = 127.2.0.253 rc = EgressRtr status = Done
Hops: 11.2.0.2

ldp-treetrace discovery state: Done
ldp-treetrace discovery status: ' OK '
Total number of discovered paths: 4
Total number of failed traces: 0

```

fc

Syntax `fc fc-name`
`no fc`

Context `config>test-oam>ldp-treetrace`

Description **Platforms Supported:** 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-R6, and 7210 SAS-R12.

This command indicates the forwarding class of the MPLS echo request packet.

When an MPLS echo request packet is generated in CPM and is forwarded to the outgoing interface, the packet is queued in the egress network queue corresponding to the specified fc parameter values. The marking of the packet's EXP is dictated by the LSP-EXP mappings on the outgoing interface.

When the MPLS echo request packet is received on the responding node, The fc parameter values are dictated by the LSP-EXP mappings of the incoming interface.

When an MPLS echo reply packet is generated in CPM and is forwarded to the outgoing interface, the packet is queued in the egress network queue corresponding to the fc parameter values determined by the classification of the echo request packet, which is being replied to, at the incoming interface. The marking of the packet's EXP is dictated by the LSP-EXP mappings on the outgoing interface. The TOS byte is not modified. The following table summarizes this behavior:

Table 21: Request Packet and Behavior

| | |
|-------------------|---|
| cpm (sender node) | echo request packet: <ul style="list-style-type: none"> • packet{tos=1, fc1} • fc1 and profile1 are as entered by user in OAM command or default values • tos1 as per mapping of {fc1} to IP precedence in network egress QoS policy of outgoing interface |
|-------------------|---|

Table 21: Request Packet and Behavior (Continued)

| | |
|-------------------------------------|---|
| outgoing interface (sender node) | <p>echo request packet:</p> <ul style="list-style-type: none"> • pkt queued as {fc1} • ToS field=tos1 not remarked • EXP=exp1, as per mapping of {fc1} to EXP in network egress QoS policy of outgoing interface |
| Incoming interface (responder node) | <p>echo request packet:</p> <ul style="list-style-type: none"> • packet{tos1, exp1} • exp1 mapped to {fc2} as per classification in network QoS policy of incoming interface |
| cpm (responder node) | <p>echo reply packet:</p> <ul style="list-style-type: none"> • packet{tos=1, fc2} |
| outgoing interface (responder node) | <p>echo reply packet:</p> <ul style="list-style-type: none"> • pkt queued as {fc2} • ToS filed= tos1 not remarked (reply inband or out-of-band) • EXP=exp2, if reply is inband, remarked as per mapping of {fc2, profile2} to EXP in network egress QoS policy of outgoing interface |
| Incoming interface (sender node) | <p>echo reply packet:</p> <ul style="list-style-type: none"> • packet{tos1, exp2} • exp2 mapped to {fc1, profile1} as per classification in network QoS policy of incoming interface |

Default be

Parameters *fc-name* — Specifies the forwarding class of the MPLS echo request packets.

Values be, l2, af, ll, h2, ef, h1, nc

path-discovery

Syntax path-discovery

Context config>test-oam>ldp-tree-trace

Description **Platforms Supported:** 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-R6, and 7210 SAS-R12.

This command creates the context to configure the LDP ECMP OAM path discovery.

The ingress LER builds the ECMP tree for a given FEC (egress LER) by sending LSP Trace messages and including the LDP IPv4 Prefix FEC TLV as well as the downstream mapping TLV. It inserts an IP address range drawn from the 127/8 space. When received by the downstream LSR, it uses this range to determine which ECMP path is exercised by any IP address or a sub-range of addresses within that range based on its internal hash routine. When the MPLS Echo reply is received by the ingress LER, it records this information and proceeds with the next echo request message targeted for a node downstream of the first LSR node along one of the ECMP paths. The sub-range of IP addresses indicated in the initial reply is used since the objective is to have the LSR downstream of the ingress LER pass this message to its downstream node along the first ECMP path.

The user configures the frequency of running the tree discovery using the command **config>test-oam>ldp-treetrace>path-discovery> interval**.

The ingress LER gets the list of FECs from the LDP FEC database. New FECs will be added to the discovery list at the next tree discovery and not when they are learnt and added into the FEC database. The user can configure FECs he/she wishes to include or exclude using a policy profile by applying the command **config>test-oam>ldp-treetrace>path-discovery>policy-statement**.

interval

Syntax **interval** *minutes*
no interval

Context config>test-oam>ldp-treetrace>path-discovery

Description **Platforms Supported:** 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-R6, and 7210 SAS-R12.

This command configures the frequency of the LDP ECMP OAM path discovery. Every interval, the node will send LSP trace messages to attempt to discover the entire ECMP path tree for a given destination FEC.

The **no** option resets the interval to its default value.

Default 60

Parameters *minutes* — Specifies the number of minutes to wait before repeating the LDP tree auto discovery process.

Values 60 — 1440

max-path

Syntax **max-path** *max-paths*

Context config>test-oam>ldp-treetrace>path-discovery

Description **Platforms Supported:** 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-R6, and 7210 SAS-R12.

This command configures the maximum number of ECMP paths the path discovery will attempt to discover for each run every **interval** minutes.

LDP TreeTrace Commands

The **no** option resets the timeout to its default value.

Default 16

Parameters *max-paths* — Specifies the tree discovery maximum path.

Values 1 — 16

max-ttl

Syntax **max-ttl** *ttl-value*

Context config>test-oam>ldp-treetrace>path-discovery

Description **Platforms Supported:** 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-R6, and 7210 SAS-R12.

This command configures the maximum number of hops the path discovery will trace in the path of each FEC to be discovered.

The **no** option resets the timeout to its default value.

Default 255

Parameters *ttl-value* — Specifies the maximum label time-to-live value for an LSP trace request during the tree discovery.

Values 1 — 255

policy-statement

Syntax **policy-statement** *policy-name* [...(up to 5 max)]

Context config>test-oam>ldp-treetrace>path-discovery

Description **Platforms Supported:** 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-R6, and 7210 SAS-R12.

This command configures the FEC policy to determine which routes are imported from the LDP FEC database for the purpose of discovering its paths and probing them.

If no policy is specified, the ingress LER imports the full list of FECs from the LDP FEC database. New FECs will be added to the discovery list at the next path discovery and not when they are learnt and added into the FEC database. The maximum number of FECs to be discovered with path discovery is limited to 500.

The user can configure FECs he/she wishes to include or exclude.

Policies are configured in the **config>router>policy-options** context. A maximum of five policy names can be specified.

The **no** form of the command removes the policy from the configuration.

Default no policy-statement

Parameters *policy-name* — Specifies the route policy name to filter LDP imported address FECs. Allowed values are any string up to 32 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. The specified policy name(s) must already be defined.

retry-count

Syntax **retry-count** *retry-count*
no **retry-count**

Context config>oam-test>ldp-treetrace>path-discovery
config>oam-test>ldp-treetrace>path-probing

Description **Platforms Supported:** 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-R6, and 7210 SAS-R12.

In the path discovery phase of the LDP tree trace feature, this command configures the number of retransmissions of an LSP trace message to discover the path of an LDP FEC when no response is received within the **timeout** parameter.

In the path-probing phase of the LDP tree trace, this command configures the number of retransmissions of an LSP ping message to probe the path of an LDP FEC when no response is received within the **timeout** parameter.

The **no** option resets the retry count to its default value

Default 3

Parameters *retry-count* — Specifies the maximum number of consecutive timeouts allowed before failing a path probe (ping).

Values 1 — 10

timeout

Syntax **timeout** *timeout*
no **timeout**

Context config>test-oam>ldp-treetrace>path-discovery

Description **Platforms Supported:** 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-R6, and 7210 SAS-R12.

This command configures the time the node waits for the response to an LSP Trace message discovering the path of an LDP FEC before it declares failure. After consecutive failures equal to the **retry-count** parameter, the node gives up.

The **no** option resets the timeout to its default value.

LDP Treetrace Commands

Default 30

Parameters *timeout* — Specifies the timeout parameter, in seconds, within a range of 1 to 60, expressed as a decimal integer.

Values 1—60

path-probing

Syntax **path-probing**

Context config>test-oam>ldp-treeTrace

Description **Platforms Supported:** 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-R6, and 7210 SAS-R12.

This command creates the context to configure the LDP tree trace path probing phase.

The periodic path exercising runs in the background to test the LDP ECMP paths discovered by the path discovery capability. The probe used is an LSP Ping message with an IP address drawn from the sub-range of 127/8 addresses indicated by the output of the tree discovery for this FEC.

The user configures the frequency of running the path probes using the command **config>test-oam>ldp-treeTrace> path-probing> interval**. If an I/F is down on the ingress LER performing the LDP tree trace, then LSP Ping probes that normally go out this interface will not be sent but the ingress LER node will not raise alarms.

The LSP Ping routine should update the content of the MPLS echo request message, specifically the IP address, as soon as the LDP ECMP path discovery phase has output the results of a new computation for the path in question.

interval

Syntax **interval** *minutes*
no interval

Context config>test-oam>ldp-treeTrace>path-probing

Description **Platforms Supported:** 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-R6, and 7210 SAS-R12.

This command configures the frequency of the LSP Ping messages used in the path probing phase to probe the paths of all LDP FECs discovered by by the LDP tree trace path discovery.

The no option resets the interval to its default value.

Default 1

Parameters *minutes* — Specifies the number of minutes to probe all active ECMP paths for each LDP FEC.

Values 1 — 60

timeout

| | |
|--------------------|---|
| Syntax | timeout <i>timeout</i>
no timeout |
| Context | config>test-oam>ldp-treetrace>path-probing |
| Description | <p>Platforms Supported: 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-R6, and 7210 SAS-R12.</p> <p>This command configures the time the node waits for the response to an LSP Ping message probing the path of an LDP FEC before it declares failure. After consecutive failures equal to the retry-count parameter, the node gives up.</p> <p>The no option resets the timeout to its default value.</p> |
| Default | 1 |
| Parameters | <p><i>timeout</i> — Specifies the timeout parameter, in minutes, with a range of 1 to 3 minutes, expressed as a decimal integer.</p> <p>Values 1—3</p> |

mpls-echo-request-downstream-map

| | |
|--------------------|--|
| Syntax | mpls-echo-request-downstream-map { dsmap ddmap }
no mpls-echo-request-downstream-map |
| Context | config>test-oam |
| Description | <p>Platforms Supported: 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.</p> <p>This command specifies which format of the downstream mapping TLV to use in all LSP trace packets and LDP tree trace packets originated on this node. The Downstream Mapping (DSMAP) TLV is the original format in RFC 4379 and is the default value. The new Downstream Detailed Mapping (DDMAP) TLV is the new enhanced format specified in RFC 6424.</p> <p>This command applies to LSP trace of an RSVP P2P LSP, a MPLS-TP LSP, or LDP unicast FEC, and to LDP tree trace of a unicast LDP FEC. It does not apply to LSP trace of an RSVP P2MP LSP which always uses the DDMAP TLV.</p> <p>The global DSMAP/DDMAP setting impacts the behavior of both OAM LSP trace packets and SAA test packets of type lsp-trace and is used by the sender node when one of the following events occurs:</p> <ol style="list-style-type: none"> 1. An SAA test of type lsp-trace is created (not modified) and no value is specified for the per-test downstream-map-tlv {dsmap ddmap none} option. In this case, the SAA test downstream-map-tlv value defaults to the global mpls-echo-request-downstream-map value. 2. An OAM test of type lsp-trace test is executed and no value is specified for the per-test downstream-map-tlv {dsmap ddmap none} option. In this case, the OAM test downstream-map-tlv value defaults to the global mpls-echo-request-downstream-map value. |

A consequence of the rules above is that a change to the value of **mpls-echo-request-downstream-map** option does not affect the value inserted in the downstream mapping TLV of existing tests.

Following are the details of the processing of the new DDMAP TLV:

1. When either the DSMAP TLV or the DDMAP TLV is received in an echo request message, the responder node will include the same type of TLV in the echo reply message with the proper downstream interface information and label stack information.
2. If an echo request message without a Downstream Mapping TLV (DSMAP or DDMAP) expires at a node which is not the egress for the target FEC stack, the responder node always includes the DSMAP TLV in the echo reply message. This can occur in the following cases:
 - a. The user issues a LSP trace from a sender node with a **min-ttl** value higher than 1 and a **max-ttl** value lower than the number of hops to reach the egress of the target FEC stack. This is the sender node behavior when the global configuration or the per-test setting of the DSMAP/DDMAP is set to DSMAP.
 - b. The user issues a LSP ping from a sender node with a **tth** value lower than the number of hops to reach the egress of the target FEC stack. This is the sender node behavior when the global configuration of the DSMAP/DDMAP is set to DSMAP.
 - c. The behavior in (a) is changed when the global configuration or the per-test setting of the Downstream Mapping TLV is set to DDMAP. The sender node will include in this case the DDMAP TLV with the Downstream IP address field set to the all-routers multicast address as per Section 3.3 of RFC 4379. The responder node then bypasses the interface and label stack validation and replies with a DDMAP TLV with the correct downstream information for the target FEC stack.
3. A sender node never includes the DSMAP or DDMAP TLV in an lsp-ping message.

In addition to performing the same features as the DSMAP TLV, the new DDMAP TLV addresses the following scenarios:

1. Full validation of an LDP FEC stitched to a BGP IPv4 label route. In this case, the LSP trace message is inserted from the LDP LSP segment or from the stitching point.
2. Full validation of a BGP IPv4 label route stitched to an LDP FEC.
3. Full validation of an LDP FEC which is stitched to a BGP LSP and stitched back into an LDP FEC. In this case, the LSP trace message is inserted from the LDP segments or the or from the stitching points.
4. Full validation of an LDP FEC tunneled over an RSVP LSP using LSP trace.

In order to properly check a target FEC which is stitched to another FEC (stitching FEC) of the same or a different type, or which is tunneled over another FEC (tunneling FEC), it is necessary for the responding nodes to provide details about the FEC manipulation back to the sender node. This is achieved via the use of the new FEC stack change sub-TLV in the Downstream Detailed Mapping TLV (DDMAP) defined in RFC 6424.


When the user configures the use of the DDMAP TLV on a trace for an LSP that does not undergo stitching or tunneling operation in the network, the procedures at the sender and responder nodes are the same as in the case of the DSMAP TLV.

This feature however introduces changes to the target FEC stack validation procedures at the sender and responder nodes in the case of LSP stitching and LSP hierarchy. These changes pertain to the processing of the new FEC stack change sub-TLV in the new DDMAP TLV and the new return code of value 15 Label switched with FEC change.

The **no** form of this command reverts to the default behavior of using the DSMAP TLV in a LSP trace packet and LDP tree trace packet.

Default **dsmap**

Output **LDP-over-RSVP**



Testing LDP FEC of Node C with DSMAP TLV

```
-----
*A:Dut-A#
*A:Dut-A# oam lsp-trace prefix 10.20.1.3/32 downstream-map-tlv dsmap detail
lsp-trace to 10.20.1.3/32: 0 hops min, 0 hops max, 104 byte packets
1 10.20.1.2 rtt=3.90ms rc=8(DSRtrMatchLabel) rsc=1
   DS 1: ipaddr=10.10.4.4 ifaddr=10.10.4.4 iftype=ipv4Numbered MRU=1500
        label[1]=131068 protocol=3 (LDP)
2 10.20.1.4 rtt=5.69ms rc=8(DSRtrMatchLabel) rsc=1
   DS 1: ipaddr=10.10.9.6 ifaddr=10.10.9.6 iftype=ipv4Numbered MRU=1500
        label[1]=131066 protocol=3 (LDP)
3 10.20.1.6 rtt=7.88ms rc=8(DSRtrMatchLabel) rsc=1
   DS 1: ipaddr=10.10.10.5 ifaddr=10.10.10.5 iftype=ipv4Numbered MRU=1500
        label[1]=131060 protocol=3 (LDP)
4 10.20.1.5 rtt=23.2ms rc=8(DSRtrMatchLabel) rsc=1
   DS 1: ipaddr=10.10.5.3 ifaddr=10.10.5.3 iftype=ipv4Numbered MRU=1496
        label[1]=131071 protocol=3 (LDP)
5 10.20.1.3 rtt=12.0ms rc=3(EgressRtr) rsc=1
*A:Dut-A#
```

Testing LDP FEC of Node C with DDMAP TLV


```
-----
*A:Dut-A# oam lsp-trace prefix 10.20.1.3/32 downstream-map-tlv ddmmap detail
lsp-trace to 10.20.1.3/32: 0 hops min, 0 hops max, 136 byte packets
1 10.20.1.2 rtt=4.00ms rc=3(EgressRtr) rsc=2
1 10.20.1.2 rtt=3.48ms rc=8(DSRtrMatchLabel) rsc=1
   DS 1: ipaddr=10.10.4.4 ifaddr=10.10.4.4 iftype=ipv4Numbered MRU=1500
        label[1]=131068 protocol=3 (LDP)
2 10.20.1.4 rtt=5.34ms rc=8(DSRtrMatchLabel) rsc=1
   DS 1: ipaddr=10.10.9.6 ifaddr=10.10.9.6 iftype=ipv4Numbered MRU=1500
        label[1]=131066 protocol=3 (LDP)
3 10.20.1.6 rtt=7.78ms rc=8(DSRtrMatchLabel) rsc=1
   DS 1: ipaddr=10.10.10.5 ifaddr=10.10.10.5 iftype=ipv4Numbered MRU=1500
        label[1]=131060 protocol=3 (LDP)
4 10.20.1.5 rtt=12.8ms rc=15(LabelSwitchedWithFecChange) rsc=1
   DS 1: ipaddr=10.10.5.3 ifaddr=10.10.5.3 iftype=ipv4Numbered MRU=1496
        label[1]=131054 protocol=4 (RSVP-TE)
```

LDP Treetrace Commands

```

label[2]=131071 protocol=3(LDP)
fecchange[1]=PUSH fectype=RSVP IPv4 prefix=10.20.1.3 remotepeer=10.10.5.3
5 10.20.1.3 rtt=12.8ms rc=3(EgressRtr) rsc=2
5 10.20.1.3 rtt=13.4ms rc=3(EgressRtr) rsc=1
*A:Dut-A#

```



Testing LDP FEC of Node B with DDMAP TLV

```

-----
*A:Dut-D#
*A:Dut-D# oam lsp-trace prefix 10.20.1.2/32 downstream-map-tlv ddmmap detail
lsp-trace to 10.20.1.2/32: 0 hops min, 0 hops max, 108 byte packets
1 10.20.1.6 rtt=3.17ms rc=8(DSRtrMatchLabel) rsc=1
   DS 1: ipaddr=10.10.10.5 ifaddr=10.10.10.5 iftype=ipv4Numbered MRU=1500
       label[1]=131065 protocol=3(LDP)
2 10.20.1.5 rtt=8.27ms rc=15(LabelSwitchedWithFecChange) rsc=1
   DS 1: ipaddr=10.10.5.3 ifaddr=10.10.5.3 iftype=ipv4Numbered MRU=1496
       label[1]=131068 protocol=4(RSVP-TE)
       label[2]=131065 protocol=3(LDP)
       fecchange[1]=PUSH fectype=RSVP IPv4 prefix=10.20.1.1 remotepeer=10.10.5.3
3 10.20.1.3 rtt=9.50ms rc=8(DSRtrMatchLabel) rsc=2
   DS 1: ipaddr=10.10.2.1 ifaddr=10.10.2.1 iftype=ipv4Numbered MRU=1500
       label[1]=131068 protocol=4(RSVP-TE)
4 10.20.1.1 rtt=10.4ms rc=3(EgressRtr) rsc=2
4 10.20.1.1 rtt=10.2ms rc=15(LabelSwitchedWithFecChange) rsc=1
   DS 1: ipaddr=10.10.1.2 ifaddr=10.10.1.2 iftype=ipv4Numbered MRU=1496
       label[1]=131066 protocol=4(RSVP-TE)
       label[2]=131071 protocol=3(LDP)
       fecchange[1]=PUSH fectype=RSVP IPv4 prefix=10.20.1.2 remotepeer=10.10.1.2
5 10.20.1.2 rtt=13.7ms rc=3(EgressRtr) rsc=2
5 10.20.1.2 rtt=13.6ms rc=3(EgressRtr) rsc=1
*A:Dut-D#

```

Testing LDP FEC of Node F with DSMAP TLV

```

-----
*A:Dut-A# *A:Dut-A# oam lsp-trace prefix 10.20.1.6/32 downstream-map-tlv dsmap detail lsp-
trace to 10.20.1.6/32: 0 hops min, 0 hops max, 104 byte packets
1 10.20.1.2 rtt=2.65ms rc=8(DSRtrMatchLabel) rsc=1
2 10.20.1.3 rtt=4.89ms rc=8(DSRtrMatchLabel) rsc=1
3 10.20.1.4 rtt=6.49ms rc=5(DSMappingMismatched) rsc=1
*A:Dut-A#

```

Testing LDP FEC of Node F with DDMAP TLV

```


-----
*A:Dut-A# oam lsp-trace prefix 10.20.1.6/32 downstream-map-tlv ddmmap detail lsp-trace to
10.20.1.6/32: 0 hops min, 0 hops max, 108 byte packets
1 10.20.1.2 rtt=3.50ms rc=15(LabelSwitchedWithFecChange) rsc=1
   DS 1: ipaddr=10.10.3.3 ifaddr=10.10.3.3 iftype=ipv4Numbered MRU=1496
       label[1]=131068 protocol=3(LDP)

```

```

label[2]=131060 protocol=2 (BGP)
fecchange[1]=POP fectype=LDP IPv4 prefix=10.20.1.6 remotepeer=0.0.0.0
(Unknown)
fecchange[2]=PUSH fectype=BGP IPv4 prefix=10.20.1.6 remotepeer=10.20.1.5
fecchange[3]=PUSH fectype=LDP IPv4 prefix=10.20.1.5 remotepeer=10.10.3.3
2 10.20.1.3 rtt=6.53ms rc=15(LabelSwitchedWithFecChange) rsc=2
DS 1: ipaddr=10.10.11.4 ifaddr=10.10.11.4 iftype=ipv4Numbered MRU=1496
label[1]=131060 protocol=4 (RSVP-TE)
label[2]=131070 protocol=3 (LDP)
label[3]=131060 protocol=2 (BGP)
fecchange[1]=PUSH fectype=RSVP IPv4 prefix=10.20.1.4 remotepeer=10.10.11.4
3 10.20.1.4 rtt=7.94ms rc=3(EgressRtr) rsc=3
3 10.20.1.4 rtt=6.69ms rc=8(DSRtrMatchLabel) rsc=2
DS 1: ipaddr=10.10.6.5 ifaddr=10.10.6.5 iftype=ipv4Numbered MRU=1500
label[1]=131071 protocol=3 (LDP)
label[2]=131060 protocol=2 (BGP)
4 10.20.1.5 rtt=10.1ms rc=3(EgressRtr) rsc=2
4 10.20.1.5 rtt=8.97ms rc=15(LabelSwitchedWithFecChange) rsc=1
DS 1: ipaddr=10.10.10.6 ifaddr=10.10.10.6 iftype=ipv4Numbered MRU=1500
label[1]=131071 protocol=3 (LDP)
fecchange[1]=POP fectype=BGP IPv4 prefix=10.20.1.6 remotepeer=0.0.0.0
(Unknown)
fecchange[2]=PUSH fectype=LDP IPv4 prefix=10.20.1.6 remotepeer=10.10.10.6
5 10.20.1.6 rtt=11.8ms rc=3(EgressRtr) rsc=1 *A:Dut-A#

```




Testing BGP Label Route of Node E with DDMAP TLV


```

-----
*A:Dut-B# oam lsp-trace prefix 11.20.1.5/32 bgp-label downstream-map-tlv ddmapi detail lsp-
trace to 11.20.1.5/32: 0 hops min, 0 hops max, 124 byte packets
1 10.20.1.3 rtt=2.35ms rc=15(LabelSwitchedWithFecChange) rsc=2
DS 1: ipaddr=10.10.11.4 ifaddr=10.10.11.4 iftype=ipv4Numbered MRU=1496
label[1]=131060 protocol=4 (RSVP-TE)
label[2]=131070 protocol=3 (LDP)
label[3]=131070 protocol=2 (BGP)
fecchange[1]=PUSH fectype=RSVP IPv4 prefix=10.20.1.4 remotepeer=10.10.11.4
2 10.20.1.4 rtt=4.17ms rc=3(EgressRtr) rsc=3
2 10.20.1.4 rtt=4.50ms rc=8(DSRtrMatchLabel) rsc=2
DS 1: ipaddr=10.10.6.5 ifaddr=10.10.6.5 iftype=ipv4Numbered MRU=1500
label[1]=131071 protocol=3 (LDP)
label[2]=131070 protocol=2 (BGP)
3 10.20.1.5 rtt=7.78ms rc=3(EgressRtr) rsc=2
3 10.20.1.5 rtt=6.80ms rc=3(EgressRtr) rsc=1 *A:Dut-B#

```



LDP TreeTrace Commands



Testing with DDMAP TLV LDP FEC of Node F when stitched to a BGP Label Route

```
-----
*A:Dut-B# oam lsp-trace prefix 10.20.1.6/32 bgp-label downstream-map-tlv ddmapp detail lsp-
trace to 10.20.1.6/32: 0 hops min, 0 hops max, 124 byte packets
1 10.20.1.3 rtt=3.21ms rc=15(LabelSwitchedWithFecChange) rsc=2
    DS 1: ipaddr=10.10.11.4 ifaddr=10.10.11.4 iftype=ipv4Numbered MRU=1496
        label[1]=131060 protocol=4(RSVP-TE)
        label[2]=131070 protocol=3(LDP)
        label[3]=131060 protocol=2(BGP)
        fecchange[1]=PUSH fectype=RSVP IPv4 prefix=10.20.1.4 remotepeer=10.10.11.4
2 10.20.1.4 rtt=5.50ms rc=3(EgressRtr) rsc=3
2 10.20.1.4 rtt=5.37ms rc=8(DSRtrMatchLabel) rsc=2
    DS 1: ipaddr=10.10.6.5 ifaddr=10.10.6.5 iftype=ipv4Numbered MRU=1500
        label[1]=131071 protocol=3(LDP)
        label[2]=131060 protocol=2(BGP)
3 10.20.1.5 rtt=7.82ms rc=3(EgressRtr) rsc=2
3 10.20.1.5 rtt=6.11ms rc=15(LabelSwitchedWithFecChange) rsc=1
    DS 1: ipaddr=10.10.10.6 ifaddr=10.10.10.6 iftype=ipv4Numbered MRU=1500
        label[1]=131071 protocol=3(LDP)
        fecchange[1]=POP fectype=BGP IPv4 prefix=10.20.1.6 remotepeer=0.0.0.0 (Unknown)
        fecchange[2]=PUSH fectype=LDP IPv4 prefix=10.20.1.6 remotepeer=10.10.10.6
4 10.20.1.6 rtt=10.2ms rc=3(EgressRtr) rsc=1 *A:Dut-B#
```

TWAMP commands

twamp

Syntax twamp

Context config>test-oam

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command enables TWAMP functionality.

Default TWAMP is disabled.

server

Syntax **retry-count** *retry-count*

Context config>test-oam>twamp

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command configures the node for TWAMP server functionality.

Default TWAMP is disabled.

prefix

Syntax **prefix** *{ip-prefix | mask}* [**create**]
no prefix

Context config>test-oam>twamp>server

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command configures an IP address prefix containing one or more TWAMP clients. In order for a TWAMP client to connect to the TWAMP server (and subsequently conduct tests) it must establish the control connection using an IP address that is part of a configured prefix.

Default no prefix

LDP Treetrace Commands

Parameters **prefix** *ip-prefix/mask* — Specifies the address prefix and subnet mask of the destination node.
ip-prefix — An IPv4 address in dotted decimal notation.

Values a.b.c.d

Default none

mask — The prefix length.

Values 0—32

Default none

create — keyword - mandatory when creating an entry.

max-conn-prefix

Syntax **max-conn-prefix** *count*
no max-conn-prefix

Context config>test-oam>twamp>server>prefix

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command configures the maximum number of TWAMP control connections by clients with an IP address in a specific prefix. A new control connection is rejected if accepting it would cause either the prefix limit defined by this command or the server limit (max-conn-server) to be exceeded.

The no form of the command sets the default value.

Default no max-conn-prefix

Parameters *count* — The maximum number of control connections.

Values 7210 SAS-M and 7210 SAS-Sx/S 1/10GE: 16

Default 8

Values 7210 SAS-T, 7210 SAS-Sx 10/100GE , 7210 SAS-R6, and 7210 SAS-R12: 32

Default 16

Values 7210 SAS-X : 48

Default 24

max-conn-server

Syntax **max-conn-server** *count*
no max-conn-server

Context config>test-oam>twamp>server

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command configures the maximum number of TWAMP control connections from all TWAMP clients.

A new control connection is rejected if accepting it would cause either this limit or a prefix limit (max-conn-prefix) to be exceeded.

The no form of the command sets the default value.

Default no max-conn-server

Parameters *count* — The maximum number of control connections.

Values 7210 SAS-M and 7210 SAS-Sx/S 1/10GE : 16

Default 8

Values 7210 SAS-T, 7210 SAS-Sx 10/100GE ,7210 SAS-R6, and 7210 SAS-R12: 32

Default 16

Values 7210 SAS-X : 48

Default 24

inactivity-timeout

Syntax **inactivity-timeout** *seconds*
no inactivity-timeout

Context config>test-oam>twamp>server

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command configures the inactivity timeout for all TWAMP-control connections. If no TWAMP control message is exchanged over the TCP connection for this duration of time the connection is closed and all inprogress tests are terminated.

The no form of the command instructs the system to go with the default value.

Default no inactivity-timeout

LDP Treetrace Commands

Parameters *retry-count* — The duration of the inactivity timeout.

Values 60— 3600

Default 900

max-sess-prefix

Syntax **max-sess-prefix** *count*
no max-sess-prefix

Context config>test-oam>twamp>server>prefix

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command configures the maximum number of concurrent TWAMP-Test sessions by clients with an IP address in a specific prefix. A new test session (described by a Request-TW-Session message) is rejected if accepting it would cause either the limit defined by this command or the server limit (*max-sess-server*) to be exceeded.

The no form of the command instructs the system to go with the default value.

Default no max-sess-prefix

Parameters *count* — The maximum number of concurrent test sessions.

Values 7210 SAS-M and 7210 SAS-Sx/S 1/10GE : 16

Default 8

Values 7210 SAS-T, 7210 SAS-Sx 10/100GE, 7210 SAS-R6, and 7210 SAS-R12: 32

Default 16

Values 7210 SAS-X : 48

Default 24

max-sess-server

Syntax **max-sess-server** *count*
no max-sess-server

Context config>test-oam>twamp>server

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command configures the maximum number of concurrent TWAMP-Test sessions across all allowed clients.

A new test session (described by a Request-TW-Session message) is rejected if accepting it would cause either the limit defined by this command or a prefix limit (max-sess-prefix) to be exceeded.

The no form of the command instructs the system to go with the default value.

Default no max-sessions

Parameters *count* — The maximum number of concurrent test sessions.

Values 7210 SAS-M and 7210 SAS-Sx/S 1/10GE : 16

Default 8

Values 7210 SAS-T, 7210 SAS-R6, and 7210 SAS-R12: 32

Default 16

Values 7210 SAS-X : 48

Default 24

TWAMP Light Commands

twamp-light

- Syntax** `twamp-light`
- Context** `config>router`
`config>service>vprn`
`config>test-oam>twamp`
- Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.
- This command enables the context to configure TWAMP Light functionality.

reflector

- Syntax** `reflector [udp-port udp-port-number] [create]`
`no reflector`
- Context** `config>router>twamp-light`
`config>service>vprn>twamp-light`
- Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.
- This context configures TWAMP Light session reflector-specific parameters. When creating a reflector, the *udp-port-number* value must be configured and the **create** keyword must be included.
- The **no** form of the command removes the reflector.
- Parameters** *udp-port-number* — Specifies the destination UDP port that the session reflector will listen to for TWAMP Light packets. The session controller that is launching the TWAMP Light packets must have the same destination UDP port configured as part of the TWAMP Light test. IES services will use the destination UDP port that is configured under the router context. Only one UDP port may be configured per unique context. An error message will be generated if the specified UDP port is unavailable.
- Values** 64364 to 64373
- create** — Keyword used to create the reflector.

description

- Syntax** **description** *description-string*
no description
- Context** config>router>twamp-light>reflector
 config>router>twamp-light>reflector>prefix
 config>service>vprn>twamp-light>reflector
 config>service>vprn>twamp-light>reflector>prefix
- Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.
- This command creates a text description for the current configuration context that is stored in the configuration file. 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 description.
- Parameters** *description-string* — The description character string. Allowed values are any character strings up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), then the entire string must be enclosed in double quotes.

prefix

- Syntax** **prefix** *ip-prefix/prefix-length* [**create**]
no prefix *ip-prefix/prefix-length*
- Context** config>router>twamp-light>reflector
 config>service>vprn>twamp-light>reflector
- Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.
- This command configures the IP prefixes that the reflector will accept TWAMP Light packets from and respond to. Each prefix requires its own configuration entry.
- The **no** form of the command removes the specifies prefix.
- Parameters** *ip-prefix* — Specifies the IP address.
- Values** IPv4 address in the form a.b.c.d
 IPv6 address in the form x:x:x:x:x:x:x (eight 6-bit pieces) (no multicast addresses)
 x:x:x:x:x:d.d.d.d
 x [0..FFFF]H
 d [0..255]D
- prefix-length* — Specifies the length of the IP prefix.
- Values** IPv4 — 0 to 32
 IPv6 — 0 to 128

TWAMP Light Commands

create — Keyword used to create the IP prefix entry.

shutdown

Syntax [no] shutdown

Context config>router>twamp-light>reflector
config>service>vprn>twamp-light>reflector

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command disables the TWAMP Light reflector functionality within the current context.

The **no** form of the command enables the TWAMP Light reflector functionality within the current context.

Default shutdown

inactivity-timeout

Syntax inactivity-timeout *seconds*
no inactivity-timeout

Context config>test-oam>twamp>twamp-light

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command configures the length of time to maintain stale states on the session reflector. A stale state occurs when test data information has not been refreshed or updated by newly arriving probes for that specific test in a predetermined amount of time. Any single reflector can maintain an up state for a maximum of 12000 tests. If the maximum value is exceeded, the session reflector will not have memory to allocate to new tests.

The **no** form of the command disables the inactivity timer.

Default inactivity-timer 100

Parameters *seconds* — Specifies the number of seconds to maintain a stale state.

Values 10 to 100

session

Syntax **session** *session-name* [**test-family** {**ethernet** | **ip**} [**session-type** {**proactive** | **on-demand**}] **create**]
no session *session-name*

Context config>oam-pm

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command configures the individual session containers that will house the test-specific configuration parameters. Since this session context provides only a container abstract to house the individual test functions, it cannot be shut down. Only individual tests sessions within the container may be shut down. No values, parameters, or configuration within this context may be changed if any individual test is active. Changes may only be made when all tests within the context are shut down, with the exception of the **description**.

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

Parameters *session-name* — Specifies the name of the session container. 32 characters maximum.

ethernet — Specifies that the test will be based on the Ethernet layer.

ip — Specifies that the test will be based on the IP layer.

proactive — Specifies that the test will always be on, with no stop. Tests are proactive by default.

on-demand — Specifies that the test will run on demand, with an immediate start and no stop, or a stop based on offset.

create — Keyword to create the session container.

ip

Syntax **ip**

Context config>oam-pm>session

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command enables the context to configure the IP-specific source and destination information, the priority, and the IP test tools on the launch point.

destination

| | |
|--------------------|---|
| Syntax | destination <i>ip-address</i>
no destination |
| Context | config>oam-pm>session>ip |
| Description | <p>Platforms Supported: 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.</p> <p>This command configures the destination IP address to which the TWAMP Light packets are addressed. The destination address must be included in the prefix list on the session reflector within the context to allow the reflector to process the inbound TWAMP Light packets.</p> <p>The no form of the command removes the destination parameters.</p> |
| Default | no destination |
| Parameters | <p><i>ip-address</i> — Specifies the IP address of the peer to which the packets are directed.</p> <p>Values</p> <ul style="list-style-type: none"> IPv4 address in the form a.b.c.d IPv6 address in the form x:x:x:x:x:x:x (eight 6-bit pieces) (no multicast addresses) x:x:x:x:x:d.d.d.d x [0..FFFF]H d [0..255]D |

dest-udp-port

| | |
|--------------------|---|
| Syntax | dest-udp-port <i>udp-port-number</i>
no dest-udp-port |
| Context | config>oam-pm>session>ip |
| Description | <p>Platforms Supported: 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.</p> <p>This command configures the destination UDP port to which the TWAMP Light packets are sent from the session controller. This value must match the udp-port <i>udp-port number</i> configured on the TWAMP Light reflector that will be responding to this specific TWAMP Light test.</p> <p>The no form of the command removes the destination UDP port configuration.</p> |
| Parameters | <p><i>udp-port-number</i> — Specifies the destination UDP port.</p> <p>Values</p> <ul style="list-style-type: none"> 64364 to 64373 |

fc

Syntax **fc** *fc-name*
no fc

Context config>oam-pm>session>ip

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command configures the forwarding class designation for TWAMP Light packets that are sent through the node and exposed to the various QoS functions on the network element.

The no form of the command restores the default value.

Default fc be

Parameters *fc-name* — Specifies the forwarding class.

Values

- be — best effort
- l2 — low-2
- af — assured
- l1 — low-1
- h2 — high-2
- ef — expedited
- h1 — high-1
- nc — network control

forwarding

Syntax **forwarding bypass-routing**
forwarding interface *interface-name*
forwarding next-hop *ip-address*
no forwarding

Context config>oam-pm>session>ip

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command configures influence for the forwarding decision of the TWAMP Light packet. When this command is used, only one of the forwarding options can be enabled at any time.

The **no** form of the command removes the configured influence and enables the default forwarding logic.

Default no forwarding

Parameters **bypass-routing** — Specifies that packets should be sent to a host on a directly attached network, bypassing the routing table.

interface-name — Specifies the name of the interface from which the packet will be sent. The name must already exist in the **config>router>interface** context or within the appropriate **config>service** context.

TWAMP Light Commands

32 characters maximum.

ip-address — Specifies the IP address of the next hop.

Values IPv4 address in the form a.b.c.d
IPv6 address in the form x:x:x:x:x:x:x:x (eight 6-bit pieces) (no multicast addresses)
x:x:x:x:x:d.d.d.d
x [0..FFFF]H
d [0..255]D

profile

Syntax **profile {in | out}**
no profile

Context config>oam-pm>session>ip

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command configures whether TWAMP Light PDUs are treated as in-profile or out-of-profile.

The **no** form of the command restores the default value. The default has been selected because the forwarding class defaults to best effort.

Default profile out

Parameters **in** — Specifies that TWAMP Light PDU packets will be treated as in-profile.
out — Specifies that TWAMP Light PDU packets will be treated as out-of-profile.

router

Syntax **router router-instance**
router service-name service-name
no router

Context config>oam-pm>session>ip

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command configures the source context from which TWAMP Light packets are launched. The routing instance and service name must be a VPRN instance.

NOTE: VPRN instances may only be specified on 7210 SAS platforms that support VPRN services. Refer to the platform-specific 7210 SAS OS Service Guide for information about platform support for VPRN services.

The **no** form of the command restores the default value.

Default router base

Parameters *router-instance* — Specifies the routing instance from which the TWAMP Light packets are launched.

Values *router-name* | *service-id*
router-name — “base”
service-id — 1 to 2147483647

service-name — Specifies the name of the service from which the TWAMP Light packets are launched. 64 characters maximum.

SOURCE

Syntax **source** *ip-address*
no source

Context config>oam-pm>session>ip

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command configures the source IP address that the session controller (launch point) will use for the test. The source address must be a local resident IP address in the context; otherwise, the response packets will not be processed by the TWAMP Light application. Only source addresses configured as part of TWAMP tests will be able to process the reflected TWAMP packets from the session reflector.

The **no** form of the command removes the source address parameters.

Parameters *ip-address* — Specifies the source IP address.

Values IPv4 address in the form a.b.c.d
 IPv6 address in the form x:x:x:x:x:x:x:x (eight 6-bit pieces) (no multicast addresses)
 x:x:x:x:x:d.d.d.d
 x [0..FFFF]H
 d [0..255]D

source-udp-port

Syntax **source-udp-port** *udp-port-number*
no source-udp-port

Context config>oam-pm>session>ip

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command should only be used if a TWAMP Client is used to establish a TCP connection and communicate the test parameters to a TWAMP Server over TWAMP TCP Control, and the test is launched from OAM-PM (Session-Sender). This command should not be used when the reflection point is a TWAMP Light reflector that does not require TCP TWAMP Control. When this command is configured, the source UDP

TWAMP Light Commands

range is restricted. When this command is omitted, the source UDP port is dynamically allocated by the system.

The **no** form of the command removes the source UDP port configuration and enables default allocation.

Default no source-udp-port

Parameters *udp-port-number* — Specifies the source UDP port.

Values 64374 to 64383

tll

Syntax **tll** *time-to-live*
no tll

Context config>oam-pm>session>ip

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command configures the value of the TTL (time to live) field in the IP header.

The **no** form of the command restores the default value.

Default tll 255

Parameters *time-to-live* — Specifies the numerical value to place in the TTL field.

Values 1 to 255

twamp-light

Syntax **twamp-light** [*test-id test-id*] [**create**]
no twamp-light

Context config>oam-pm>session>ip

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command assigns an identifier to the TWAMP Light test and creates the individual test.

The **no** form of the command removes the TWAMP Light test function from the OAM-PM session.

Default no twamp-light

Parameters *test-id* — Specifies the value of the 4-byte local test identifier that is not sent in TWAMP Light packets.

Values 0 to 2147483647

create — Keyword to create the test.

interval

Syntax **interval** *milliseconds*
no interval

Context config>oam-pm>session>ip>twamp-light

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command configures the message period, or probe spacing, for the transmission of TWAMP Light frames.

The **no** form of the command restores the default value.

Parameters *milliseconds* — Specifies the number of milliseconds between the transmission of TWAMP Light frames.

Values 100, 1000, 10000

pad-size

Syntax **pad-size** *octets*
no pad-size

Context config>oam-pm>session>ip>twamp-light

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command configures the amount by which the TWAMP Light packets are padded. TWAMP session-controller packets are 27 bytes smaller than TWAMP session-reflector packets. If symmetrical packet sizes in the forward and backward direction are required, a minimum padding of 27 bytes must be configured.

The **no** form of the command removes all padding.

Default pad-size 0

Parameters *padding* — Specifies the size of the padding, in octets.

Values 0 to 2000

shutdown

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

Context config>oam-pm>session>ip>twamp-light

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command is used to start or stop the test.

TWAMP Light Commands

Default shutdown

test-duration

Syntax **test-duration** *seconds*
no test-duration

Context config>oam-pm>session>ip>twamp-light

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This optional command configures the length of time that the test will run before stopping automatically. This command is only a valid option when a **session-type** is configured as **on-demand**. This command is not an option when the **session-type** is configured as **proactive**.

The **no** form of the command removes a previously configured **test-duration** value and allows the TWAMP Light test to execute until it is stopped manually.

Default test-duration 0

Parameters *seconds* — Specifies the length of time, in seconds, that the TWAMP Light test will run.

Values 1 to 86400

Show Commands

twamp-light

Syntax twamp-light

Context show>router
show>service

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command is used to display TWAMP Light information.

Output

Sample Output

```
show router twamp-light
-----
TWAMP-Light Reflector
-----
Admin State           : Up                UDP Port           : 15000
Description           : (Not Specified)
Up Time               : 0d 00:02:24
Test Frames Received  : 0                Test Frames Sent   : 0
-----

-----
TWAMP-Light Reflector Prefixes
-----
Prefix                Description
-----
172.16.1.0/24
-----
No. of TWAMP-Light Reflector Prefixes: 1
-----

show service id 500 twamp-light
-----
TWAMP-Light Reflector
-----
Admin State           : Up                UDP Port           : 15000
Description           : TWAMP Light reflector VPRN 500
Up Time               : 0d 01:47:12
Test Frames Received  : 6431             Test Frames Sent   : 6431
-----

-----
```

TWAMP Light Commands

```
TWAMP-Light Reflector Prefixes
-----
Prefix                               Description
-----
10.2.1.1/32                           Process only 10.2.1.1 TWAMP Light
Packets
172.16.1.0/24                          Process all 172.16.1.0 TWAMP
Light packets
-----
No. of TWAMP-Light Reflector Prefixes: 2
-----
```

saa

Syntax `saa [test-name] [owner test-owner]`

Context `show>saa`

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

Use this command to display information about the SAA test.

If no specific test is specified a summary of all configured tests is displayed.

If a specific test is specified then detailed test results for that test are displayed for the last three occurrences that this test has been executed, or since the last time the counters have been reset via a system reboot or clear command.

Parameters *test-name* — Enter the name of the SAA test for which the information needs to be displayed. The test name must already be configured in the `config>saa>test` context.

This is an optional parameter.

owner test-owner — Specifies the owner of an SAA operation up to 32 characters in length.

Values 32 characters maximum.

Default If a *test-owner* value is not specified, tests created by the CLI have a default owner "TiMOS CLI".

Output **SAA Output** — The following table provides SAA field descriptions.

| Label | Description |
|-------------------|--|
| Test Name | Specifies the name of the test. |
| Owner Name | Specifies the owner of the test. |
| Description | Specifies the description for the test type. |
| Accounting policy | Specifies the associated accounting policy ID. |

| Label | Description (Continued) |
|----------------------------|--|
| Administrative status | Specifies whether the administrative status is enabled or disabled. |
| Test type | Specifies the type of test configured. |
| Trap generation | Specifies the trap generation for the SAA test. |
| Test runs since last clear | Specifies the total number of tests performed since the last time the tests were cleared. |
| Number of failed tests run | Specifies the total number of tests that failed. |
| Last test run | Specifies the last time a test was run. |
| Threshold type | Indicates the type of threshold event being tested, jitter-event, latency-event, or loss-event, and the direction of the test responses received for a test run:
in — inbound
out — outbound
rt — roundtrip |
| Direction | Indicates the direction of the event threshold, rising or falling. |
| Threshold Value | Displays the configured threshold value.

Displays the measured crossing value that triggered the threshold crossing event. |
| Last event | Indicates the time that the threshold crossing event occurred. |
| Run # | Indicates what test run produced the specified values. |

```
*A:7210 SAS>show# saa
```

```
=====
SAA Test Information
=====
Test name           : abc
Owner name          : TiMOS CLI
Description          : test
Accounting policy   : None
Administrative status : Disabled
Test type           : Not configured
Trap generation     : None
Test runs since last clear : 0
Number of failed test runs : 0
Last test result    : Undetermined
-----
Threshold
Type      Direction Threshold Value      Last Event      Run #
-----
Jitter-in Rising      None      None      Never      None
```

TWAMP Light Commands

| | | | | | |
|-------------|---------|------|------|-------|------|
| Jitter-out | Falling | None | None | Never | None |
| | Rising | None | None | Never | None |
| Jitter-rt | Falling | None | None | Never | None |
| | Rising | 100 | None | Never | None |
| Latency-in | Falling | 10.0 | None | Never | None |
| | Rising | None | None | Never | None |
| Latency-out | Falling | None | None | Never | None |
| | Rising | None | None | Never | None |
| Latency-rt | Falling | None | None | Never | None |
| | Rising | 100 | None | Never | None |
| Loss-in | Falling | 20.0 | None | Never | None |
| | Rising | None | None | Never | None |
| Loss-out | Falling | None | None | Never | None |
| | Rising | None | None | Never | None |
| Loss-rt | Falling | None | None | Never | None |
| | Rising | 300 | None | Never | None |
| | Falling | 30 | None | Never | None |
| | Rising | | | | |

```

=====
*
=====
*A:7210 SAS>show#

```

test-oam

Syntax test-oam

Context show

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command enables the context to display Operations, Administration, and Maintenance test parameters

Sample Output

```

*A:Dut-A# show saa "Dut-A:1413:1501" owner "TiMOS"
=====
SAA Test Information
=====
Test name           : Dut-A:1413:1501
Owner name          : TiMOS
Administrative status : Enabled
Test type           : vccv-ping 1413:1501 fc "nc" timeout 10 size 200
                    : count 2
Test runs since last clear : 1
Number of failed test runs : 0
Last test result    : Success
-----
Threshold
Type      Direction Threshold Value      Last Event      Run #
-----
Jitter-in Rising      None      None      Never      None

```

| | | | | | |
|-------------|---------|------|------|-------|------|
| Jitter-out | Falling | None | None | Never | None |
| | Rising | None | None | Never | None |
| | Falling | None | None | Never | None |
| Jitter-rt | Rising | None | None | Never | None |
| | Falling | None | None | Never | None |
| Latency-in | Rising | None | None | Never | None |
| | Falling | None | None | Never | None |
| Latency-out | Rising | None | None | Never | None |
| | Falling | None | None | Never | None |
| Latency-rt | Rising | 100 | None | Never | None |
| | Falling | None | None | Never | None |
| Loss-in | Rising | None | None | Never | None |
| | Falling | None | None | Never | None |
| Loss-out | Rising | None | None | Never | None |
| | Falling | None | None | Never | None |
| Loss-rt | Rising | 2 | None | Never | None |
| | Falling | None | None | Never | None |

```
=====
Test Run: 144
```

```
Total number of attempts: 2
```

```
Number of requests that failed to be sent out: 0
```

```
Number of responses that were received: 2
```

```
Number of requests that did not receive any response: 0
```

```
Total number of failures: 0, Percentage: 0
```

| (in ms) | Min | Max | Average | Jitter |
|-------------|-----|-----|---------|--------|
| Outbound : | 0 | 0 | 0 | 0 |
| Inbound : | 10 | 20 | 15 | 0 |
| Roundtrip : | 10 | 20 | 15 | 0 |

```
Per test packet:
```

| Sequence | Outbound | Inbound | RoundTrip | Result |
|----------|----------|---------|-----------|----------------------|
| 1 | 0 | 20 | 20 | EgressRtr(10.20.1.4) |
| 2 | 0 | 10 | 10 | EgressRtr(10.20.1.4) |

```
=====
*A:Dut-A#
```

eth-cfm

Syntax eth-cfm

Context show

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command enables the context to display CFM information.

association

Syntax `association [ma-index] [detail]`

Context `show>eth-cfm`

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command displays eth-cfm association information.

Parameters *ma-index* — Specifies the MA index.

Values 1— 4294967295

detail — Displays detailed information for the eth-cfm association.

Sample Output

```
A:dut-b# show eth-cfm association
```

```
=====
CFM Association Table
=====
Md-index   Ma-index   Name                               CCM-interval  Bridge-id
-----
1           1           a1                                 1              1
1           2           a2                                 1              2
2           1           a1                                 1              2
2           2           a2                                 1              1
=====
A:dut-b#
```

cfm-stack-table

Syntax `cfm-stack-tableup | down[port [port-id [vlan vlan-id]]|sdp sdp-id[:vc-id]] [level 0..7] [direction up | down]`

Context `show>eth-cfm`

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command displays stack-table information. This stack-table is used to display the various management points MEPs and MIPs that are configured on the system. These can be Service based or facility based. The various option allow the operator to be specific. If no parameters are include then the entire stack-table will be displayed.

Parameters *port port-id* — Displays the bridge port or aggregated port on which MEPs or MHFs are configured.

vlan *vlan-id* — Displays the associated VLAN ID.

Values 0 — 4094

level — Display the MD level of the maintenance point.

Values 0 — 7

direction up | down — Displays the direction in which the MP faces on the bridge port.

Sample Output

```
A:dut-b# show eth-cfm cfm-stack-table
```

```
=====
CFM SAP Stack Table
=====
Sap          Level Dir  Md-index  Ma-index  Mep-id  Mac-address
-----
1/1/9:1      6    Down 1      1          1      00:25:ba:01:c3:6a
1/1/9:1      7    Down 2      2          1      00:25:ba:01:c3:6a
1/1/9:2      6    Down 1      2          1      00:25:ba:01:c3:6a
1/1/9:2      7    Down 2      1          1      00:25:ba:01:c3:6a
=====

=====
CFM Ethernet Tunnel Stack Table
=====
Eth-tunnel   Level Dir  Md-index  Ma-index  Mep-id  Mac-address
-----
=====

=====
CFM SDP Stack Table
=====
Sdp          Level Dir  Md-index  Ma-index  Mep-id  Mac-address
-----
No Matching Entries
=====

=====
CFM Virtual Stack Table
=====
Service      Level Dir  Md-index  Ma-index  Mep-id  Mac-address
```

```
-----
No Matching Entries
=====
A:dut-b#
```

domain

Syntax `domain [md-index] [association ma-index | all-associations] [detail]`

Context `show>eth-cfm`

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command displays domain information.

Parameters *md-index* — Displays the index of the MD to which the MP is associated, or 0, if none.

Values [1..4294967295]

association *ma-index* — Displays the index to which the MP is associated, or 0, if none.

all-associations — Displays all associations to the MD.

detail — Displays detailed domain information.

Sample Output

```
A:dut-b# show eth-cfm domain

=====
CFM Domain Table
=====
Md-index   Level Name                               Format
-----
1          6      d1                                         charString
2          7      d2                                         charString
=====
A:dut-b#
```


mep

Syntax **mep** *mep-id* **domain** *md-index* **association** *ma-index* [**loopback**] [**linktrace**]
mep *mep-id* **domain** *md-index* **association** *ma-index* [**remote-mepid** *mep-id* | **all-remote-mepids**]
mep *mep-id* **domain** *md-index* **association** *ma-index* **eth-test-results** [**remote-peer** *mac-address*]
mep *mep-id* **domain** *md-index* **association** *ma-index* **one-way-delay-test** [**remote-peer** *mac-address*]
mep *mep-id* **domain** *md-index* **association** *ma-index* **two-way-delay-test** [**remote-peer** *mac-address*]
mep *mep-id* **domain** *md-index* **association** *ma-index* **two-way-slm-test** [**remote-peer** *mac-address*]

Context show>eth-cfm

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command displays Maintenance Endpoint (MEP) information.

Parameters **domain** *md-index* — Displays the index of the MD to which the MP is associated, or 0, if none.
association *ma-index* — Displays the index to which the MP is associated, or 0, if none.
loopback — Displays loopback information for the specified MEP.
linktrace — Displays linktrace information for the specified MEP.
two-way-slm-test — Includes specified MEP information for two-way-slm-test. **Sample Output**

```
A:dut-b# show eth-cfm mep 1 domain 1 association 1 linktrace
```

```
-----  
Mep Information  
-----
```

| | | | |
|-----------------------|---------------------|----------------|-----------|
| Md-index | : 1 | Direction | : Down |
| Ma-index | : 1 | Admin | : Enabled |
| MepId | : 1 | CCM-Enable | : Enabled |
| IfIndex | : 35946496 | PrimaryVid | : 1 |
| FngState | : fngReset | ControlMep | : False |
| LowestDefectPri | : macRemErrXcon | HighestDefect | : none |
| Defect Flags | : None | | |
| Mac Address | : 00:25:ba:01:c3:6a | CcmLtmPriority | : 7 |
| CcmTx | : 0 | CcmSequenceErr | : 0 |
| Eth-1Dm Threshold | : 3(sec) | | |
| Eth-Ais: | : Disabled | | |
| Eth-Tst: | : Disabled | | |
| CcmLastFailure Frame: | None | | |
| XconCcmFailure Frame: | None | | |

```
-----  
Mep Linktrace Message Information  
-----
```

| | | | |
|-----------------|-----|----------------|-----|
| LtRxUnexplained | : 0 | LtNextSequence | : 2 |
|-----------------|-----|----------------|-----|

TWAMP Light Commands

```
LtStatus          : False          LtResult          : False
TargIsMepId      : False          TargMepId        : 0
TargMac          : 00:00:00:00:00:00  TTL              : 64
EgressId        : 00:00:00:25:ba:01:c3:6a  SequenceNum     : 1
LtFlags         : useFDBonly
```

Mep Linktrace Replies

```
-----
SequenceNum      : 1              ReceiveOrder      : 1
Ttl              : 63            Forwarded         : False
LastEgressId    : 00:00:00:25:ba:01:c3:6a  TerminalMep     : True
NextEgressId    : 00:00:00:25:ba:00:5e:bf  Relay           : rlyHit
ChassisIdSubType : unknown value (0)
ChassisId:
  None
ManAddressDomain:
  None
ManAddress:
  None
IngressMac      : 00:25:ba:00:5e:bf      Ingress Action   : ingOk
IngrPortIdSubType : unknown value (0)
IngressPortId:
  None
EgressMac       : 00:00:00:00:00:00      Egress Action   : egrNoTlv
EgrPortIdSubType : unknown value (0)
EgressPortId:
  None
Org Specific TLV:
  None
A:dut-b#
A:dut-b#
```

```
A:dut-b# show eth-cfm mep 1 domain 1 association 1 loopback
```

Mep Information

```
-----
Mg-index        : 1              Direction        : Down
Ma-index        : 1              Admin            : Enabled
MepId           : 1              CCM-Enable      : Enabled
IfIndex         : 35946496      PrimaryVid      : 1
FngState        : fngReset      ControlMep      : False
LowestDefectPri : macRemErrXcon  HighestDefect   : none
Defect Flags    : None
Mac Address     : 00:25:ba:01:c3:6a  CcmLtmPriority  : 7
CcmTx           : 0              CcmSequenceErr  : 0
Eth-1Dm Threshold : 3(sec)
Eth-Ais:        : Disabled
Eth-Tst:        : Disabled
CcmLastFailure Frame:
  None
XconCcmFailure Frame:
  None
```

Mep Loopback Information

```
-----
LbRxReply       : 1              LbRxBadOrder    : 0
LbRxBadMsdu    : 0              LbTxReply       : 0
```

```

LbSequence      : 2                      LbNextSequence  : 2
LbStatus        : False                  LbResultOk      : True
DestIsMepId     : False                  DestMepId       : 0
DestMac         : 00:00:00:00:00:00     SendCount       : 0
VlanDropEnable  : True                   VlanPriority    : 7
Data TLV:
  None
A:dut-b#
  
```

```

*A:dut-b# show eth-cfm mep 1 domain 4 association 4 two-way-delay-test remote-peer
00:25:ba:00:5e:bf
  
```

```

=====
Eth CFM Two-way Delay Test Result Table
=====
Peer Mac Addr      Delay (us)      Delay Variation (us)
-----
00:25:ba:00:5e:bf  507            507
=====
  
```

```

*A:dut-b#
*A:dut-b# show eth-cfm mep 1 domain 4 association 4 two-way-delay-test
  
```

```

=====
Eth CFM Two-way Delay Test Result Table
=====
Peer Mac Addr      Delay (us)      Delay Variation (us)
-----
00:25:ba:00:5e:bf  507            507
=====
  
```

```

*A:dut-b#
*A:dut-a# show eth-cfm mep 2 domain 4 association 4 eth-test-results remote-peer
00:25:ba:01:c3:6a
  
```

```

=====
Eth CFM ETH-Test Result Table
=====
Peer Mac Addr      FrameCount      Current          Accumulate
                   ByteCount       ErrBits         ErrBits
                   CrcErrs        CrcErrs        CrcErrs
-----
00:25:ba:01:c3:6a  6               0               0
                   384             0               0
=====
  
```

```

*A:dut-a#
*A:dut-a# show eth-cfm mep 2 domain 4 association 4 eth-test-results
  
```

```

=====
Eth CFM ETH-Test Result Table
=====
Peer Mac Addr      FrameCount      Current          Accumulate
                   ByteCount       ErrBits         ErrBits
                   CrcErrs        CrcErrs        CrcErrs
-----
00:25:ba:01:c3:6a  6               0               0
                   384             0               0
=====
  
```

```

*A:dut-a# show eth-cfm mep 2 domain 4 association 4 one-way-delay-test remote-peer
  
```

TWAMP Light Commands

```
00:25:ba:01:c3:6a

=====
Eth CFM One-way Delay Test Result Table
=====
Peer Mac Addr          Delay (us)          Delay Variation (us)
-----
00:25:ba:01:c3:6a     402                 402
=====
*A:dut-a#

*A:dut-a# show eth-cfm mep 2 domain 4 association 4 one-way-delay-test

=====
Eth CFM One-way Delay Test Result Table
=====
Peer Mac Addr          Delay (us)          Delay Variation (us)
-----
00:25:ba:01:c3:6a     402                 402
=====
*A:dut-a#
```

mip

Syntax mip

Context show>eth-cfm

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command displays SAPs/bindings provisioned for allowing the default MIP creation.

```
*A:node-1# show eth-cfm mip
=====
CFM SAP MIP Table
=====
Sap Mip-Enabled Mip Mac Address
-----
1/1/1:1.1 yes Not Configured
=====
CFM SDP MIP Table
=====
Sdp Mip-Enabled Mip Mac Address
-----
No Matching Entries
```

system-config

Syntax system-config**Context** show>eth-cfm**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command shows various system level configuration parameters. These global eth-cfm commands are those which are configured directly under the config>eth-cfm context.

```
# show eth-cfm system-config
=====
CFM System Configuration
=====
Redundancy
MC-LAG Standby MEP Shutdown: true
MC-LAG Hold-Timer : 1 second(s)
Synthetic Loss Measurement
Inactivity Timer : 100 second(s)
=====
```

ldp-treetrace

Syntax ldp-treetrace [*prefix ip-prefix/mask*] [*detail*]**Context** show>test-oam**Description** **Platforms Supported:** 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-R6, and 7210 SAS-R12.

This command displays OAM LDP treetrace information.

Parameters **prefix** *ip-prefix/mask* — Specifies the address prefix and subnet mask of the destination node.
detail — Displays detailed information.

Sample Output

```
*A:ALA-48# show test-oam ldp-treetrace
Admin State           : Up           Discovery State       : Done
Discovery-intvl (min) : 60           Probe-intvl (min)    : 2
Probe-timeout (min)   : 1            Probe-retry          : 3
Trace-timeout (sec)   : 60           Trace-retry          : 3
Max-TTL               : 30           Max-path             : 128
Forwarding-class (fc) : be           Profile              : Out
Total Fecs            : 400         Discovered Fecs     : 400
Last Discovery Start   : 12/19/2006 05:10:14
Last Discovery End     : 12/19/2006 05:12:02
Last Discovery Duration : 00h01m48s
Policy1                : policy-1
Policy2                : policy-2
```

TWAMP Light Commands

```

*A:ALA-48# show test-oam ldp-treetrace detail
Admin State           : Up           Discovery State       : Done
Discovery-intvl (min) : 60           Probe-intvl (min)   : 2
Probe-timeout (min)   : 1            Probe-retry          : 3
Trace-timeout (sec)   : 60           Trace-retry          : 3
Max-TTL               : 30           Max-path             : 128
Forwarding-class (fc) : be           Profile              : Out
Total Fecs            : 400           Discovered Fecs     : 400
Last Discovery Start   : 12/19/2006 05:10:14
Last Discovery End     : 12/19/2006 05:12:02
Last Discovery Duration : 00h01m48s
Policy1                : policy-1
Policy2                : policy-2
=====
Prefix (FEC) Info
=====
Prefix                Path Last           Probe Discov   Discov
                    Num  Discovered        State State     Status
-----
11.11.11.1/32         54  12/19/2006 05:10:15 OK   Done      OK
11.11.11.2/32         54  12/19/2006 05:10:15 OK   Done      OK
11.11.11.3/32         54  12/19/2006 05:10:15 OK   Done      OK
.....
14.14.14.95/32        72  12/19/2006 05:11:13 OK   Done      OK
14.14.14.96/32        72  12/19/2006 05:11:13 OK   Done      OK
14.14.14.97/32        72  12/19/2006 05:11:15 OK   Done      OK
14.14.14.98/32        72  12/19/2006 05:11:15 OK   Done      OK
14.14.14.99/32        72  12/19/2006 05:11:18 OK   Done      OK
14.14.14.100/32       72  12/19/2006 05:11:20 OK   Done      OK
=====
Legend: uP - unexplored paths, tO - trace request timed out
        mH - max hop exceeded, mP - max path exceeded
        nR - no internal resource

*A:ALA-48# show test-oam ldp-treetrace prefix 12.12.12.10/32
Discovery State       : Done           Last Discovered    : 12/19/2006 05:11:02
Discovery Status     : ' OK '
Discovered Paths     : 54
Probe State          : OK
Failed Hops          : 0
Failed Probes        : 0

*A:ALA-48# show test-oam ldp-treetrace prefix 12.12.12.10/32 detail
Discovery State       : Done           Last Discovered    : 12/19/2006 05:11:02
Discovery Status     : ' OK '
Discovered Paths     : 54
Probe State          : OK
Failed Hops          : 0
Failed Probes        : 0
=====
Discovered Paths
=====
PathDest              Egr-NextHop          Remote-RtrAddr       Discovery-time
  DiscoveryTtl         ProbeState            ProbeTmOutCnt         RtnCode
-----
127.1.1.0.5           10.10.1.2            12.12.12.10          12/19/2006 05:11:01
                        7 OK                  0                     EgressRtr
127.1.1.0.9           10.10.1.2            12.12.12.10          12/19/2006 05:11:01

```

```

127.1.0.15      7      OK      0      EgressRtr
                10.10.1.2  12.12.12.10  12/19/2006 05:11:01
127.1.0.19      7      OK      0      EgressRtr
                10.10.1.2  12.12.12.10  12/19/2006 05:11:01
127.1.0.24      7      OK      0      EgressRtr
                10.10.1.2  12.12.12.10  12/19/2006 05:11:01
127.1.0.28      7      OK      0      EgressRtr
                10.10.1.2  12.12.12.10  12/19/2006 05:11:01
.....
127.1.0.252     7      OK      0      EgressRtr
                10.10.1.2  12.12.12.10  12/19/2006 05:11:01
127.1.0.255     7      OK      0      EgressRtr
                10.10.1.2  12.12.12.10  12/19/2006 05:11:01
=====
*A:ALA-48#

```

twamp

Syntax twamp

Context show>test-oam

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command enables the context to request TWAMP information.

twamp-light

Syntax twamp-light

Context show>test-oam>twamp

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command enables the context to request TWAMP Light information.

reflectors

Syntax reflectors

Context show>test-oam>twamp>twamp-light

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

TWAMP Light Commands

This command displays TWAMP Light reflector information.

Output

Sample Output

```
show test-oam twamp twamp-light reflectors
=====
TWAMP-Light Reflectors
=====
Router/VPRN      Admin      UDP Port      Prefixes      Frames Rx      Frames Tx
-----
Base              Up          15000          1              0              0
500               Up          15000          2              6340           6340
-----
No. of TWAMP-Light Reflectors: 2
=====
```

twamp server

Syntax twamp server

Context show>test-oam

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

Used to obtain information about the TWAMP server. It displays summary information for the ip-prefix in use.

Sample Output

```
*A:Dut-G>show>test-oam# twamp server

=====
TWAMP Server
=====
Admin State           : Down                Operational State    : Down
Up Time               : 0d 00:00:00
Current Connections   : 0                  Max Connections     : 8
Connections Rejected : 0                  Inactivity Time Out : 900 seconds
Current Sessions      : 0                  Max Sessions        : 8
Sessions Rejected     : 0                  Sessions Aborted    : 0
Sessions Completed    : 0
Test Packets Rx       : 0                  Test Packets Tx     : 0
=====

TWAMP Server Prefix Summary
```



```

=====
Prefix          Current    Current  Description
                Connections Sessions
-----
No. of TWAMP Server Prefixes: 0
=====
*A:Dut-G>show>test-oam#

```

server all

Syntax server all

Context show>test-oam twamp server

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

Used to display detailed information about the TWAMP server and TWAMP clients using different IP prefix.

Sample Output

```

7210SASM# show test-oam twamp server all

=====
TWAMP Server
=====
Admin State      : Up                Operational State : Up
Up Time         : 0d 08:17:34
Current Connections : 0                Max Connections   : 16
Connections Rejected : 0            Inactivity Time Out : 900 seconds
Current Sessions  : 0                Max Sessions      : 16
Sessions Rejected : 0                Sessions Aborted   : 0
Sessions Completed : 0
Test Packets Rx   : 0                Test Packets Tx    : 0
=====

TWAMP Server Prefix 30.1.1.0/24
=====
Description      : (Not Specified)
Current Connections : 0                Max Connections   : 16
Connections Rejected : 0            Max Sessions      : 16
Current Sessions  : 0                Sessions Aborted   : 0
Sessions Rejected : 0
Sessions Completed : 0
Test Packets Rx   : 0                Test Packets Tx    : 0
=====

Connection information for TWAMP server prefix 30.1.1.0/24
=====
Client          State      Curr Sessions  Sessions Rejected  Sessions Completed

```

TWAMP Light Commands

```

                                     Idle Time (s)   Test Packets Rx   Test Packets Tx
-----
No. of TWAMP Server Connections for Prefix 30.1.1.0/24: 0
=====
TWAMP Server Prefix 60.1.1.0/24
=====
Description           : (Not Specified)
Current Connections   : 0                   Max Connections     : 16
Connections Rejected  : 0
Current Sessions      : 0                   Max Sessions        : 16
Sessions Rejected     : 0                   Sessions Aborted    : 0
Sessions Completed    : 0
Test Packets Rx       : 0                   Test Packets Tx     : 0
=====

Connection information for TWAMP server prefix 60.1.1.0/24
=====
Client                State      Curr Sessions  Sessions Rejected  Sessions Completed
                   Idle Time (s)  Test Packets Rx  Test Packets Tx
-----
No. of TWAMP Server Connections for Prefix 60.1.1.0/24: 0
=====
No. of TWAMP Server Prefixes: 2
=====
```

server prefix

Syntax **server prefix** *ip-prefix/mask*

Context show>test-oam twamp server

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

Display information about the TWAMP clients using the specified prefix.

Sample output

```
*A:7210SAS# show test-oam twamp server prefix 60.1.1.0/24
```

```

=====
TWAMP Server Prefix 60.1.1.0/24
=====
Description           : (Not Specified)
Current Connections   : 0                   Max Connections     : 16
Connections Rejected  : 0
```

```

Current Sessions      : 0                      Max Sessions        : 16
Sessions Rejected    : 0                      Sessions Aborted     : 0
Sessions Completed   : 0
Test Packets Rx      : 0                      Test Packets Tx      : 0
=====
=====
Connection information for TWAMP server prefix 60.1.1.0/24
=====
Client              State      Curr Sessions  Sessions Rejected  Sessions Completed
                   Idle Time (s)  Test Packets Rx  Test Packets Tx
-----
No. of TWAMP Server Connections for Prefix 60.1.1.0/24: 0
=====

```

Parameters *ip-prefix* — The destination address of the static route.

Values [18 chars max]

mask — The prefix length.

testhead-profile

Syntax `testhead-profile profile-id`

Context `show>test-oam`

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

Specifies the testhead profile ID to use with this run/session of testhead invocation. Testhead profile must be configure beforehand using the commands under `config> test-oam> testhead-profile>`.

Output **Testhead-profile Output** — The following table provides testhead-profile field descriptions.

| Label | Description |
|-----------------|---|
| Description | Displays the description configured by the user for the test. |
| Profile Id | Displays the profile identifier. |
| CIR Configured | Displays the value of the CIR configured. |
| PIR Configured | Displays the value of the PIR configured. |
| Frame Size | Displays the size of the frame. |
| CIR Operational | Displays the value of the CIR operational rate configured. |
| PIR Operational | Displays the value of the PIR operational rate configured. |

| Label | Description (Continued) |
|------------------------------|--|
| CIR Rule | Displays the adaptation rule configured by the user. |
| InPrf Dot1p | Displays the dot1p value used to identify green or in-profile packets. |
| Ref. Count | Displays the total number of testhead (completed or running) sessions pointing to a profile or acceptance criteria or a frame payload. |
| OutPrf Dot1p | Displays the dot1p value used to identify green or out-of-profile packets. |
| Duration Hrs, mins, and secs | Displays the test duration in hours, minutes, and seconds. |
| Loss TH | Displays the user configured loss threshold value for comparison with measured value. |
| Jitter TH | Displays the user configured jitter threshold value for comparison with measured value. |
| InProf Loss TH | Displays the user configured in-profile loss threshold value for comparison with measured value. |
| OutProf Loss TH | Displays the user configured out-of-profile loss threshold value for comparison with measured value. |
| Latency TH | Displays the user configured latency threshold value for comparison with measured value. |
| InProf Latency TH | Displays the user configured in-profile latency threshold value for comparison with measured value. |
| OutProf Latency TH | Displays the user configured out-of-profile latency threshold value for comparison with measured value. |
| InProf Jitter TH | Displays the user configured in-profile jitter threshold value for comparison with measured value. |
| OutProf Jitter TH | Displays the user configured out-of-profile jitter threshold value for comparison with measured value. |
| CIR TH | Displays the user configured CIR threshold value for comparison with measured value. |
| PIR TH | Displays the user configured PIR threshold value for comparison with measured value. |
| Payload Type | Identifies the type of the payload. |
| Dst Mac | Displays the value of destination MAC configured by the user to use in the frame generated by the testhead tool |

| Label | Description (Continued) |
|--------------|---|
| Src Mac | Displays the value of source MAC configured by the user to use in the frame generated by the testhead tool |
| Vlan Tag 1 | Displays the values of the outermost vlan-tag configured by the user to use in the frame generated by the testhead tool. |
| Vlan Tag 2 | Displays the values of the second vlan-tag configured by the user to use in the frame generated by the testhead tool. |
| Ethertype | Displays the values of the ethertype configured by the user to use in the frame generated by the testhead tool. |
| TOS | Displays the values of the IP TOS (Type of Service) configured by the user to use in the frame generated by the testhead tool. |
| Src. IP | Displays the values of the source IPv4 address configured by the user to use in the frame generated by the testhead tool. |
| L4 Dst Port | Displays the values of the TCP header configured by the user to use in the frame generated by the testhead tool. |
| Protocol | Displays the values of the IP protocol value configured by the user to use in the frame generated by the testhead tool. |
| Data Pattern | Displays the values of the data pattern configured by the user to use in the frame generated by the testhead tool. |
| DSCP | Displays the values of the DSCP configured by the user to use in the frame generated by the testhead tool. |
| TTL | Displays the values of the IP TTL (Time-to-Live) value configured by the user to use in the frame generated by the testhead tool. |
| Dst. IP | Displays the values of the destination IPv4 address configured by the user to use in the frame generated by the testhead tool. |
| L4 Src Port | Displays the values of the source port configured by the user to use in the frame generated by the testhead tool. |

Sample Output

```
*A:7210SAS>config>test-oam># show test-oam testhead-profile 1
=====
Y.1564 Testhead Profile
=====
Description      : Testhead_Profile_1
Profile Id       : 1
Frame Size       : 512
```

TWAMP Light Commands

```
CIR Configured      : 100                CIR Operational    : 96
PIR Configured      : 200                PIR Operational    : 200
CIR Rule            : max                 Ref. Count         : 0
InPrf Dot1p        : 2                   OutPrf Dot1p       : 4
Duration Hrs        : 0
Duration Mins       : 3
Duration Secs       : 0
```

Acceptance Criteria Id 1

```
Loss TH             : 0.000100           Jitter TH          : 100
InProf Loss TH      : 0.000100           InProf Jitter TH   : 100
OutProf Loss TH     : 0.000100           OutProf Jitter TH  : 100

Latency TH          : 100                Ref. Count         : 0
InProf Latency TH   : 100                CIR TH             : 1000
OutProf Latency TH  : 100                PIR TH            : 200
```

Frame Payload Id 1

```
Payload Type        : tcp-ipv4
Description          : Frame_Payload_1
Dst Mac             : 00:00:00:00:00:02
Src Mac             : 00:00:00:00:00:01
Vlan Tag 1          : Not configured
Vlan Tag 2          : Not configured
Ethertype           : 0x0800              DSCP                : af11
TOS                 : 8                   TTL                 : 64
Src. IP             : 1.1.1.1             Dst. IP             : 2.2.2.2
L4 Dst Port         : 50                  L4 Src Port         : 40
Protocol            : 6                   Ref. Count          : 0
Data Pattern        : a1b2c3d4e5f6
```

=====
*A:7210SAS>config>test-oam>#

testhead

Syntax `testhead test-name owner test-owner`

Context show

Parameters **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

test-name — Name of the SAA test. The test name must already be configured in the `config>saa>test` context.

owner test-owner — Specifies the owner of an SAA operation up to 32 characters in length.

Default If a *test-owner* value is not specified, tests created by the CLI have a default owner “TiMOS CLI”.

Output Testhead Output — The following table provides “testhead test-me owner owner-me” field descriptions.

| Label | Description |
|--------------------|---|
| Owner | Displays the owner of the test. |
| Name | Displays the name of the test. |
| Description | Displays the description for the test type. |
| Profile Id | Displays the associated profile ID. |
| Accept. Crit. Id | Displays the test acceptance criteria ID to be used by the testhead OAM tool to declare the PASS/FAIL result at the completion of the test. |
| Frame Payload Id | Displays frame payload ID, that determines the frame content of the frames generated by the tool. |
| Frame Payload Type | Displays the type of frame payload to be used in frames generated by testhead tool. |
| Color Aware Test | Displays if color aware tests need to be executed. |
| SAP | Displays the SAP ID configured. |
| Completed | Displays if the test has been completed. |
| Stopped | Displays if the test has been stopped. |
| FC | Displays the forwarding class (FC) to use to send the frames generated by the testhead tool. |
| Start Time | Displays the start time of the test. |
| End Time | Displays the end time of the test. |
| Total time taken | Displays the total time taken to execute the test. |
| total pkts in us | Displays the total packets in microseconds. |
| OutPrf pkts in us | Displays the out-of-profile packets in microseconds. |
| InPrf pkts in us | Displays the in-profile packets in microseconds. |
| Total Injected | Displays the running count of total injected packets, including marker packets. |
| Total Received | Displays the running count of total received packets, including marker packets. |
| OutPrf Injected | Displays the running count of total out-of-profile packets, excluding marker packets. |

| Label | Description (Continued) |
|--------------------|---|
| OutPrf Received | Displays the running count of total out-of-profile packets received, including marker packets. |
| InPrf Injected | Displays the running count of total in-profile packets, excluding marker packets. |
| InPrf Received | Displays the running count of total in-profile packets received, including marker packets. |
| Throughput Configd | Displays the CIR Throughput rate Threshold Configured (in Kbps). |
| Throughput Oper | Displays the operational rate used for the configured rate. Operational rate is arrived considering the adaptation rule configured by the user and supported hardware rate. |
| Throughput Measurd | Displays the CIR Throughput Measured Value (in Kbps). |
| PIR Tput Threshld | Displays the PIR Throughput rate Threshold Configured (in Kbps). |
| PIR Tput Meas | Displays the PIR Throughput rate Measured Value (in Kbps). |
| FLR Configured | Displays the Frame Loss Ratio Threshold Configured (in-profile). |
| FLR Measurd | Displays the Frame Loss Ratio Measured (in-profile). |
| FLR Acceptance | Displays Pass, Fail, or Not Applicable. It displays Pass, if the measured value is less than or equal to the configured threshold and displays "Fail" otherwise, and displays "Not Applicable", if the FLR criteria is not used to determine whether the test is in Passed or Failed status. |
| OutPrf FLR Conf | Displays the out-of-profile Frame Loss Ratio configured. |
| OutPrf FLR Meas | Displays the out-of-profile Frame Loss Ratio measured. |
| OutPrf FLR Acep | Displays Pass, Fail, or Not Applicable. It displays Pass, if the measured value is less than or equal to the configured threshold and displays "Fail" otherwise, and displays "Not Applicable", if the out-of-profile FLR criteria is not used to determine whether the test is in Passed or Failed status. |
| InPrf FLR Conf | Displays the in-profile Frame Loss Ratio configured. |
| InPrf FLR Meas | Displays the in-profile Frame Loss Ratio measured. |

| Label | Description (Continued) |
|----------------------|---|
| InPrf FLR Acep | Displays Pass, Fail, or Not Applicable. It displays Pass, if the measured value is less than or equal to the configured threshold and displays "Fail" otherwise, and displays "Not Applicable", if the in-profile FLR criteria is not used to determine whether the test is in Passed or Failed status. |
| Latency Conf (us) | Displays the Latency Threshold configured (in microseconds) |
| Latency Meas (us) | Displays the Average Latency measured (in microseconds) |
| Latency Acceptance | Displays Pass, Fail, or Not Applicable. It displays Pass, if the measured value is less than or equal to the configured threshold and displays "Fail" otherwise, and displays "Not Applicable", if the latency criteria is not used to determine whether the test is in Passed or Failed status. |
| OutPrf Lat Conf (us) | Displays the out-of-profile latency configured. |
| OutPrf Lat Meas (us) | Displays the out-of-profile latency measured. |
| OutPrf Lat Acep | Displays Pass, Fail, or Not Applicable. It displays Pass, if the measured value is less than or equal to the configured threshold and displays "Fail" otherwise, and displays "Not Applicable", if the out-of-profile latency criteria is not used to determine whether the test is in Passed or Failed status. |
| InPrf Lat Conf (us) | Displays the in-profile latency configured. |
| InPrf Lat Meas (us) | Displays the in-profile latency measured. |
| InPrf Lat Acep | Displays Pass, Fail, or Not Applicable. It displays Pass, if the measured value is less than or equal to the configured threshold and displays "Fail" otherwise, and displays "Not Applicable", if the in-profile latency criteria is not used to determine whether the test is in Passed or Failed status. |
| Jitter Conf (us) | Displays the Jitter Threshold Configured (in microseconds). |
| Jitter Meas (us) | Displays the Jitter Measured (in microseconds). |

| Label | Description (Continued) |
|------------------------|--|
| Jitter Accep-
tance | Displays Pass, Fail, or Not Applicable. It displays Pass, if the measured value is less than or equal to the configured threshold and displays "Fail" otherwise, and displays "Not Applicable", if the jitter criteria is not used to determine whether the test is in Passed or Failed status. |
| OutPrf Jit
Conf(us) | Displays the out-of-profile Jitter configured. |
| OutPrf Jit
Meas(us) | Displays the out-of-profile Jitter measured. |
| OutPrf Jit Acep | Displays Pass, Fail, or Not Applicable. It displays Pass, if the measured value is less than or equal to the configured threshold and displays "Fail" otherwise, and displays "Not Applicable", if the out-of-profile jitter criteria is not used to determine whether the test is in Passed or Failed status. |
| InPrf Jit
Conf(us) | Displays the in-profile Jitter configured. |
| InPrf Jit
Meas(us) | Displays the in-profile Jitter measured. |
| InPrf Jit Acep | Displays Pass, Fail, or Not Applicable. It displays Pass, if the measured value is less than or equal to the configured threshold and displays "Fail" otherwise, and displays "Not Applicable", if the in-profile jitter criteria is not used to determine whether the test is in Passed or Failed status. |
| Total Pkts. Tx. | Total number of packets (that is, data and marker) transmitted by the testhead session for the duration of the test. |
| OutPrf Latency
Pkt* | Total number of out-of-profile marker packets received by the testhead session for the duration of the test. |
| Total Tx. Fail | Total number of failed transmission attempts by the testhead session for the duration of the test. |
| Latency Pkts. Tx | Total number of marker packets transmitted by the testhead session for the duration of the test. |
| InPrf Latency
Pkt* | Total number of in-profile marker packets received by the testhead session for the duration of the test. |

Sample output

```
*A:7210SAS# show testhead test-me owner owner-me
```

=====
Y.1564 Testhead Session
=====

```

Owner          : owner-me
Test           : test-me
Profile Id     : 1                      SAP           : 1/1/2:100
Accept. Crit. Id : 0                    Completed      : Yes
Frame Payload Id : 1                    Stopped       : No
Frame Payload Type : tcp-ipv4           FC            : be
Color Aware Test : Yes
Start Time    : 08/08/2001 19:37:11
End Time      : 08/08/2001 19:40:16
Total time taken : 0d 00:03:05

```

Latency Results

| (total pkts in us): | Min | Max | Average | Jitter |
|----------------------|-----|-----|---------|--------|
| Roundtrip : | 0 | 0 | 0 | 0 |
| (OutPrf pkts in us): | Min | Max | Average | Jitter |
| Roundtrip : | 0 | 0 | 0 | 0 |
| (InPrf pkts in us): | Min | Max | Average | Jitter |
| Roundtrip : | 0 | 0 | 0 | 0 |

Packet Count

```

Total Injected   : 42273637
Total Received   : 0

OutPrf Injected  : 16898179
OutPrf Received  : 0

InPrf Injected   : 25375450
InPrf Received   : 0

```

Test Compliance Report

```

Throughput Configd : 962388
Throughput Oper    : 962384
Throughput Measurd : 0

PIR Tput Threshld : Not configured
PIR Tput Meas     : 0

CIR Tput Threshld : Not configured
CIR Tput Meas     : 0

FLR Configured    : None
FLR Measurd       : Not Applicable
FLR Acceptance    : Fail

OutPrf FLR Conf   : None
OutPrf FLR Meas   : Not Applicable

```

TWAMP Light Commands

```
OutPrf FLR Acep      : Not Applicable

InPrf FLR Conf      : None
InPrf FLR Meas      : Not Applicable
InPrf FLR Acep      : Not Applicable

Latency Configd(us) : None
Latency Measurd(us) : None
Latency Acceptance  : Not Applicable

OutPrf Lat Conf(us) : None
OutPrf Lat Meas(us) : None
OutPrf Lat Acep     : Not Applicable

InPrf Lat Conf(us)  : None
InPrf Lat Meas(us)  : None
InPrf Lat Acep     : Not Applicable

Jitter Configd(us) : None
Jitter Measurd(us) : None
Jitter Acceptance  : Not Applicable

OutPrf Jit Conf(us) : None
OutPrf Jit Meas(us) : None
OutPrf Jit Acep     : Not Applicable

InPrf Jit Conf(us)  : None
InPrf Jit Meas(us)  : None
InPrf Jit Acep     : Not Applicable

Total Pkts. Tx.      : 13                      Latency Pkts. Tx. : 8
OutPrf Latency Pkt* : 0                        InPrf Latency Pkt* : 0
Total Tx. Fail       : 0
=====
*A:7210SAS# show testhead test-me  owner owner-me
```

bin-group

Syntax `bin-group [bin-group-number]`

Context `show>oam-pm`

Description This command displays data for one or all OAM-PM bin groups.

Parameters *bin-group-number* — Specifies an OAM-PM bin group.

Values 1 to 255

Output

Sample Output

```
show oam-pm bin-group
```

```

-----
Configured Lower Bounds for Delay Measurement (DMM) Tests, in microseconds
-----
Group Description                Admin Bin   FD(us)   FDR(us)  IFDV(us)
-----
1   OAM PM default bin group (not*  Up   0         0         0         0
                                  1       5000      5000      5000
                                  2      10000      -         -
-----
2                               Up   0         0         0         0
                                  1       1000      5000      100
                                  2       2000      -         200
                                  3       3000      -         300
                                  4       4000      -         400
                                  5       5000      -         500
                                  6       6000      -         600
                                  7       7000      -         700
                                  8       8000      -         800
                                  9      10000      -        1000
-----
3                               Down  0         0         0         0
                                  1       6000      5000      8000
                                  2      10000     10000     10000
                                  3      15000     15000      -
                                  4      22000      -         -
-----
10  base                            Up   0         0         0         0
                                  1       5000      5000      5000
                                  2      10000     10000     10000
-----

```

* indicates that the corresponding row element may have been truncated.

show oam-pm bin-group 2

```

-----
Configured Lower Bounds for Delay Measurement (DMM) Tests, in microseconds
-----
Group Description                Admin Bin   FD(us)   FDR(us)  IFDV(us)
-----
2                               Up   0         0         0         0
                                  1       1000      5000      100
                                  2       2000      -         200
                                  3       3000      -         300
                                  4       4000      -         400
                                  5       5000      -         500
                                  6       6000      -         600
                                  7       7000      -         700
                                  8       8000      -         800
                                  9      10000      -        1000
-----

```

bin-group-using

Syntax `bin-group-using [bin-group bin-group-number]`

Context `show>oam-pm`

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command displays the list of sessions configured against one or all OAM-PM bin groups.

Parameters *bin-group-number* — Specifies an OAM-PM bin group.

Values 1 to 255

Output

Sample Output

```
show oam-pm bin-group-using
=====
OAM Performance Monitoring Bin Group Configuration for Sessions
=====
Bin Group      Admin  Session                               Session State
-----
2              Up     eth-vpls-00005                         Inact
              eth-pm-service-4                       Act
-----
3              Down   eth-epipe-000001                       Inact
-----
10             Up     eth-epipe-00002                         Inact
-----
=====
Admin: State of the bin group
Session State: The state of session referencing the bin-group

show oam-pm bin-group-using bin-group 2
=====
OAM Performance Monitoring Bin Group Configuration for Sessions
=====
Bin Group      Admin  Session                               Session State
-----
2              Up     eth-vpls-00005                         Inact
              eth-pm-service-4                       Act
-----
=====
Admin: State of the bin group
Session State: The state of session referencing the bin-group
```

session

Syntax `session session-name` [{**all** | **base** | **bin-group** | **event-mon** | **meas-interval**}]

Context `show>oam-pm`

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command displays the configuration and status information for an OAM-PM session.

Parameters *session-name* — Specifies the name of the session. 32 characters maximum.

all — Keyword to display all attributes.

base — Keyword to display the base configuration option for the session.

bin-group — Keyword to display the associated bin group and its attributes.

event-mon — Keyword to display configured event monitoring and last TCA information.

meas-interval — Keyword to display the associated measured interval and its attributes.

Output

Sample Output

```
show oam-pm session "eth-pm-service-4" all
-----
Basic Session Configuration
-----
Session Name       : eth-pm-service-4
Description        : (Not Specified)
Test Family        : ethernet           Session Type       : proactive
Bin Group          : 2
-----

Ethernet Configuration
-----
Source MEP         : 28                 Priority            : 0
Source Domain      : 12                 Dest MAC Address   : 00:00:00:00:00:30
Source Assoc'n     : 4
-----

DMM Test Configuration and Status
-----
Test ID           : 10004                Admin State        : Up
Oper State        : Up                   Data TLV Size      : 1000 octets
On-Demand Duration: Not Applicable       On-Demand Remaining: Not Applicable
Interval          : 1000 ms
-----

SLM Test Configuration and Status
```

TWAMP Light Commands

```
-----  
Test ID           : 10004           Admin State       : Up  
Oper State        : Up              Data TLV Size    : 1000 octets  
On-Demand Duration: Not Applicable  On-Demand Remaining: Not Applicable  
Interval          : 100 ms  
CHLI Threshold   : 4 HLIs          Frames Per Delta-T : 10 SLM frames  
Consec Delta-Ts  : 10              FLR Threshold    : 50%  
-----
```

15-mins Measurement Interval Configuration

```
-----  
Duration          : 15-mins         Intervals Stored  : 32  
Boundary Type     : clock-aligned    Clock Offset      : 0 seconds  
Accounting Policy : none  
-----
```

Configured Lower Bounds for Delay Measurement (DMM) Tests, in microseconds

```
-----  
Group Description      Admin Bin   FD(us)   FDR(us)  IFDV(us)  
-----  
2                     Up         0         0         0         0  
                     1         1000     5000     -         100  
                     2         2000     -         -         200  
                     3         3000     -         -         300  
                     4         4000     -         -         400  
                     5         5000     -         -         500  
                     6         6000     -         -         600  
                     7         7000     -         -         700  
                     8         8000     -         -         800  
                     9         10000    -         -         1000  
-----
```

```
?  
show oam-pm session "eth-pm-service-4" base  
-----
```

Basic Session Configuration

```
-----  
Session Name       : eth-pm-service-4  
Description        : (Not Specified)  
Test Family        : ethernet       Session Type      : proactive  
Bin Group          : 2  
-----
```

Ethernet Configuration

```
-----  
Source MEP         : 28              Priority          : 0  
Source Domain      : 12              Dest MAC Address : 00:00:00:00:00:30  
Source Assoc'n     : 4  
-----
```

DMM Test Configuration and Status

```
-----  
Test ID           : 10004           Admin State       : Up  
Oper State        : Up              Data TLV Size    : 1000 octets  
-----
```



```
On-Demand Duration: Not Applicable      On-Demand Remaining: Not Applicable
Interval          : 1000 ms
```

```
-----
SLM Test Configuration and Status
-----
```

```
Test ID           : 10004                Admin State       : Up
Oper State        : Up                   Data TLV Size    : 1000 octets
On-Demand Duration: Not Applicable      On-Demand Remaining: Not Applicable
Interval          : 100 ms
CHLI Threshold   : 4 HLIs                Frames Per Delta-T : 10 SLM frames
Consec Delta-Ts  : 10                     FLR Threshold    : 50%
```

```
show oam-pm session "eth-pm-service-4" bin-group
```

```
-----
Configured Lower Bounds for Delay Measurement (DMM) Tests, in microseconds
-----
```

| Group Description | Admin Bin | FD(us) | FDR(us) | IFDV(us) |
|-------------------|-----------|--------|---------|----------|
| 2 | Up | 0 | 0 | 0 |
| | 1 | 1000 | 5000 | 100 |
| | 2 | 2000 | - | 200 |
| | 3 | 3000 | - | 300 |
| | 4 | 4000 | - | 400 |
| | 5 | 5000 | - | 500 |
| | 6 | 6000 | - | 600 |
| | 7 | 7000 | - | 700 |
| | 8 | 8000 | - | 800 |
| | 9 | 10000 | - | 1000 |

```
show oam-pm session "eth-pm-service-4" meas-interval
```

```
-----
15-mins Measurement Interval Configuration
-----
```

```
Duration          : 15-mins                Intervals Stored  : 32
Boundary Type     : clock-aligned          Clock Offset      : 0 seconds
Accounting Policy : none
```

sessions

Syntax `sessions [test-family {ethernet | ip}] [event-mon]`

Context `show>oam-pm`

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command displays a summary of the OAM-PM sessions.

TWAMP Light Commands

- Parameters**
- ethernet** — Keyword to display Ethernet sessions only.
 - ip** — Keyword to display IP sessions only.
 - event-mon** — Keyword to display a summary of all event monitoring information, and the current state for each session.

Output

Sample Output

```
show oam-pm sessions
=====
OAM Performance Monitoring Session Summary for the Ethernet Test Family
=====
Session                State   Bin Group   Sess Type   Test Types
-----
eth-vpls-00005         Inact      2   proactive   DMM SLM
eth-epipe-00002        Inact     10   proactive   DMM SLM
eth-epipe-000001       Inact      3   proactive   DMM
eth-pm-service-4       Act        2   proactive   DMM SLM
=====
```

```
show oam-pm sessions test-family ethernet
=====
OAM Performance Monitoring Session Summary for the Ethernet Test Family
=====
Session                State   Bin Group   Sess Type   Test Types
-----
eth-vpls-00005         Inact      2   proactive   DMM SLM
eth-epipe-00002        Inact     10   proactive   DMM SLM
eth-epipe-000001       Inact      3   proactive   DMM
eth-pm-service-4       Act        2   proactive   DMM SLM
=====
```

statistics

Syntax **statistics**

Context show>oam-pm

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command enables the context to display OAM-PM delay or synthetic loss statistics.

session

Syntax `session session-name`

Context `show>oam-pm>statistics`

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command displays OAM-PM session statistics.

Parameters *session-name* — Specifies the session name. 32 characters maximum.

dmm

Syntax `dmm`

Context `show>oam-pm>statistics>session`

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command enables the context to display DMM test statistics.

meas-interval

Syntax `meas-interval raw [{all | bins | summary}]
meas-interval {5-mins | 15-mins | 1-hour | 1-day} interval-number interval-number [{all | bins | summary}]`

Context `show>oam-pm>statistics>session>dmm`

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command displays measured interval statistics for DMM tests in the specified session.

Parameters **raw** — Keyword to display raw information.

5-mins — Keyword to display information for 5-min intervals.

15-mins — Keyword to display information for 15-min intervals.

1-hour — Keyword to display information for 1-hour intervals.

1-day — Keyword to display information for 1-day intervals.

interval-number — Specifies the interval number.

Values 1 to 97

all — Keyword to display all information for the interval.

bins — Keyword to display bin information for the interval.

summary — Keyword to display summarized information for the interval.

Output

Sample Output

```
show oam-pm statistics session "eth-pm-service-4" dmm meas-interval 15-mins all interval-
number 2
```

```
-----
Start (UTC)           : 2014/02/01 10:15:00           Status           : completed
Elapsed (seconds)    : 900                          Suspect          : no
Frames Sent          : 900                          Frames Received  : 900
-----
```

```
-----
```

| Bin Type | Direction | Minimum (us) | Maximum (us) | Average (us) |
|----------|------------|--------------|--------------|--------------|
| FD | Forward | 0 | 11670 | 779 |
| FD | Backward | 0 | 7076 | 1746 |
| FD | Round Trip | 1109 | 13222 | 2293 |
| FDR | Forward | 0 | 11670 | 779 |
| FDR | Backward | 0 | 7076 | 1738 |
| FDR | Round Trip | 0 | 12104 | 1178 |
| IFDV | Forward | 0 | 10027 | 489 |
| IFDV | Backward | 0 | 5444 | 742 |
| IFDV | Round Trip | 0 | 11853 | 1088 |

```
-----
```

```
-----
Frame Delay (FD) Bin Counts
```

```
-----
```

| Bin | Lower Bound | Forward | Backward | Round Trip |
|-----|-------------|---------|----------|------------|
| 0 | 0 us | 625 | 244 | 0 |
| 1 | 1000 us | 194 | 356 | 465 |
| 2 | 2000 us | 50 | 153 | 244 |
| 3 | 3000 us | 11 | 121 | 119 |
| 4 | 4000 us | 10 | 17 | 40 |
| 5 | 5000 us | 5 | 6 | 20 |
| 6 | 6000 us | 4 | 2 | 5 |
| 7 | 7000 us | 0 | 1 | 3 |
| 8 | 8000 us | 0 | 0 | 3 |
| 9 | 10000 us | 1 | 0 | 1 |

```
-----
```

```
-----
Frame Delay Range (FDR) Bin Counts
```

```
-----
```

| Bin | Lower Bound | Forward | Backward | Round Trip |
|-----|-------------|---------|----------|------------|
| 0 | 0 us | 890 | 891 | 889 |
| 1 | 5000 us | 10 | 9 | 11 |

```
-----
```

Inter-Frame Delay Variation (IFDV) Bin Counts

| Bin | Lower Bound | Forward | Backward | Round Trip |
|-----|-------------|---------|----------|------------|
| 0 | 0 us | 398 | 255 | 102 |
| 1 | 100 us | 82 | 88 | 89 |
| 2 | 200 us | 79 | 57 | 59 |
| 3 | 300 us | 60 | 63 | 61 |
| 4 | 400 us | 39 | 37 | 54 |
| 5 | 500 us | 31 | 24 | 42 |
| 6 | 600 us | 26 | 30 | 43 |
| 7 | 700 us | 29 | 20 | 34 |
| 8 | 800 us | 54 | 47 | 67 |
| 9 | 1000 us | 102 | 279 | 349 |

?

```
show oam-pm statistics session "eth-pm-service-4" dmm meas-interval 15-mins bins interval-
number 2
```

```
-----
Start (UTC)           : 2014/02/01 10:30:00           Status           : completed
Elapsed (seconds)    : 900                          Suspect          : no
Frames Sent          : 900                          Frames Received  : 900
-----
```

Frame Delay (FD) Bin Counts

| Bin | Lower Bound | Forward | Backward | Round Trip |
|-----|-------------|---------|----------|------------|
| 0 | 0 us | 699 | 167 | 0 |
| 1 | 1000 us | 169 | 312 | 456 |
| 2 | 2000 us | 24 | 228 | 274 |
| 3 | 3000 us | 3 | 136 | 111 |
| 4 | 4000 us | 3 | 48 | 41 |
| 5 | 5000 us | 1 | 7 | 10 |
| 6 | 6000 us | 1 | 1 | 3 |
| 7 | 7000 us | 0 | 1 | 2 |
| 8 | 8000 us | 0 | 0 | 3 |
| 9 | 10000 us | 0 | 0 | 0 |

Frame Delay Range (FDR) Bin Counts

| Bin | Lower Bound | Forward | Backward | Round Trip |
|-----|-------------|---------|----------|------------|
| 0 | 0 us | 898 | 891 | 892 |
| 1 | 5000 us | 2 | 9 | 8 |

Inter-Frame Delay Variation (IFDV) Bin Counts

| Bin | Lower Bound | Forward | Backward | Round Trip |
|-----|-------------|---------|----------|------------|
| 0 | 0 us | 462 | 217 | 107 |

TWAMP Light Commands

```

1          100 us          63          99          80
2          200 us          64          85          71
3          300 us          63          74          53
4          400 us          34          53          45
5          500 us          37          24          50
6          600 us          34          17          41
7          700 us          35          23          57
8          800 us          46          32          60
9         1000 us          62          276         336
-----

```

?

```
show oam-pm statistics session "eth-pm-service-4" dmm meas-interval 15-mins summary inter-
val-number 2
-----
```

```

Start (UTC)      : 2014/02/01 10:30:00      Status           : completed
Elapsed (seconds) : 900                    Suspect          : no
Frames Sent      : 900                    Frames Received  : 900
-----

```

```

-----
Bin Type   Direction   Minimum (us)   Maximum (us)   Average (us)
-----
FD         Forward     0              6379           518
FD         Backward    0              7856           2049
FD         Round Trip  1118           9879           2241
FDR        Forward     0              6379           518
FDR        Backward    0              7856           2049
FDR        Round Trip  9              8770           1132
IFDV      Forward     0              6021           328
IFDV      Backward    0              5800           732
IFDV      Round Trip  2              7758           984
-----

```

?

```
show oam-pm statistics session "eth-pm-service-4" dmm meas-interval raw
-----
```

```

Start (UTC)      : 2014/02/01 09:43:58      Status           : in-progress
Elapsed (seconds) : 3812                    Suspect          : yes
Frames Sent      : 3812                    Frames Received  : 3812
-----

```

```

-----
Bin Type   Direction   Minimum (us)   Maximum (us)   Average (us)
-----
FD         Forward     0              11670          629
FD         Backward    0              11710          2156
FD         Round Trip  1109           14902          2497
FDR        Forward     0              11670          617
FDR        Backward    0              11710          2156
FDR        Round Trip  0              13784          1360
IFDV      Forward     0              10027          404
IFDV      Backward    0              10436          768
IFDV      Round Trip  0              13542          1056
-----

```

```
-----
Frame Delay (FD) Bin Counts
-----
```

| Bin | Lower Bound | Forward | Backward | Round Trip |
|-----|-------------|---------|----------|------------|
| 0 | 0 us | 2815 | 661 | 0 |
| 1 | 1000 us | 803 | 1287 | 1591 |
| 2 | 2000 us | 127 | 971 | 1227 |
| 3 | 3000 us | 21 | 639 | 623 |
| 4 | 4000 us | 25 | 181 | 232 |
| 5 | 5000 us | 12 | 42 | 72 |
| 6 | 6000 us | 7 | 14 | 28 |
| 7 | 7000 us | 0 | 4 | 13 |
| 8 | 8000 us | 1 | 12 | 19 |
| 9 | 10000 us | 1 | 1 | 7 |

Frame Delay Range (FDR) Bin Counts

| Bin | Lower Bound | Forward | Backward | Round Trip |
|-----|-------------|---------|----------|------------|
| 0 | 0 us | 3792 | 3740 | 3751 |
| 1 | 5000 us | 21 | 73 | 62 |

Inter-Frame Delay Variation (IFDV) Bin Counts

| Bin | Lower Bound | Forward | Backward | Round Trip |
|-----|-------------|---------|----------|------------|
| 0 | 0 us | 1815 | 884 | 410 |
| 1 | 100 us | 338 | 439 | 354 |
| 2 | 200 us | 280 | 313 | 282 |
| 3 | 300 us | 241 | 313 | 268 |
| 4 | 400 us | 162 | 193 | 231 |
| 5 | 500 us | 134 | 141 | 202 |
| 6 | 600 us | 126 | 102 | 178 |
| 7 | 700 us | 127 | 97 | 153 |
| 8 | 800 us | 208 | 165 | 276 |
| 9 | 1000 us | 381 | 1165 | 1458 |

slm

Syntax `slm`

Context `show>oam-pm>statistics>session`

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command enables the context to display SLM test statistics.

meas-interval

Syntax **meas-interval raw**
meas-interval {5-mins | 15-mins | 1-hour | 1-day} interval-number interval-number

Context show>oam-pm>statistics>session>slm

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command displays measured interval statistics for SLM tests in the specified session

Parameters **raw** — Keyword to display raw information.

5-mins — Keyword to display information for 5-min intervals.

15-mins — Keyword to display information for 15-min intervals.

1-hour — Keyword to display information for 1-hour intervals.

1-day — Keyword to display information for 1-day intervals.

interval-number — Specifies the interval number.

Values 1 to 97

Output

Sample Output

```
show oam-pm statistics session "eth-pm-service-4" slm meas-interval 15-mins interval-number
2
-----
Start (UTC)           : 2014/02/01 10:30:00           Status           : completed
Elapsed (seconds)    : 900                               Suspect          : no
Frames Sent          : 9000                               Frames Received  : 9000
-----

-----
                Frames Sent           Frames Received
-----
Forward                9000                9000
Backward               9000                9000
-----

-----
Frame Loss Ratios
-----
                Minimum           Maximum           Average
-----
Forward          0.000%           0.000%           0.000%
Backward         0.000%           0.000%           0.000%
-----

-----
Availability Counters (Und = Undetermined)
```



```
-----
                Available  Und-Avail Unavailable Und-Unavail      HLI      CHLI
-----
Forward          900          0          0          0          0          0
Backward         900          0          0          0          0          0
-----
```

```
show oam-pm statistics session "eth-pm-service-4" slm meas-interval raw
```

```
-----
Start (UTC)      : 2014/02/01 09:44:03      Status      : in-progress
Elapsed (seconds) : 4152                Suspect     : yes
Frames Sent     : 41523                Frames Received : 41523
-----
```

```
-----
                Frames Sent      Frames Received
-----
Forward          41369                41369
Backward         41369                41369
-----
```

```
-----
Frame Loss Ratios
-----
```

```
-----
                Minimum      Maximum      Average
-----
Forward          0.000%      0.000%      0.000%
Backward         0.000%      0.000%      0.000%
-----
```

```
-----
Availability Counters (Und = Undetermined)
-----
```

```
-----
                Available  Und-Avail Unavailable Und-Unavail      HLI      CHLI
-----
Forward          4137          0          0          0          0          0
Backward         4137          0          0          0          0          0
-----
```

```
Ib
```

twamp-light

Syntax twamp-light

Context show>oam-pm>statistics>session

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command enables the context to display TWAMP Light test statistics.

meas-interval

Syntax **meas-interval raw delay** [{all | bins | summary}]
meas-interval raw [loss]
meas-interval {5-mins | 15-mins | 1-hour | 1-day} **interval-number** *interval-number* **delay** [{all | bins | summary}]
meas-interval {5-mins | 15-mins | 1-hour | 1-day} **interval-number** *interval-number* [loss]

Context show>oam-pm>statistics>session>twamp-light

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command displays measured interval statistics for TWAMP-Light tests in the specified session

Parameters **raw** — Keyword to display raw information.
5-mins — Keyword to display information for 5-min intervals.
15-mins — Keyword to display information for 15-min intervals.
1-hour — Keyword to display information for 1-hour intervals.
1-day — Keyword to display information for 1-day intervals.
interval-number — Specifies the interval number.
 Values 1 to 97
delay — Keyword to display TWAMP Light delay statistics only.
loss — Keyword to display TWAMP Light loss statistics only.
all — Keyword to display all information for the interval.
bins — Keyword to display bin information for the interval.
summary — Keyword to display summarized information for the interval.

Monitor Commands

session

Syntax `session session-name [{dmm | slm | twamp-light}]`

Context `monitor>oam-pm`

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command monitors the raw measurement interval for the specified session.

Parameters *session-name* — Specifies the session name. 32 characters maximum.

dmm — Keyword to display monitoring information for DMM tests only.

slm — Keyword to display monitoring information for SLM tests only.

twamp-light — Keyword to display monitoring information for TWAMP Light tests only.

Output

Sample Output

```
monitor oam-pm session "eth-pm-service-4" dmm
```

```
-----  
At time t = 0 sec (Base Statistics)  
-----
```

```
-----  
Frame Delay (FD) Bin Counts  
-----
```

| Bin | Lower Bound | Forward | Backward | Round Trip |
|-----|-------------|---------|----------|------------|
| 0 | 0 us | 3928 | 1125 | 0 |
| 1 | 1000 us | 1197 | 1855 | 2611 |
| 2 | 2000 us | 183 | 1361 | 1565 |
| 3 | 3000 us | 36 | 762 | 778 |
| 4 | 4000 us | 30 | 214 | 280 |
| 5 | 5000 us | 14 | 45 | 81 |
| 6 | 6000 us | 8 | 17 | 35 |
| 7 | 7000 us | 1 | 5 | 16 |
| 8 | 8000 us | 5 | 15 | 26 |
| 9 | 10000 us | 1 | 4 | 11 |

```
-----  
Frame Delay Range (FDR) Bin Counts  
-----
```

| Bin | Lower Bound | Forward | Backward | Round Trip |
|-----|-------------|---------|----------|------------|
|-----|-------------|---------|----------|------------|

Monitor Commands

```

0          0 us          5374          5317          5321
1          5000 us         29           86           82
-----

```

Inter-Frame Delay Variation (IFDV) Bin Counts

```

-----
Bin      Lower Bound      Forward      Backward      Round Trip
-----
0          0 us          2475          1268          625
1          100 us          516           676          554
2          200 us          395           479          417
3          300 us          338           451          398
4          400 us          224           291          340
5          500 us          185           212          280
6          600 us          187           137          234
7          700 us          185           134          208
8          800 us          315           223          392
9         1000 us          582          1531          1954
-----

```

```

-----
At time t = 10 sec (Mode: Delta)
-----

```

Frame Delay (FD) Bin Counts

```

-----
Bin      Lower Bound      Forward      Backward      Round Trip
-----
0          0 us           0            7            0
1         1000 us          10           2            6
2         2000 us           0            1            3
3         3000 us           0            0            1
4         4000 us           0            0            0
5         5000 us           0            0            0
6         6000 us           0            0            0
7         7000 us           0            0            0
8         8000 us           0            0            0
9        10000 us           0            0            0
-----

```

Frame Delay Range (FDR) Bin Counts

```

-----
Bin      Lower Bound      Forward      Backward      Round Trip
-----
0          0 us           10           10           10
1         5000 us           0            0            0
-----

```

Inter-Frame Delay Variation (IFDV) Bin Counts

```

-----
Bin      Lower Bound      Forward      Backward      Round Trip
-----
0          0 us           5            4            2

```

| | | | | |
|---|---------|---|---|---|
| 1 | 100 us | 2 | 2 | 2 |
| 2 | 200 us | 2 | 1 | 1 |
| 3 | 300 us | 1 | 0 | 0 |
| 4 | 400 us | 0 | 0 | 1 |
| 5 | 500 us | 0 | 0 | 0 |
| 6 | 600 us | 0 | 0 | 0 |
| 7 | 700 us | 0 | 0 | 1 |
| 8 | 800 us | 0 | 0 | 0 |
| 9 | 1000 us | 0 | 3 | 3 |

 At time t = 20 sec (Mode: Delta)

 Frame Delay (FD) Bin Counts

| Bin | Lower Bound | Forward | Backward | Round Trip |
|-----|-------------|---------|----------|------------|
| 0 | 0 us | 9 | 0 | 0 |
| 1 | 1000 us | 0 | 7 | 6 |
| 2 | 2000 us | 0 | 3 | 3 |
| 3 | 3000 us | 1 | 0 | 0 |
| 4 | 4000 us | 0 | 0 | 0 |
| 5 | 5000 us | 0 | 0 | 1 |
| 6 | 6000 us | 0 | 0 | 0 |
| 7 | 7000 us | 0 | 0 | 0 |
| 8 | 8000 us | 0 | 0 | 0 |
| 9 | 10000 us | 0 | 0 | 0 |

 Frame Delay Range (FDR) Bin Counts

| Bin | Lower Bound | Forward | Backward | Round Trip |
|-----|-------------|---------|----------|------------|
| 0 | 0 us | 10 | 10 | 10 |
| 1 | 5000 us | 0 | 0 | 0 |

 Inter-Frame Delay Variation (IFDV) Bin Counts

| Bin | Lower Bound | Forward | Backward | Round Trip |
|-----|-------------|---------|----------|------------|
| 0 | 0 us | 5 | 3 | 2 |
| 1 | 100 us | 0 | 2 | 2 |
| 2 | 200 us | 0 | 1 | 0 |
| 3 | 300 us | 0 | 3 | 1 |
| 4 | 400 us | 2 | 0 | 0 |
| 5 | 500 us | 1 | 0 | 0 |
| 6 | 600 us | 0 | 1 | 2 |
| 7 | 700 us | 0 | 0 | 0 |
| 8 | 800 us | 0 | 0 | 0 |
| 9 | 1000 us | 2 | 0 | 3 |

Monitor Commands

```
?
monitor oam-pm session "eth-pm-service-4" slm
-----
At time t = 0 sec (Base Statistics)
-----

-----
                Frames Sent      Frames Received
-----
Forward                54749            54749
Backward               54749            54749
-----

-----
Availability Counters (Und = Undetermined)
-----
                Available  Und-Avail Unavailable Und-Unavail      HLI      CHLI
-----
Forward                5475             0             0             0             0             0
Backward               5475             0             0             0             0             0
-----

-----
At time t = 10 sec (Mode: Delta)
-----

-----
                Frames Sent      Frames Received
-----
Forward                100             100
Backward              100             100
-----

-----
Availability Counters (Und = Undetermined)
-----
                Available  Und-Avail Unavailable Und-Unavail      HLI      CHLI
-----
Forward                10             0             0             0             0             0
Backward               10             0             0             0             0             0
-----

-----
At time t = 20 sec (Mode: Delta)
-----

-----
                Frames Sent      Frames Received
-----
Forward                100             100
Backward              100             100
-----

-----
Availability Counters (Und = Undetermined)
-----
```

OAM and SAA Command Reference

| | Available | Und-Avail | Unavailable | Und-Unavail | HLI | CHLI |
|----------|-----------|-----------|-------------|-------------|-----|------|
| Forward | 10 | 0 | 0 | 0 | 0 | 0 |
| Backward | 10 | 0 | 0 | 0 | 0 | 0 |

Clear Commands

saa

| | |
|--------------------|---|
| Syntax | saa-test [<i>test-name</i> [<i>owner test-owner</i>]] |
| Context | clear |
| Description | <p>Platforms Supported: 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.</p> <p>Clear the SAA results for the latest and the history for this test. If the test name is omitted, all the results for all tests are cleared.</p> |
| Parameters | <p><i>test-name</i> — Name of the SAA test. The test name must already be configured in the config>saa>test context.</p> <p><i>owner test-owner</i> — Specifies the owner of an SAA operation up to 32 characters in length.</p> <p>Default If a <i>test-owner</i> value is not specified, tests created by the CLI have a default owner “TIMOS CLI”.</p> |

session

| | |
|--------------------|---|
| Syntax | session <i>session-name</i> { dmm slm twamp-light } |
| Context | clear>oam-pm |
| Description | <p>Platforms Supported: 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.</p> <p>This command clears the raw measurement interval for the specified session and test.</p> |
| Parameters | <p><i>session-name</i> — Specifies the name of the session. 32 characters maximum.</p> <p>dmm — Keyword to clear the raw measurement interval for DMM tests.</p> <p>slm — Keyword to clear the raw measurement interval for SLM tests.</p> <p>twamp-light — Keyword to clear the raw measurement interval for TWAMP Light tests.</p> |

test-oam

Syntax test-oam

Context clear

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.
Clears the Operations, Administration, and Maintenance test parameters

twamp server

Syntax twamp server

Context clear>test-oam

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.
Clear TWAMP server statistics.

testhead

Syntax testhead result [test-name] [owner test-owner]

Context oam>clear

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.
Clear the testhead results identified by the test-name and test-owner.

Parameters *result* — Clears the test results from the latest history for the test.

test-name — Name of the test

Values ASCII string upto 32 characters in length

owner test-owner — Specifies the owner of a testhead operation.

Values ASCII string upto 32 characters in length

testhead

Syntax `testhead testhead-profile profile-id`

Context `oam>clear`

Description **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

Clear the testhead results identified by the testhead-profile.

Parameters **testhead-profile** *profile-id* — Clears the testhead profile ID to use with this run/session of testhead invocation.

Values [1..10]

Tools Command Reference

Command Hierarchies

- [Tools Dump Commands on page 443](#)
- [Tools Perform Commands on page 446](#)

Configuration Commands

Tools Dump Commands

- ```

tools
 — dump
 — accounting-policy acct-policy-id flash-write-count [clear]
 — eth-ring ring-index [clear]
 — eth-ring control-sap-tag port-id [list-in-use|next-available] (Supported only on 7210 SAS-X)
 — lag lag-id lag-id
 — ldp-treetrace {prefix ip-prefix/mask} manual-prefix ip-prefix/mask {path-destination ip-address}
 [trace-tree]
 — redundancy
 — multi-chassis
 — mc-endpoint peer ip-address
 — sync-database [peer ip-address] [port port-id | lag-id] [sync-tag sync-tag]
 [application application] [detail] [type type]
 — router router-instance
 — dintf [ip-address]
 — filter-info [verbose]
 — l3info
 — l3-stats [clear]
 — service-name service-name
 — ldp
 — fec prefix ip-prefix/mask
 — fec p2mp-id id root ip-address [detail]
 — fec vc-type {vc-type} agi agi
 — fec vc-type {ethernet|vlan} vc-id vc-id
 — instance
 — interface [ip-int-name | ip-address]
 — memory-usage
 — peer ip-address
 — session [ip-addr[:label-space]] [connection|peer|adjacency]
 — sockets
 — timers
 — mpls
 — csnf to ip-addr [from ip-addr] [strict-srlg] [srlg-group grp-id...(up to 8 max)]
 [bandwidth bandwidth] [include-bitmap bitmap] [exclude-bitmap bitmap]
 [hop-limit limit] [exclude-address excl-addr [excl-addr...(upto 8 max)]] [use-

```

- **te-metric** [exclude-node *excl-node-id* [*excl-node-id...*(upto 8 max)]] [skip-interface *interface-name*]
- **force-switch-path** lsp *lsp-name* path *path-name*
- **no force-switch-path** lsp *lsp-name*
- **ftn** [endpoint *endpoint* | sender *sender* | nexthop *nexthop* | lsp-id *lsp-id* | tunnel-id *tunnel-id* | label *start-label* *end-label*]
- **ilm** [endpoint *endpoint* | sender *sender* | nexthop *nexthop* | lsp-id *lsp-id* | tunnel-id *tunnel-id* | label *start-label* *end-label*]
- **lspinfo** [*lsp-name*] [detail]
- **memory-usage**
- **resignal** lsp *lsp-name* path *path-name* delay *minutes*
- **resignal** {p2mp-lsp *p2mp-lsp-name* p2mp-instance *p2mp-instance-name* | p2mp-delay *p2mp-minutes*}
- **te-lspinfo** [endpoint *ip-address*] [sender *ip-address*] [lspid *lsp-id*] [detail]
- **te-lspinfo** [endpoint *ip-address*] [sender *ip-address*] [lspid *lsp-id*] [detail]
- **switch-path** lsp *lsp-name* path *path-name*
- **tp-tunnel**
  - **clear** {*lsp-name* | id *tunnel-id*}
  - **force** {*lsp-name* | id *tunnel-id*}
  - **manual** {*lsp-name* | id *tunnel-id*}
  - **lockout** {*lsp-name* | id *tunnel-id*}
- **trap-suppress** *number-of-traps* *time-interval*
- **memory-usage** lsp *lsp-name* path *path-name* new-path *path-name*
- **ospf** *ospf-instance*
  - **abr** [detail]
  - **asbr** [detail]
  - **bad-packet** *interface-name*
  - **leaked-routes** [summary | detail]
  - **memory-usage** [detail]
  - **request-list** [neighbor *ip-address*] [detail]
  - **request-list** virtual-neighbor *ip-address* area-id *area-id* [detail]
  - **retransmission-list** [neighbor *ip-address*] [detail]
  - **retransmission-list** virtual-neighbor *ip-address* area-id *area-id* [detail]
  - **route-summary**
  - **route-table** *ip-prefix/mask* [type] [detail]
- **ospf3**
- **rsvp**
  - **psb** [endpoint *endpoint-address*] [sender *sender-address*] [tunnelid *tunnel-id*] [lspid *lsp-id*]
  - **rsb** [endpoint *endpoint-address*] [sender *sender-address*] [tunnelid *tunnel-id*] [lspid *lsp-id*]
  - **tcsb**[endpoint *endpoint-address*] [sender *sender-address*] [tunnelid *tunnel-id*] [lspid *lsp-id*]
  - **neighbor** [*ip-address*] [detail]
- **service**
  - **base-stats** [clear]
  - **dpipe** *service-id*
  - **dtls** *service-id*
  - **iom-stats** [clear]
  - **l2pt-diags**
  - **l2pt-diags** clear
  - **l2pt-diags** detail

- **vpls-fdb-stats** [clear]
- **vpls-mfib-stats** [clear]
- **system**
  - **cpu-pkt-stats**
- **system-resources** *slot-number*
- **system-resources** *sap-ingress-qos*

Tools Perform Commands

- tools**
  - **perform**
    - **cron**
      - **action**
        - **stop** *[action-name]* [**owner** *action-owner*] [**all**]
      - **tod**
        - **re-evaluate**
          - **customer** *customer-id* [**site** *customer-site-name*]
          - **filter ip-filter** *[filter-id]*
          - **filter ipv6-filter** *[filter-id]*
          - **filter mac-filter** *[filter-id]*
          - **service id** *service-id* [**sap** *sap-id*]
          - **tod-suite** *tod-suite-name*
  - **eth-ring**
    - **clear** *ring-index*
    - **force** *ring-index* **path** {**a**|**b**}
    - **manual** *ring-index* **path** {**a**|**b**}
  - **lag**
    - **clear-force all-mc**
    - **clear-force lag-id** *lag-id* [**sub-group** *sub-group-id*]
    - **clear-force peer-mc** *ip-address*
    - **force all-mc** {**active** | **standby**}
    - **force lag-id** *lag-id* [**sub-group** *sub-group-id*] {**active** | **standby**}
    - **force peer-mc** *peer-ip-address* {**active** | **standby**}
  - **log**
  - **test-event**
  - **redundancy**
    - **issu-post-process**
  - **router** [*router-instance*]
    - **isis** (Not supported on 7210 SAS-M device configured in Access uplink mode)
    - **mpls** (Not supported on 7210 SAS-M device configured in Access uplink mode)
      - **cspf to ip-addr** [**from** *ip-addr*] [**bandwidth** *bandwidth*] [**include-bitmap** *bit-map*] [**exclude-bitmap** *bitmap*] [**hop-limit** *limit*] [**exclude-address** *excl-addr* [*excl-addr...*(up to 8 max)]] [**use-te-metric**] [**skip-interface** *interface-name*]
      - **resignal lsp** *lsp-name* **path** *path-name* **delay** *minutes*
      - **resignal** {**p2mp-lsp** *p2mp-lsp-name* **p2mp-instance** *p2mp-instancename* | **p2mp-delay** *p2mp-minutes*}
      - **trap-suppress** *number-of-traps* *time-interval*
    - **ospf** [*ospf-instance*] (Not supported on 7210 SAS-M device configured in Access uplink mode)
      - **ldp-sync-exit**
      - **refresh-lsas**
      - **run-manual-spf**
    - **ospf3**
      - **ldp-sync-exit**
      - **refresh-lsas**
      - **run-manual-spf**
  - **service**
    - **eval-pw-template** *policy-id* [**allow-service-impact**]

- **id** *service-id*
  - **endpoint** *endpoint-name*
    - **force-switchover** *sdp-id:vc-id*
    - **no force-switchover**
  - **eval-pw-template** *policy-id* [**allow-service-impact**]
- **pw-routing**
  - **eval-expired-fec** *spoke-sdp-fec-id*
  - **eval-expired-fec** **all**
- **spoke-sdp-fec-release** *global-id[:prefix[:ac-id]]*





---

## Tools Configuration Commands

---

### Generic Commands

#### tools

**Syntax** `tools`

**Context** `root`

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command enables the context to enable useful tools for debugging purposes.

**Default** `none`

**Parameters** **dump** — Enables dump tools for the various protocols.

**perform** — Enables tools to perform specific tasks.

---

## Dump Commands

### dump

**Syntax** `dump router-name`

**Context** tools

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.  
The context to display information for debugging purposes.

**Default** none

**Parameters** *router-name* — Specify a router name, up to 32 characters in length.

**Default** Base

### accounting-policy

**Syntax** `accounting-policy acct-policy-id flash-write-count [clear]`

**Context** tools>dump

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.  
The above command dumps the total count of flash writes for the accounting policy specified by the user. The 'clear' option allows the user to clear the count maintained per accounting policy and starts the counter afresh.

**Parameters** **flash-write-count** — This is a keyword used to dump the total number of flash writes up to the present for the accounting policy specified by accounting-policy 'id'.

*acct-policy-id* — Identifies the Accounting policy.

**Values** 1 - 99

**clear** — This keyword clears statistics.

## eth-ring

**Syntax** eth-ring *ring-index* [clear]  
**eth-ring control-sap-tag** *port-id* [**list-in-use**|**next-available**] (Supported only on 7210 SAS-X)

**Context** tools>dump

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

The command displays Ethernet-ring information.

**Parameters** *ring-index* — Specify ring index.

**Values** 1 —128

**clear** — This keyword clears statistics.

## eth-ring

**Syntax** eth-ring control-sap-tag *port-id* [**list-in-use**|**next-available**] (Supported only on 7210 SAS-X)

**Context** tools>dump

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command allows the user to list the VLAN IDs on the given port in use as G.8032 control SAP tags using the 'list-in-use' option. These VLAN IDs cannot be used with any other SAP on the same port (even the Default SAP, will not receive packets with this VLAN ID).

User can obtain the next available VLAN ID to use as a control SAP tag using the 'next-available' option. This option returns a VLAN ID, after ensuring the following:

- The VLAN ID is in the range of VLAN IDs reserved for G.8032 control SAP tags.
- The VLAN ID is not being used by any service SAP on the same port.
- This VLAN ID is not being used as a control-sap-tag on any other eth-ring instance configured on the same port.

**Parameters** *port-id* — Specify the port ID.

**Values** slot/mda/port

**list-in-use** — Lists the in-use control SAP VLAN IDs on the port.

**next-available** — Lists the next available VLAN ID which can be used as control SAP TAG.

## Dump Commands

### lag

**Syntax** `lag lag-id lag-id`

**Context** `tools>dump`

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This tool displays LAG information.

**Parameters** `lag-id` — Specify an existing LAG id.

**Values** 1 — 12

```
*A:kiran3>tools>dump# lag lag-id 1
Port state : Up
Selected subgrp : 1
NumActivePorts : 2
ThresholdRising : 2
ThresholdFalling: 0
IOM bitmask : 2
Config MTU : 1522
Oper. MTU : 1522
Bandwidth : 200000

multi-chassis : NO
```

```

Indx PortId RX pkts TX pkts State Active Port Cfg Oper Speed BW AP CS
 Pri Mtu Mtu

0 1/1/1 1 1 Up yes 32768 1522 1522 1000 100000 0 2
1 1/1/2 0 0 Up yes 32768 1522 1522 1000 100000 0 2
```

### ldp-treetrace

**Syntax** `ldp-treetrace {prefix ip-prefix/mask| manual-prefix ip-prefix/mask}[path-destination ip-address] [trace-tree]`

**Context** `tools>dump`

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command displays TreeTrace information.

**Parameters** `prefix ip-prefix/mask` — Specifies the IP prefix and host bits.

**Values** host bits: must be 0  
mask: 0 — 32

## Sample Output

Automated ldp-treetrace:

Note that the **tools dump ldp-treetrace prefix** command displays entries only if **ldp-treetrace** is enabled (**configure test-oam ldp-treetrace no shutdown**).

```
*A:Dut-B# tools dump ldp-treetrace prefix 10.20.1.6/32
Discovered Paths:
=====
Id PathDst EgrNextHop ReplyRtrAddr DiscoveryTime
 DiscoveryTtl ProbeState ProbeTmOutCnt RtnCode
=== ===== ===== ===== =====
001 127.1.0.255 10.10.41.2 10.10.9.6 11/09/2010 16:15:54
 002 OK 00 EgressRtr
002 127.2.0.255 10.10.42.2 10.10.9.6 11/09/2010 16:15:54
 002 OK 00 EgressRtr
003 127.3.0.255 10.10.43.2 10.10.9.6 11/09/2010 16:15:54
 002 OK 00 EgressRtr
004 127.4.0.255 10.10.44.2 10.10.9.6 11/09/2010 16:15:54
 002 OK 00 EgressRtr
005 127.5.0.255 10.10.45.2 10.10.9.6 11/09/2010 16:15:54
 002 OK 00 EgressRtr

ldp-treetrace discovery state: Done
ldp-treetrace discovery status: ' OK '
Total number of discovered paths: 5
Total number of probe-failed paths: 0
Total number of failed traces: 0
*A:Dut-B#
```

## eth-ring

**Syntax** eth-ring *ring-index* [**clear**]

**Context** tools>dump

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This tool displays eth-ring information.

**Parameters** *ring-index* — Specifies the ring index.

**Values** 1 — 128

*clear* — Clears the eth-ring statistics.

## Dump Commands

### redundancy

**Syntax** `redundancy`

**Context** `tools>dump`

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command enables the context to dump tools for redundancy.

### multi-chassis

**Syntax** `multi-chassis`

**Context** `tools>dump>redundancy>multi-chassis`

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command enables the context to dump tools for multi-chassis redundancy.

### mc-endpoint

**Syntax** `multi-chassis`

**Context** `tools>dump>redundancy>multi-chassis`

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command dumps multi-chassis endpoint information.

**Parameters** `peer ip-address` — Specifies the peer's IP address.

### sync-database

**Syntax** `sync-database [peer ip-address] [port port-id | lag-id] [sync-tag sync-tag] [application application] [detail] [type type]`

**Context** `tools>dump>redundancy>multi-chassis`

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command dumps MCS database information.

**peer** *ip-address* — Specifies the peer's IP address.

**port** *port-id* | *lag-id* — Indicates the port or LAG ID to be synchronized with the multi-chassis peer.

**Values** *slot/mda/port* or *lag-lag-id*

**sync-tag** *sync-tag* — Specifies a synchronization tag to be used while synchronizing this port with the multi-chassis peer.

**application** *application* — Specifies a particular multi-chassis peer synchronization protocol application.

**Values** igmp-snooping: igmp-snooping  
 mc-ring: multi-chassis ring  
 sub-host-trk: subscriber host tracking

**type** *type* — Indicates the locally deleted or alarmed deleted entries in the MCS database per multi-chassis peer.

**Values** alarm-deleted, local-deleted

**detail** — Displays detailed information.

## system

**Syntax** **system**

**Context** tools>dump>system

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command dumps tools for system information.

## cpu-pkt-stats

**Syntax** **cpu-pkt-stats**

**Context** tools>dump>system

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command dumps statistics for CPU traffic.

## system-resources

**Syntax** **system-resources** *slot-number*  
**system-resources** *sap-ingress-qos*  
**system-resources** *slot-number* **g8032-control-sap-tags**

**Context** tools>dump

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command displays system resource information.

**Default** none

**Parameters** *slot-number* — Specify a specific slot to view system resources information.

**Values** 1

*sap-ingress-qos* — This command provides details on usage of resources allocated for QoS classification and different match criteria under QoS classification.

### Sample Output

**tools dump system-resources sap-ingress-qos** — The following table describes tools dump system-resource sap-ingress-qos output fields:

**Table 22: Output fieldstools dump system-resource sap-ingress-qos**

Labels	Descriptions
Total Chunks Configured	Displays the total number of chunks configured for use by SAP ingress QoS classification across all the match criteria.
Total Chunks Available	Displays the total number of chunks allotted by software for use by SAP ingress QoS classification across all the match criteria.
Number of Chunks in Use	Displays the total number of chunks in use by SAP for SAP ingress QoS classification.
Number of Free Chunks	Displays the total number of chunks available for use by SAP for SAP ingress QoS classification.
Number of Chunks in use for IP match	Displays the total number of chunks in use for by SAP that use IP classification match criteria in the SAP ingress QoS policy.



**Table 22: Output fieldstools dump system-resource sap-ingress-qos**

Labels	Descriptions
Number of Chunks in use for IPv6 match	Displays the total number of chunks in use for by SAP that use IPv6 classification match criteria in the SAP ingress QoS policy.
Number of Chunks in use for MAC match	Displays the total number of chunks in use for by SAP that use MAC classification match criteria in the SAP ingress QoS policy.
Classification Entries	The total number of Classification entries that are available/allocated/free per chunk. Information is displayed only for chunks that are in use. <b>Meters</b> - The total number of Meters that are available/allocated/free per chunk. Information is displayed only for chunks that are in use.
Number of Chunks available for use with IP match criteria	Displays the total number of chunks in use for by SAP that use IP classification match criteria in the SAP ingress QoS policy. This assumes all of the free chunks are allotted to IP classification match criteria.
Number of Chunks available for use with IPv6 match criteria	Displays the total number of chunks in use for by SAP that use IPv6 classification match criteria in the SAP ingress QoS policy. This assumes all of the free chunks are allotted to IPv6 classification match criteria.
Number of Chunks available for use with MAC match criteria	Displays the total number of chunks in use for by SAP that use MAC classification match criteria in the SAP ingress QoS policy. This assumes all of the free chunks are allotted to MAC classification match criteria.

```
*A:7210-SAS>tools>dump# system-resources tools dump system-resources sap-ingress-qos
Sap Resource Manager info at 001 d 10/11/12 04:42:00.043:
```

```
Sap Ingress Resource Usage for Slot #1, Cmplx #0:
```

```
Total Chunks Configured : 6
Total Chunks Available : 6
Number of Chunks in Use : 1
Number of Free Chunks : 5
Number of Chunks in use for IP match : 0
Number of Chunks in use for IPv6 match : 0
Number of Chunks in use for MAC match : 1
 | Classification Entries | Meters
Chunk | Type | Total |Allocated| Free | Total |Allocated| Free
-----+-----+-----+-----+-----+-----+-----+-----
```

## Dump Commands

```
0| Mac| 512| 2| 510| 256| 1| 255
```

```
Number of Chunks available for use with IP match* : 5
Number of Chunks available for use with IPv6 match* : 0
Number of Chunks available for use with MAC match* : 5
```

```
* - Assumes all remaining chunks are used
*A:Dut-A>tools>dump#
```

### The following sample displays the resource utilization for SAP egress aggregate meter on 7210 SAS-M and 7210 SAS-T:

```
*7210SAS>config>service# /tools dump system-resources
Resource Manager info at 022 h 07/27/01 22:50:44.676:
```

Hardware Resource Usage per Node:

	Total	Allocated	Free
Max System Ecmp Routes		1	

Hardware Resource Usage for Slot #1, CardType iom-sas, Cmplx #0:

	Total	Allocated	Free
SAP Ingress QoS Policies	1791	1	1790
Access Egr. QoS Policies	255	1	254
SAP Ingress Aggregate-Meter	256	0	256
Shared Qos Ingress Meters	1280	256	1024
Shared Qos Ingress CAM Entries	2560	512	2048
Mac Qos Ingress Meters	256	2	254
Mac Qos Ingress CAM Entries	512	4	508
IPv4 Qos Ingress Meters	0	0	0
IPv4 Qos Ingress CAM Entries	0	0	0
IPv6 Qos Ingress Meters	0	0	0
IPv6 Qos Ingress CAM Entries	0	0	0
IP + Mac Qos Ingress Meters	0	0	0
IP + Mac Qos Ingress CAM Entries	0	0	0
SAP Egress Aggregate-Meter Entries	512	2	510
Network Ingress Meters	256	0	256
Network Ing CAM Entries	512	0	512
Port Scheduler Policies	1023	1	1022
Ingress Shared ACL Entries	2560	0	2560
Ingress Mac ACL Entries	0	0	0
Ingress IPv4 ACL Entries	0	0	0
Ing IPv6 128 bit ACL Entries	0	0	0
Ing IPv6 64 bit ACL Entries	0	0	0
Egress Shared ACL Entries	0	0	0
Egress Mac only ACL Entries	0	0	0
Egress Mac+IPv4 ACL Entries	0	0	0
Egr IPv6 128 bit ACL Entries	0	0	0
Egr Mac+IPv6 64 bit ACL Entries	0	0	0
Ingress SAP Lookup Entries	1496	4	1492
Num VLAN-ID/Range in Con Profile	4096	0	4096

=====  
**The following sample displays the resource utilization for SAP egress aggregate meter on 7210 SAS-R6:**

A:7210SAS-053>tools>dump# system-resources  
 Resource Manager info at 001 h 02/29/16 10:03:45.035:

Hardware Resource Usage per Node:

	Total	Allocated	Free
Max System Ecmp Routes		1	
Ldp Ecmp percent value		0	
Eth-Ring Fast Flood Entries	1280	1	1279

Hardware Resource Usage for Slot #1, CardType imm-sas-b-10sfp-1sfp+, Cmplx #0:

	Total	Allocated	Free
SAP Ingress QoS Policies	1791	2	1789
Access Egr. QoS Policies	255	2	253
SAP Egress QoS Policies	2047	2	2045
SAP Ingress Aggregate-Meter	384	0	384
Shared Qos Ingress Meters	640	128	512
Shared Qos Ingress CAM Entries	1280	256	1024
Mac Qos Ingress Meters	128	6	122
Mac Qos Ingress CAM Entries	256	12	244
IPv4 Qos Ingress Meters	0	0	0
IPv4 Qos Ingress CAM Entries	0	0	0
IPv6 Qos Ingress Meters	0	0	0
IPv6 Qos Ingress CAM Entries	0	0	0
IP + Mac Qos Ingress Meters	0	0	0
IP + Mac Qos Ingress CAM Entries	0	0	0
SAP Egress Agg-Meter Entries	0	0	0
Network Ing Port Meters	128	10	118
Network Ing Port CAM Entries	256	10	246
Network Ing IpIntf Meters	128	4	124
Network Ing IpIntf CAM Entries	256	4	252
Network MPLS Exp Profile Table	31	0	31
Port Scheduler Policies	0	0	0
Queue Management Policies	500	1	499
Remark Policies	256	5	251
Shared Egr QOS MAP Entries	14	1	13
Egress QOS CAM Entries	400	6	394
Ingress Shared ACL Entries	768	0	768
Ingress Mac ACL Entries	0	0	0
Ingress IPv4 ACL Entries	0	0	0
Ing IPv6 128 bit ACL Entries	0	0	0
Ing IPv6 64 bit ACL Entries	0	0	0
Egress Shared ACL Entries	128	0	128
Egress Mac only ACL Entries	0	0	0
Egress Mac+IPv4 ACL Entries	0	0	0
Egr IPv6 128 bit ACL Entries	0	0	0
Egr Mac+IPv6 64 bit ACL Entries	0	0	0

## Dump Commands

```

 Ingress SAP Lookup Entries | 706 | 8 | 698
 Egress Sap Counter Entries | 127 | 0 | 127
 MEP Lookup CAM Entries | 256 | 0 | 256
 DN MEP Entries | 256 | 38 | 218
 Port Scheduler Mode | | 0 |
Num VLAN-ID/Range in Con Profile | 4096 | 0 | 4096

```

**The following sample displays the resource utilization for SAP egress aggregate meter on 7210 SAS-Sx/S 1/10GE:**

```

7210SAS_duth>tools>dump# system-resources sap-ingress-qos
Sap Resource Manager info at 002 d 01/15/00 02:20:05.333:

```

Sap Ingress Resource Usage for Slot #1, Cmplx #0:

```

Total Chunks Configured : 4
Total Chunks Available : 4
Number of Chunks in Use : 0
Number of Free Chunks : 4
Number of Chunks in use for IP match : 0
Number of Chunks in use for IPv6 match : 0
Number of Chunks in use for IP + Mac match : 0
Number of Chunks in use for MAC match : 0

```

Chunk	Type	Classification Entries			Meters		
		Total	Allocated	Free	Total	Allocated	Free
-----+-----+-----+-----+-----+-----+-----							
Number of Chunks available for use with IP match* : 4							
Number of Chunks available for use with IPv6 match* : 0							
Number of Chunks available for use with Ip + Mac match* : 4							
Number of Chunks available for use with MAC match* : 4							

\* - Assumes all remaining chunks are used

```

7210SAS_duth>tools>dump#

```

**The following sample displays the resource utilization for SAP egress aggregate meter on 7210 SAS-Sx 10/100GE:**

```

7210SAS>tools>dump# system-resources
Resource Manager info at 008 d 01/09/00 06:24:56.148:

```

Hardware Resource Usage per Node:

	Total	Allocated	Free
-----+-----+-----+-----			
Max System Ecmp Routes		1	

```

Hardware Resource Usage for Slot #1, CardType iom-sas, Cmplx #0:
 Total | Allocated | Free

```

SAP Ingress QoS Policies	1791	1	1790
Access Egr. QoS Policies	255	1	254
Network QoS Policies	255	2	253
Network Queue QoS Policies	255	1	254
Shared Qos Ingress Meters	3072	0	3072
Shared Qos Ingress CAM Entries	6144	0	6144
Mac Qos Ingress Meters	0	0	0
Mac Qos Ingress CAM Entries	0	0	0
IPv4 Qos Ingress Meters	0	0	0
IPv4 Qos Ingress CAM Entries	0	0	0
IPv6 Qos Ingress Meters	0	0	0
IPv6 Qos Ingress CAM Entries	0	0	0
IP + Mac Qos Ingress Meters	0	0	0
IP + Mac Qos Ingress CAM Entries	0	0	0
SAP Egress Agg-Meter Entries	0	0	0
Network Ing Meters	512	68	444
Network Ing CAM Entries	1024	68	956
Network MPLS Exp Profile Table	31	0	31
Port Scheduler Policies	1023	1	1022
Remark Policies	256	2	254
Shared Egr QOS MAP Entries	13	0	13
Ingress Shared ACL Entries	2048	0	2048
Ingress Mac ACL Entries	0	0	0
Ingress IPv4 ACL Entries	0	0	0
Ing IPv6 128 bit ACL Entries	0	0	0
Ing IPv6 64 bit ACL Entries	0	0	0
Egress Shared ACL Entries	256	0	256
Egress Mac only ACL Entries	0	0	0
Egress Mac+IPv4 ACL Entries	0	0	0
Egr IPv6 128 bit ACL Entries	0	0	0
Egr Mac+IPv6 64 bit ACL Entries	0	0	0
Ingress SAP Lookup Entries	703	0	703
TWAMP/LT Entries	32	0	32
MEP Lookup CAM Entries	256	0	256
DN MEP Entries	1024	42	982
Num VLAN-ID/Range in Con Profile	4096	0	4096
Egress TLS Mcast Entries	147455	0	147455

---

## Service Commands

### service

**Syntax** `service`

**Context** `tools>dump`

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

Use this command to configure tools to display service dump information.

### base-stats

**Syntax** `base-stats [clear]`

**Context** `tools>dump>service`

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

Use this command to display internal service statistics.

**Default** `none`

**Parameters** `clear` — Clears stats after reading.

### dpipe

**Syntax** `dpipe service-id`

**Context** `tools>dump>service`

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command displays debug information for specified service.

**Parameters** `service-id` — Displays specified service ID details.

**Values** 1 - 2147483647

## dtls

<b>Syntax</b>	<b>dtls</b> <i>service-id</i>
<b>Context</b>	tools>dump>service
<b>Description</b>	<b>Platforms Supported:</b> 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12. Use this command to display TLS service statistics.
<b>Default</b>	none
<b>Parameters</b>	<i>service-id</i> — Displays specified service ID details.

## iom-stats

<b>Syntax</b>	<b>iom-stats</b> [ <b>clear</b> ]
<b>Context</b>	tools>dump>service
<b>Description</b>	<b>Platforms Supported:</b> 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12. Use this command to display IOM message statistics.
<b>Default</b>	none
<b>Parameters</b>	<b>clear</b> — Clears stats after reading.

## l2pt-diags

<b>Syntax</b>	<b>l2pt-diags</b> <b>l2pt-diags clear</b> <b>l2pt-diags detail</b>
<b>Context</b>	tools>dump>service
<b>Description</b>	<b>Platforms Supported:</b> 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12. Use this command to display L2pt diagnostics.
<b>Default</b>	none
<b>Parameters</b>	<b>clear</b> — Clears the diags after reading. <b>detail</b> — Displays detailed information.

**Sample Output**

## Service Commands

```
A:ALA-48>tools>dump>service# l2pt-diags
[l2pt/bpdu error diagnostics]
 Error Name | Occurence | Event log
-----+-----+-----
[l2pt/bpdu forwarding diagnostics]

 Rx Frames | Tx Frames | Frame Type
-----+-----+-----
A:ALA-48>tools>dump>service#

A:ALA-48>tools>dump>service# l2pt-diags detail
[l2pt/bpdu error diagnostics]
 Error Name | Occurence | Event log
-----+-----+-----
[l2pt/bpdu forwarding diagnostics]

 Rx Frames | Tx Frames | Frame Type
-----+-----+-----
[l2pt/bpdu config diagnostics]
WARNING - service 700 has l2pt termination enabled on all access points :
 consider translating further down the chain or turning it off.
WARNING - service 800 has l2pt termination enabled on all access points :
 consider translating further down the chain or turning it off.
WARNING - service 9000 has l2pt termination enabled on all access points :
 consider translating further down the chain or turning it off.
WARNING - service 32806 has l2pt termination enabled on all access points :
 consider translating further down the chain or turning it off.
WARNING - service 90001 has l2pt termination enabled on all access points :
 consider translating further down the chain or turning it off.
A:ALA-48>tools>dump>service#
```

## vpls-fdb-stats

<b>Syntax</b>	<b>vpls-fdb-stats [clear]</b>
<b>Context</b>	tools>dump>service
<b>Description</b>	<b>Platforms Supported:</b> 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12. Use this command to display VPLS FDB statistics.
<b>Default</b>	none
<b>Parameters</b>	<b>clear</b> — Clears stats after reading.



## vpls-mfib-stats

**Syntax** `vpls-mfib-stats [clear]`

**Context** `tools>dump>service`

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

Use this command to display VPLS MFIB statistics.

**Default** none

**Parameters** `clear` — Clears stats after reading.

---

## Router Commands

### router

<b>Syntax</b>	<b>router</b> <i>router-instance</i>
<b>Context</b>	tools>dump tools>perform
<b>Description</b>	<b>Platforms Supported:</b> 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.  This command enables tools for the router instance.
<b>Default</b>	none
<b>Parameters</b>	<b>router</b> <i>router-instance</i> — Specifies the router name or service ID.  <b>Values</b> <i>router-name:</i> Base, management <b>Default</b> Base

### dintf

<b>Syntax</b>	<b>dintf</b> [ <i>ip-address</i> ]
<b>Context</b>	tools>dump>router
<b>Description</b>	<b>Platforms Supported:</b> 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.  This command displays the internal IP interface details.
<b>Parameters</b>	<i>ip-address</i> — Only displays the internal IP interface details.

### filter-info

<b>Syntax</b>	<b>filter-info</b> [ <i>verbose</i> ]
<b>Context</b>	tools>dump>router
<b>Description</b>	<b>Platforms Supported:</b> 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.  This command dumps the hardware-specific filter information.

**Parameters** `verbose` — Displays the hardware information of the filter.

## l3info

**Syntax** `lag`

**Context** `tools>dump>router`

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.  
This command dumps the hardware-specific L3 information.

## l3-stats

**Syntax** `l3-stats [clear]`

**Context** `tools>dump>router`

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.  
This command dumps the hardware-specific L3 statistics.

**Parameters** `clear` — Clears the hardware information of the filter.

## eth-ring

**Syntax** `eth-ring`

**Context** `tools>perform`

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.  
This command configures tools to control Ethernet rings.

## clear

**Syntax** `clear ring-index`

**Context** `tools>perform>eth-ring`

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

## Router Commands

This command removes all switching operational commands. The clear command is used for the following operations:

- Clears an active local administrative command (for example the Force switch or Manual switch commands).
- Triggers reversion before the WTR or WTB timer expires in case of revertive operation.
- Triggers reversion in case of non-revertive operation.

**Parameters** *ring-index* — Specifies the ring index of the Ethernet ring.

**Values** 1—128

### force

**Syntax** **force** *ring-index* **path** {a|b}

**Context** tools>perform>eth-ring

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command forces the specified path into a blocked state.

**Parameters** *ring-index* — Specifies the ring index of the Ethernet ring.

**Values** 1—128

**path** {a|b} — Specifies the path of the Ethernet ring.

### manual

**Syntax** **manual** *ring-index* **path** {a|b}

**Context** tools>perform>eth-ring

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command sets the specified eth-ring path into a blocked state.

**Parameters** *ring-index* — Specifies the ring index of the Ethernet ring.

**Values** 1—128

**path** {a|b} — Specifies the path of the Ethernet ring.

## lag

**Syntax** lag

**Context** tools>perform

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command configures tools to control LAG.

## clear-force

**Syntax** clear-force lag-id *lag-id* [**sub-group** *sub-group-id*]  
clear-force all-mc  
clear-force peer-mc *ip-address*

**Context** tools>perform>lag

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command clears a forced status.

**Parameters** lag-id *lag-id* — Specify an existing LAG id.

**Values** 1-200

all-mc — Clears all multi-chassis LAG information.

## force

**Syntax** force lag-id *lag-id* [**sub-group** *sub-group-id*] {**active** | **standby**}  
force all-mc {**active** | **standby**}  
force peer-mc *peer-ip-address* {**active** | **standby**}

**Context** tools>perform>lag

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command forces an active or standby status.

**Parameters** active — If active is selected, then all drives on the active CPM are forced.

all-mc — Clears all multi-chassis LAG information.

peer-mc — Clears mutichassis LAG peer information.

standby — If standby is selected, then all drives on the standby CPM are forced.

## Router Commands

**lag-id** *lag-id* — Specify an existing LAG id.

**Values** 1 — 6

### log

**Syntax** log

**Context** tools>perform

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.  
Tools for event logging.

### test-event

**Syntax** test-event

**Context** tools>perform>log

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.  
This command causes a test event to be generated. The test event is LOGGER event #2011 and maps to the tmnxEventSNMP trap in the TIMETRA-LOG-MIB.

### redundancy

**Syntax** redundancy

**Context** tools>perform

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.  
This command enables the context to configure redundancy parameters.

### issu-post-process

**Syntax** issu-post-process

**Context** tools>perform>redundancy

**Description** **Platforms Supported:** 7210 SAS-R6 and 7210 SAS-R12

This command forces the MPLS module to exit maintenance mode, and thus resume signaling new LSP paths, before major or minor ISSU is completed.

When the system starts major or minor ISSU procedures, MPLS will automatically be put into maintenance mode such that existing LSP paths will continue to operate normally while the node will not issue new LSP paths or a Make-Before-Break (MBB) path for existing LSPs. It will also reject requests for new LSP paths or MBB paths of existing LSPs coming from RSVP neighbors.

The MPLS module will automatically exit maintenance mode when the major or minor ISSU is completed. As such this command must only be used if the user encounters MPLS issues during the ISSU process.

## interface

**Syntax** `interface [ip-int-name | ip-address]`

**Context** `tools>dump>router>ldp`

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command displays information for an LDP interface.

**Default** none

**Parameters** *ip-int-name* — Specifies the interface name.  
*ip-address* — Specifies the IP address.

## ldp

**Syntax** `ldp`

**Context** `tools>dump>router`

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command enables dump tools for LDP.

**Default** none

## peer

**Syntax** `peer ip-address`

**Context** `tools>dump>router>ldp`

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

## Router Commands

This command displays information for an LDP peer.

**Default** none

**Parameters** *ip-address* — Specifies the IP address.

### fec

**fec p2mp-id** *id* **root** *ip-address* [**detail**]  
**fec prefix** [*ip-prefix/mask*]  
**fec vc-type** {**ethernet** | **vlan**} **vc-id** *vc-id*

**Context** tools>dump>router>ldp

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command displays information for an LDP FEC.

**Default** none

**Parameters** **fec p2mp-id** *id* — Dumps LDP active P2MP identifier bindings information.

**Values** 0 to 4294967295

**root** *ip-address* — Dumps root IP address information.

**detail** — Dumps detailed LDP active P2MP identifier bindings information or detailed root IP address information.

*ip-prefix/mask* — Specifies the IP prefix and host bits.

**Values** host bits: must be 0  
mask: 0 — 32

**vc-type** — Specifies the VC type signaled for the spoke or mesh binding to the far end of an SDP. The VC type is a 15 bit-quantity containing a value which represents the type of VC. The actual signaling of the VC type depends on the signaling parameter defined for the SDP. If signaling is disabled, the **vc-type** command can still be used to define the Dot1q value expected by the far-end provider equipment. A change of the binding's VC type causes the binding to signal the new VC type to the far end when signaling is enabled.

VC types are derived according to IETF *draft-martini-l2circuit-trans-mpls*.

- Ethernet — The VC type value for Ethernet is 0x0005.
- VLAN — The VC type value for an Ethernet VLAN is 0x0004.

*vc-id* — Specifies the virtual circuit identifier.

**Values** 1 — 4294967295



## memory-usage

**Syntax** **memory-usage**

**Context** tools>dump>router>ldp

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command displays memory usage information for LDP.

**Default** none

## session

**Syntax** **session** [*ip-address* [:*label space* ] [*connection* | *peer* | *adjacency* ]

**Context** tools>dump>router>ldp

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command displays information for an LDP session.

**Default** none

**Parameters** *ip-address* — Specifies the IP address of the LDP peer.

*label-space* — Specifies the label space identifier that the router is advertising on the interface.

**connection** — Displays connection information.

**peer** — Displays peer information.

**adjacency** — Displays hello adjacency information.

## sockets

**Syntax** **sockets**

**Context** tools>dump>router>ldp

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command displays information for all sockets being used by the LDP protocol.

**Default** none

## timers

<b>Syntax</b>	<b>timers</b>
<b>Context</b>	tools>dump>router>ldp
<b>Description</b>	<b>Platforms Supported:</b> 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12. This command displays timer information for LDP.
<b>Default</b>	none

## mpls

<b>Syntax</b>	<b>mpls</b>
<b>Context</b>	tools>dump>router
<b>Description</b>	<b>Platforms Supported:</b> 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12. This command enables the context to display MPLS information.
<b>Default</b>	none

## ftn

<b>Syntax</b>	<b>ftn</b>
<b>Context</b>	tools>dump>router>mpls
<b>Description</b>	<b>Platforms Supported:</b> 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12. This command displays FEC-to-NHLFE (FTN) dump information for MPLS. (NHLFE is the acronym for Next Hop Label Forwarding Entry.)
<b>Default</b>	none

## ilm

<b>Syntax</b>	<b>ilm</b>
<b>Context</b>	tools>dump>router>mpls
<b>Description</b>	<b>Platforms Supported:</b> 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command displays incoming label map (ILM) information for MPLS.

**Default** none

## lspinfo

**Syntax** **lspinfo** [*lsp-name*] [**detail**]

**Context** tools>dump>router>mpls

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command displays label-switched path (LSP) information for MPLS.

**Default** none

**Parameters** *lsp-name* — Specifies the name that identifies the LSP. The LSP name can be up to 32 characters long and must be unique.

**detail** — Displays detailed information about the LSP.

## cspf

**Syntax** **cspf to** ip-addr [**from** ip-addr] [**strict-srlg**] [**srlg-group** grp-id...(up to 8 max)] [**bandwidth** bandwidth] [**include-bitmap** bitmap] [**exclude-bitmap** bitmap] [**hop-limit** limit] [**exclude-address** excl-addr [excl-addr...(upto 8 max)]] [**use-te-metric**] [**exclude-node** excl-node-id] [**excl-node-id**...(upto 8 max)]] [**skip-interface** interface-name]

**Context** tools>perform>router>mpls

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command computes a CSPF path with specified user constraints.

**Default** none

**Parameters** **to** *ip-addr* — Specify the destination IP address.

**from** *ip-addr* — Specify the originating IP address.

**srlg-group** *grp-id* — Specify the

**bandwidth** *bandwidth* — Specifies the amount of bandwidth in mega-bits per second (Mbps) to be reserved.

**include-bitmap** *bitmap* — Specifies to include a bit-map that specifies a list of admin groups that should be included during setup.

**exclude-bitmap** *bitmap* — Specifies to exclude a bit-map that specifies a list of admin groups that should be included during setup.

**hop-limit** *limit* — Specifies the total number of hops a detour LSP can take before merging back onto the main LSP path.

**exclude-address** *ip-addr* — Specifies an IP address to exclude from the operation.

**use-te-metric** — Specifies whether the TE metric would be used for the purpose of the LSP path computation by CSPF.

**skip-interface** *interface-name* — Specifies a local interface name, instead of the interface address, to be excluded from the CSPF computation.

### force-switch-path

**Syntax** **force-switch-path** **lsp** *lsp-name path path-name*  
**no force-switch-path** **lsp** *lsp-name*

**Context** tools>perform>router>mpls

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command is used to configure a forced path switch.

### resignal

**Syntax** **resignal** **lsp** *lsp-name path path-name delay minutes*  
**resignal** {**p2mp-lsp** *p2mp-lsp-name p2mp-instance p2mp-instance-name* | **p2mp-delay** *p2mp-minutes*}

**Context** tools>perform>router>mpls

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

Use this command to resignal a specific LSP path.

**Default** none

**Parameters** **lsp** *lsp-name* — Specifies the name that identifies the LSP. The LSP name can be up to 32 characters long and must be unique.

**path** *path-name* — Specifies the name for the LSP path up to 32 characters in length.

**delay** *minutes* — Specifies the resignal delay in minutes.

**Values** 0 — 30

**p2mp-lsp** *p2mp-lsp-name* — Specifies an existing point-to-multipoint LSP name.

**p2mp-instance** *p2mp-instance-name* — Specifies a name that identifies the P2MP LSP instance

**p2mp-delay** *p2mp-minutes* — Specifies the delay time, in minutes.

**Values** 0 — 60

## switch-path

**Syntax** **switch-path lsp** *lsp-name path path-name*

**Context** tools>perform>router>mpls

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

Use this command to move from a standby (or an active secondary) to another standby of the same priority. If a new standby path with a higher priority or a primary path comes up after the tools perform command is executed, the path re-evaluation command runs and the path is moved to the path specified by the outcome of the re-evaluation.

**Parameters** *lsp-name* — Specifies an existing LSP name to move.

*path-name* — Specifies the path name to which to move the specified LSP.

## trap-suppress

**Syntax** **trap-suppress** [*number-of-traps*] [*time-interval*]

**Context** tools>perform>router>mpls

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command modifies thresholds for trap suppression.

**Default** none

**Parameters** *number-of-traps* — Specify the number of traps in multiples of 100. An error messages is generated if an invalid value is entered.

**Values** 100 to 1000

*time-interval* — Specify the timer interval in seconds.

**Values** 1 — 300

## tp-tunnel

**Syntax** `tp-tunnel`

**Context** `tools>perform>router>mpls`

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command enables the context to perform Linear Protection operations on an MPLS-TP LSP.

## clear

**Syntax** `clear {lsp-name | id tunnel-id}`

**Context** `tools>perform>router>mpls>tp-tunnel`

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

Clears all the MPLS-TP linear protection operational commands for the specified LSP that are currently active.

**Parameters** *lsp-name* — Specifies the name of the MPLS-TP LSP.

**Values** up to 32 characters in text

**id** *tunnel-id* — Specifies the tunnel number of the MPLS-TP LSP

**Values** 1 — 61440

## force

**Syntax** `force {lsp-name | id tunnel-id}`

**Context** `tools>perform>router>mpls>tp-tunnel`

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

Performs a force switchover of the MPLS-TP LSP from the active path to the protect path.

**Parameters** *lsp-name* — Specifies the name of the MPLS-TP LSP.

**Values** up to 32 characters in text

**id** *tunnel-id* — Specifies the tunnel number of the MPLS-TP LSP

**Values** 1 — 61440

## manual

**Syntax** `manual {/sp-name | id tunnel-id}`

**Context** `tools>perform>router>mpls>tp-tunnel`

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

Performs a manual switchover of the MPLS-TP LSP from the active path to the protect path.

**Parameters** *lsp-name* — Specifies the name of the MPLS-TP LSP.

**Values** up to 32 characters in text

*id tunnel-id* — Specifies the tunnel number of the MPLS-TP LSP

**Values** 1 — 61440

## lockout

**Syntax** `lockout {/sp-name | id tunnel-id}`

**Context** `tools>perform>router>mpls>tp-tunnel`

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

Performs a lockout of protection for an MPLS-TP LSP. This prevents a switchover to the protect path.

**Parameters** *lsp-name* — Specifies the name of the MPLS-TP LSP.

**Values** up to 32 characters in text

*id tunnel-id* — Specifies the tunnel number of the MPLS-TP LSP

**Values** 1 — 61440

## memory-usage

**Syntax** `memory-usage`

**Context** `tools>dump>router>mpls`

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command displays memory usage information for MPLS.

**Default** none

## te-lspinfo

**Syntax** **te-lspinfo** [**endpoint** *ip-address*] [**sender** *ip-address*] [**lspid** *lsp-id*] [**detail**]  
**te-lspinfo** [**endpoint** *ip-address*] [**sender** *ip-address*] [**lspid** *lsp-id*] [**detail**]

**Context** tools>dump>router>mpls

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command displays TE LSP information for MPLS.

**Default** none

## ospf

**Syntax** **ospf** [*ospf-instance*]

**Context** tools>dump>router

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command enables the context to display tools information for OSPF.

### CONFIGURATION NOTES:

- 7210 SAS-M network mode devices allow for configuration of a single instance at any given time. The instance ID can be anything other than zero. It allows 7210 SAS-M devices to be used in a network, where Multi-Instance OSPF is deployed and the node needs to use an instance ID other than the default instance ID of 0.
- Number of OSPF instances supported on different 7210 platforms are different. Kindly talk to a Nokia representative to know the supported scaling limits.

**Default** none

**Parameters** **ospf-instance** — OSPF instance.

**Values** 0— 31

## abr

**Syntax** **abr** [**detail**]

**Context** tools>dump>router>ospf

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command displays area border router (ABR) information for OSPF.



**Default** none

**Parameters** **detail** — Displays detailed information about the ABR.

## asbr

**Syntax** **asbr** [**detail**]

**Context** tools>dump>router>ospf

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command displays autonomous system border router (ASBR) information for OSPF.

**Default** none

**Parameters** **detail** — Displays detailed information about the ASBR.

## bad-packet

**Syntax** **bad-packet** [*interface-name*]

**Context** tools>dump>router>ospf

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command displays information about bad packets for OSPF.

**Default** none

**Parameters** *interface-name* — Display only the bad packets identified by this interface name.

## leaked-routes

**Syntax** **leaked-routes** [**summary** | **detail**]

**Context** tools>dump>router>ospf

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command displays information about leaked routes for OSPF.

**Default** summary

**Parameters** **summary** — Display a summary of information about leaked routes for OSPF.

**detail** — Display detailed information about leaked routes for OSPF.

## memory-usage

**Syntax** **memory-usage [detail]**

**Context** tools>dump>router>ospf

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.  
This command displays memory usage information for OSPF.

**Default** none

**Parameters** **detail** — Displays detailed information about memory usage for OSPF.

## request-list

**Syntax** **request-list [neighbor ip-address] [detail]**  
**request-list virtual-neighbor ip-address area-id area-id [detail]**

**Context** tools>dump>router>ospf

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.  
This command displays request list information for OSPF.

**Default** none

**Parameters** **neighbor ip-address** — Display neighbor information only for neighbor identified by the IP address.

**detail** — Displays detailed information about the neighbor.

**virtual-neighbor ip-address** — Displays information about the virtual neighbor identified by the IP address.

**area-id area-id** — The OSPF area ID expressed in dotted decimal notation or as a 32-bit decimal integer.

## retransmission-list

**Syntax** **retransmission-list [neighbor ip-address] [detail]**  
**retransmission-list virtual-neighbor ip-address area-id area-id [detail]**

**Context** tools>dump>router>ospf

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command displays dump retransmission list information for OSPF.

**Default** none

**Parameters** **neighbor** *ip-address* — Display neighbor information only for neighbor identified by the IP address.  
*detail* — Displays detailed information about the neighbor.  
**virtual-neighbor** *ip-address* — Displays information about the virtual neighbor identified by the IP address.  
**area-id** *area-id* — The OSPF area ID expressed in dotted decimal notation or as a 32-bit decimal integer.

## route-summary

**Syntax** **route-summary**

**Context** tools>dump>router>ospf

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command displays dump route summary information for OSPF.

**Default** none

## route-table

**Syntax** **route-table** *ip-prefix/mask* [**type**] [**detail**]

**Context** tools>dump>router>ospf

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command displays dump information about routes learned through OSPF.

**Default** none

**Parameters** **type** — Specify the type of route table to display information.

**Values** intra-area, inter-area, external-1, external-2, nssa-1, nssa-2

**detail** — Displays detailed information about learned routes.

## ospf3

<b>Syntax</b>	<b>ospf3</b>
<b>Context</b>	tools>dump>router
<b>Description</b>	<b>Platforms Supported:</b> 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.  This command enables the context to display tools information for OSPF3.
<b>Default</b>	none

## refresh-lsas

<b>Syntax</b>	<b>refresh-lsas [lsa-type] [area-id]</b>
<b>Context</b>	tools>perform>router>ospf tools>perform>router>ospf3
<b>Description</b>	<b>Platforms Supported:</b> 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.  This command refreshes LSAs for OSPF.
<b>Default</b>	none
<b>Parameters</b>	<b>lsa-type</b> — Specify the LSA type using allow keywords. <b>Values</b> router, network, summary, asbr, extern, nssa, opaque <b>area-id</b> — The OSPF area ID expressed in dotted decimal notation or as a 32-bit decimal integer. <b>Values</b> 0 — 4294967295

## run-manual-spf

<b>Syntax</b>	<b>run-manual-spf externals-only</b>
<b>Context</b>	tools>perform>router>ospf tools>perform>router>ospf3
<b>Description</b>	<b>Platforms Supported:</b> 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.  This command runs the Shortest Path First (SPF) algorithm.
<b>Default</b>	none

**Parameters** `externals-only` — Specify the route preference for OSPF external routes.

## rsvp

**Syntax** `rsvp`

**Context** `tools>dump>router`

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command enables the context to display RSVP information.

**Default** `none`

## psb

**Syntax** `psb [endpoint endpoint-address] [sender sender-address] [tunnelid tunnel-id] [lspid lsp-id]`

**Context** `tools>dump>router>rsvp`

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command displays path state block (PSB) information for RSVP.

When a PATH message arrives at an LSR, the LSR stores the label request in the local PSB for the LSP. If a label range is specified, the label allocation process must assign a label from that range.

The PSB contains the IP address of the previous hop, the session, the sender, and the TSPEC. This information is used to route the corresponding RESV message back to LSR 1.

**Default** `none`

**Parameters** `endpoint endpoint-address` — Specifies the IP address of the last hop.

`sender sender-address` — Specifies the IP address of the sender.

`tunnelid tunnel-id` — Specifies the SDP ID.

**Values** 0 — 4294967295

`lspid lsp-id` — Specifies the label switched path that is signaled for this entry.

**Values** 1 — 65535

### Sample Output

```
A:Dut-A# config>router>mpls# tools dump router rsvp psb detail
```

```

PSB:
```

## Router Commands

```
P2P: Session (To: 10.20.1.4 - 61441 - 10.20.1.1), Sender (10.20.1.1 -
2) PHop 255.255.255.255

PSB CurrState: BACKUPS_CONNECTED PrevState: BACKUPS_INIT Flags: 0x0
LocalLabel 0 OutLabel 131070
Incoming IfIndex: Interface: Local API(-1)
Refresh interval 0, Send Path refresh in 3 secs, Path Refresh timeout
0 secs
PrevHop: Ctype 1 Addr 255.255.255.255, LIH 0
DnStream Nbr: Addr-> 10.20.1.3 IfIndex ip-10.10.2.1(3)
UpStream Neighbor is NULLP
Session Attribute:
 Session Name: bypass-node10.20.1.2
 HoldPri: 0 SetupPri: 7 Flags: 0x2
 Ctype: 7, IncludeGroup: 0x0 IncludeAllGroup: 0x0 ExcludeGroup: 0x0
ClassType: Absent
TSpec: Flags 0x8000 QOSC 0, PDR (infinity), PBS 0.000 bps, CDR (0.000
bps) MTU: 0
CSPF Hop List: ->
 (1) UnnumIfId 3 RtrId 10.20.1.1 EgrAdmGrp 0x0 (Strict)
 (2) UnnumIfId 2 RtrId 10.20.1.3 EgrAdmGrp 0x0 (Strict)
 (3) UnnumIfId 5 RtrId 10.20.1.4 EgrAdmGrp 0x0 (Strict)
PSB RRO : ->
 (1) * Flags : 0x0 : U
 (1) * UnInf : 10.20.1.1, 3
PSB SENT RRO : ->
 (1) * Flags : 0x0 : U
 (1) * UnInf : 10.20.1.1, 3
PSB FILTERSPEC RRO : ->
 (1) * Flags : 0x0 : U
 (1) * UnInf : 10.20.1.3, 2
 (2) * Flags : 0x1 : Global
 (2) * Label : 131070
 (3) * Flags : 0x0 : U
 (3) * UnInf : 10.20.1.4, 5
 (4) * Flags : 0x1 : Global
 (4) * Label : 131070
PSB ERO : ->
 (1) Unnumbered RouterId 10.20.1.1, LinkId 3, Strict
 (2) Unnumbered RouterId 10.20.1.3, LinkId 2, Strict
 (3) Unnumbered RouterId 10.20.1.4, LinkId 5, Strict
PSB SENT ERO : ->
 (1) Unnumbered RouterId 10.20.1.3, LinkId 2, Strict
 (2) Unnumbered RouterId 10.20.1.4, LinkId 5, Strict
SendTempl: Sender:10.20.1.1_2
AdSpec Present - Flags: 0x0
 AdSpec General
 - Service Break bit : 0x0
 - IS Hop Count : 0x0
 - Path Bandwidth Estimate : 0x0
 - Minimum Path latency : 0x0
 - Composed path MTU : 0

Num Paths Received :0
Num Paths Transmitted:5
Num Resvs Received :8
Num Resvs Transmitted:0
```

```

Num Summmary Paths Received :0
Num Summmary Paths Transmitted:0
Num Summmary Resvs Received :0
Num Summmary Resvs Transmitted:0
Created at 91359 (26 secs back)

PSB:
P2P: Session (To: 10.20.1.6 - 1 - 10.20.1.1), Sender (10.20.1.1 -
30208) PHop 0.0.0.0

PSB CurrState: PRIMARYS_CONNECTED PrevState: PRIMARYS_INIT Flags: 0x8
LocalLabel 0 OutLabel 131071
Incoming IfIndex: Interface: Local API(-1)
Refresh interval 5, Send Path refresh in 4 secs, Path Refresh timeout
0 secs
PrevHop: Ctype 1 Addr 0.0.0.0, LIH 0
DnStream Nbr: Addr-> 10.20.1.2 IfIndex ip-10.10.1.1(2)
UpStream Neighbor is NULLP
Session Attribute:
 Session Name: 1::1
 HoldPri: 0 SetupPri: 7 Flags: 0x17
 Ctype: 7, IncludeGroup: 0x0 IncludeAllGroup: 0x0 ExcludeGroup: 0x0
ClassType: Absent
TSpec: Flags 0x8000 QOSC 1, PDR (infinity), PBS 0.000 bps, CDR (0.000
bps) MTU: 0
CSPF Hop List: ->
 (1) UnnumIfId 2 RtrId 10.20.1.1 EgrAdmGrp 0x0 (Strict)
 (2) UnnumIfId 2 RtrId 10.20.1.2 EgrAdmGrp 0x0 (Strict)
 (3) UnnumIfId 2 RtrId 10.20.1.4 EgrAdmGrp 0x0 (Strict)
 (4) UnnumIfId 2 RtrId 10.20.1.6 EgrAdmGrp 0x0 (Strict)
PSB RRO : ->
 (1) * Flags : 0x9 : U LP_AVAIL NODE
 (1) * UnInf : 10.20.1.1, 2
PSB SENT RRO : ->
 (1) * Flags : 0x0 : U
 (1) * UnInf : 10.20.1.1, 2
PSB FILTERSPEC RRO : ->
 (1) * Flags : 0x9 : U LP_AVAIL NODE
 (1) * UnInf : 10.20.1.2, 2
 (2) * Flags : 0x1 : Global
 (2) * Label : 131071
 (3) * Flags : 0x1 : U LP_AVAIL
 (3) * UnInf : 10.20.1.4, 2
 (4) * Flags : 0x1 : Global
 (4) * Label : 131071
 (5) * Flags : 0x0 : U
 (5) * UnInf : 10.20.1.6, 2
 (6) * Flags : 0x1 : Global
 (6) * Label : 131071
PSB ERO : ->
 (1) Unnumbered RouterId 10.20.1.2, LinkId 2, Strict
 (2) Unnumbered RouterId 10.20.1.4, LinkId 2, Strict
 (3) Unnumbered RouterId 10.20.1.6, LinkId 2, Strict
PSB SENT ERO : ->
 (1) Unnumbered RouterId 10.20.1.2, LinkId 2, Strict

```

## Router Commands

```
(2) Unnumbered RouterId 10.20.1.4, LinkId 2, Strict
(3) Unnumbered RouterId 10.20.1.6, LinkId 2, Strict
SendTempl: Sender:10.20.1.1_30208
AdSpec not present
FRR: Flags 0x2 HopLimit 16 SetupPri 7 HoldPri 0 IncludeAny 0x0
ExcludeAny 0x0 IncludeAll 0x0
PLR: Flag (0x166) State PLRS_BYPASS_UP AvoidNodeId 10.20.1.2 inIntf -1
inLabel 0
PLR: FRRRequestCount: 1 CSPFFailures: 0 ProtectionType: NodeProtect

Num Paths Received :0
Num Paths Transmitted:5
Num Resvs Received :5
 Num Resvs Transmitted:0

Num Summmary Paths Received :0
Num Summmary Paths Transmitted:0
Num Summmary Resvs Received :0
Num Summmary Resvs Transmitted:0
Created at 91359 (28 secs back)

Total PSB Count : 2

A:Dut-A# config>router>mpls#
```

## rsb

<b>Syntax</b>	<b>rsb</b> [ <b>endpoint</b> <i>endpoint-address</i> ] [ <b>sender</b> <i>sender-address</i> ] [ <b>tunnelid</b> <i>tunnel-id</i> ] [ <b>lspid</b> <i>lsp-id</i> ]
<b>Context</b>	tools>dump>router>rsvp
<b>Description</b>	<b>Platforms Supported:</b> 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12. This command displays RSVP Reservation State Block (RSB) information.
<b>Default</b>	none
<b>Parameters</b>	<b>endpoint</b> <i>endpoint-address</i> — Specifies the IP address of the last hop. <b>sender</b> <i>sender-address</i> — Specifies the IP address of the sender. <b>tunnelid</b> <i>tunnel-id</i> — Specifies the SDP ID. <b>Values</b> 0 — 4294967295 <b>lspid</b> <i>lsp-id</i> — Specifies the label switched path that is signaled for this entry. <b>Values</b> 1 — 65535



**Sample Output**

```
A:Dut-A# config>router>mpls# tools dump router rsvp rsb detail
```

```

RSB:
 EndPt 10.20.1.4 Tid 61441 XTid 10.20.1.1 SnDR 10.20.1.1 LspId 2
 ifIndex 3 NHop 20.20.1.3
 Style FF, refresh in 0 secs
 RSVP NextHop 20.20.1.3, LIH 3 (TLV: RtrId 10.20.1.3 IntfId 2)
 CT Shared Reservation Info:
 No Reservation:
 FlowSpec :Flags 0x8000 QOSC 1, PDR (infinity), PBS 0.000 bps, CDR
 (0.000 bps)
 CBS 0, EBS 0, RSpecR 0, RSpecS 0 MTU 1500 MPU 20
 FwdFlowspec :Flags 0x0 QOSC 0, PDR (0.000 bps), PBS 0.000 bps, CDR
 (0.000 bps)
 CBS 0, EBS 0, RSpecR 0, RSpecS 0 MPU 0
 FilterSpec:
 Timeout in : 26 secs, LocLabel: 0 Sender: 10.20.1.1 lspId: 2 OutIfId:
 0
 RRO :
 (1) * Flags : 0x0 : U
 (1) * UnInf : 10.20.1.3, 2
 (2) * Flags : 0x1 : Global
 (2) * Label : 131070
 (3) * Flags : 0x0 : U
 (3) * UnInf : 10.20.1.4, 5
 (4) * Flags : 0x1 : Global
 (4) * Label : 131070

RSB:
 EndPt 10.20.1.6 Tid 1 XTid 10.20.1.1 SnDR 0.0.0.0 LspId 0 ifIndex
 2 NHop 20.20.1.2
 Style SE, refresh in 0 secs
 RSVP NextHop 20.20.1.2, LIH 2 (TLV: RtrId 10.20.1.2 IntfId 2)
 CT Shared Reservation Info:
 No Reservation:
 FlowSpec :Flags 0x8000 QOSC 1, PDR (infinity), PBS 0.000 bps, CDR
 (0.000 bps)
 CBS 0, EBS 0, RSpecR 0, RSpecS 0 MTU 1496 MPU 20
 FwdFlowspec :Flags 0x0 QOSC 0, PDR (0.000 bps), PBS 0.000 bps, CDR
 (0.000 bps)
 CBS 0, EBS 0, RSpecR 0, RSpecS 0 MPU 0
 FilterSpec:
 Timeout in : 21 secs, LocLabel: 0 Sender: 10.20.1.1 lspId: 30208
 OutIfId: 0
 RRO :
 (1) * Flags : 0x9 : U LP_AVAIL NODE
 (1) * UnInf : 10.20.1.2, 2
 (2) * Flags : 0x1 : Global
 (2) * Label : 131071
 (3) * Flags : 0x1 : U LP_AVAIL
 (3) * UnInf : 10.20.1.4, 2
 (4) * Flags : 0x1 : Global
 (4) * Label : 131071
```

## Router Commands

```
(5) * Flags : 0x0 : U
(5) * UnInf : 10.20.1.6, 2
(6) * Flags : 0x1 : Global
(6) * Label : 131071
```

---

```
Total RSB Count : 2
```

```
A:Dut-A# config>router>mpls#
```

### tcsb

**Syntax** `tcsb [endpoint endpoint-address] [sender sender-address] [tunnelid tunnel-id] [lspid lsp-id]`

**Context** `tools>dump>router>rsvp`

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command displays RSVP traffic control state block (TCSB) information.

**Default** none

**Parameters** **endpoint** *endpoint-address* — Specifies the IP address of the egress node for the tunnel supporting this session.

**sender** *sender-address* — Specifies the IP address of the sender node for the tunnel supporting this session. It is derived from the source address of the associated MPLS LSP definition.

**tunnelid** *tunnel-id* — Specifies the IP address of the ingress node of the tunnel supporting this RSVP session.

**Values** 0 — 4294967295

**lspid** *lsp-id* — Specifies the label switched path that is signaled for this entry.

**Values** 1 — 65535

### static-route

**Syntax** `static-route ldp-sync-status`

**Context** `tools>dump>router`

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command displays the sync status of LDP interfaces that static-route keeps track of.

---

## Performance Tools

### perform

**Syntax** perform

**Context** tools

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command enables the context to enable tools to perform specific tasks.

**Default** none

### cron

**Syntax** cron

**Context** tools>perform

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command enables the context to perform CRON (scheduling) control operations.

**Default** none

### action

**Syntax** action

**Context** tools>perform>cron

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command enables the context to stop the execution of a script started by CRON action. See the **stop** command.

## stop

- Syntax** `stop [action-name] [owner action-owner] [all]`
- Context** `tools>perform>cron>action`
- Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.  
This command stops execution of a script started by CRON action.
- Parameters** *action-name* — Specifies the action name.  
**Values** Maximum 32 characters.  
*owner action-owner* — Specifies the owner name.  
**Default** TiMOS CLI  
**all** — Specifies to stop all CRON scripts.

## tod

- Syntax** `tod`
- Context** `tools>perform>cron`
- Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.  
This command enables the context for tools for controlling time-of-day actions.
- Default** none

## re-evaluate

- Syntax** `re-evaluate`
- Context** `tools>perform>cron>tod`
- Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.  
This command enables the context to re-evaluate the time-of-day state.
- Default** none

## customer

**Syntax** **customer** *customer-id* [**site** *customer-site-name*]

**Context** tools>perform>cron>tod>re-eval

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command re-evaluates the time-of-day state of a multi-service site.

**Parameters** *customer-id* — Specify an existing customer ID.

**Values** 1 — 2147483647

*site customer-site-name* — Specify an existing customer site name.

## filter

**Syntax** **filter** *ip-filter* [*filter-id*]  
**filter** *ipv6-filter* [*filter-id*]  
**filter** *mac-filter* [*filter-id*]

**Context** tools>perform>cron>tod>re-eval

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command re-evaluates the time-of-day state of a filter entry.

**Parameters** *filter-type* — Specify the filter type.

**Values** ip-filter, mac-filter

*filter-id* — Specify an existing filter ID.

**Values** 1 — 65535

## service

**Syntax** **service id** *service-id* [**sap** *sap-id*]

**Context** tools>perform>cron>tod>re-eval

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command re-evaluates the time-of-day state of a SAP.

**Parameters** *r service-id* — Specify the an existing service ID.

**Values** 1 — 2147483647

*sap sap-id* — Specifies the physical port identifier portion of the SAP definition. See [Common CLI Command Descriptions on page 501](#) for CLI command syntax.

## tod-suite

**Syntax** `tod-suite tod-suite-name`

**Context** `tools>perform>cron>tod>re-eval`

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command re-evaluates the time-of-day state for the objects referring to a tod-suite.

**Parameters** *tod-suite-name* — Specify an existing TOD nfame.

## ldp-sync-exit

**Syntax** `[no] ldp-sync-exit`

**Context** `tools>perform>router>isis`  
`tools>perform>router>ospf`  
`tools>perform>router>ospf3`

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command restores the actual cost of an interface at any time. When this command is executed, IGP immediately advertises the actual value of the link cost for all interfaces which have the IGP-LDP synchronization enabled if the currently advertised cost is different.

## run-manual-spf

**Syntax** `run-manual-spf`

**Context** `tools>perform>router>isis`  
`tools>perform>router>ospf`  
`tools>perform>router>ospf3`

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command runs the Shortest Path First (SPF) algorithm or OSPF or ISIS.

## isis

Note : This command is not supported on 7210 SAS-M and 7210 SAS-T devices configured in Access uplink mode

**Syntax** **isis**

**Context** tools>perform>router

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T (network mode), 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command enables the context to configure tools to perform certain ISIS tasks.

## mpls

Note : This command is not supported on 7210 SAS-M and 7210 SAS-T devices configured in Access uplink mode

**Syntax** **mpls**

**Context** tools>perform>router

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T (network mode), 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command enables the context to perform specific MPLS tasks.

**Default** none

## cspf

**Syntax** **cspf to** *ip-addr* [**from** *ip-addr*] [**bandwidth** *bandwidth*] [**include-bitmap** *bitmap*] [**exclude-bitmap** *bitmap*] [**hop-limit** *limit*] [**exclude-address** *excl-addr* [*excl-addr...*(up to 8 max)]] [**use-te-metric**] [**skip-interface** *interface-name*] [**ds-class-type** *class-type*] [**cspf-reqtype** *req-type*]

**Context** tools>perform>router>mpls

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command computes a CSPF path with specified user constraints.

**Default** none

**Parameters** **to** *ip-addr* — Specify the destination IP address.  
**from** *ip-addr* — Specify the originating IP address.

**bandwidth** *bandwidth* — Specifies the amount of bandwidth in mega-bits per second (Mbps) to be reserved.

**Values** 1 - 100000 in Mbps

**include-bitmap** *bitmap* — Specifies to include a bit-map that specifies a list of admin groups that should be included during setup.

**Values** 0 - 4294967295 - accepted in decimal, hex(0x) or binary(0b)

**exclude-bitmap** *bitmap* — Specifies to exclude a bit-map that specifies a list of admin groups that should be included during setup.

**Values** 0 - 4294967295 - accepted in decimal, hex(0x) or binary(0b)

**hop-limit** *limit* — Specifies the total number of hops a detour LSP can take before merging back onto the main LSP path.

**Values** 1- 255

**exclude-address** *ip-addr* — Specifies an IP address to exclude from the operation.

**use-te-metric** — Specifies whether the TE metric would be used for the purpose of the LSP path computation by CSPF.

**skip-interface** *interface-name* — Specifies a local interface name, instead of the interface address, to be excluded from the CSPF computation.

## resignal

**Syntax** **resignal lsp** *lsp-name* **path** *path-name* **delay** *minutes*  
**resignal** {**p2mp-lsp** *p2mp-lsp-name* **p2mp-instance** *p2mp-instancename* | **p2mp-delay** *p2mp-minutes*}

**Context** tools>perform>router>mpls

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

Use this command to resignal a specific LSP path or point-to-multipoint instance.

Default none

**Parameters** **lsp** *lsp-name* — Specifies the name that identifies the LSP. The LSP name can be up to 32 characters long and must be unique.

**path** *path-name* — Specifies the name for the LSP path up, to 32 characters in length.

**delay** *minutes* — Specifies the resignal delay in minutes.

**Values** 0 — 30

**p2mp-lsp** *p2mp-lsp-name* — Specifies an existing point-to-multipoint LSP name.

**p2mp-instance** *p2mp-instance-name* — Specifies a name that identifies the P2MP LSP instance



**p2mp-delay** *p2mp-minutes* — Specifies the delay time, in minutes.

**Values** 0 — 60

## trap-suppress

**Syntax** **trap-suppress** [*number-of-traps*] [*time-interval*]

**Context** tools>perform>router>mpls

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command modifies thresholds for trap suppression.

**Default** none

**Parameters** *number-of-traps* — Specify the number of traps in multiples of 100. An error messages is generated if an invalid value is entered.

**Values** 100 to 1000

*time-interval* — Specify the timer interval in seconds.

**Values** 1 — 300

## ospf

**Syntax** ospf

**Context** tools>perform>router

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command enables the context to perform specific OSPF tasks.

**Note:** This command is not supported on 7210 SAS-M and 7210 SAS-T devices configured in Access uplink mode

**Default** none

## ldp-sync-exit

<b>Syntax</b>	<b>[no] ldp-sync-exit</b>
<b>Context</b>	tools>perform>router>isis tools>perform>router>ospf
<b>Description</b>	<b>Platforms Supported:</b> 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.  This command restores the actual cost of an interface at any time. When this command is executed, IGP immediately advertises the actual value of the link cost for all interfaces which have the IGP-LDP synchronization enabled if the currently advertised cost is different.

## service

<b>Syntax</b>	<b>services</b>
<b>Context</b>	tools>perform
<b>Description</b>	<b>Platforms Supported:</b> 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.  This command enables the context to configure tools for services.

## id

<b>Syntax</b>	<b>id <i>service-id</i></b>
<b>Context</b>	tools>perform>service
<b>Description</b>	<b>Platforms Supported:</b> 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.  This command enables the context to configure tools for a specific service.
<b>Parameters</b>	<i>service-id</i> — Specify an existing service ID.  <b>Values</b> 1 — 2147483647

## endpoint

<b>Syntax</b>	<b>endpoint <i>endpoint-name</i></b>
<b>Context</b>	tools>perform>service>id
<b>Description</b>	<b>Platforms Supported:</b> 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command enables the context to configure tools for a specific VLL service endpoint.

**Parameters** *endpoint-name* — Specify an existing VLL service endpoint name.

## force-switchover

**Syntax** **force-switchover** *sdp-id:vc-id*  
**no force-switchover**

**Context** tools>perform>service>id

**Description** **Platforms Supported:** 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12.

This command forces a switch of the active spoke SDP for the specified service.

**Parameters** *sdp-id:vc-id* — Specify an existing spoke SDP for the service.

### Sample Output

```
A:Dut-B# tools perform service id 1 endpoint mcep-t1 force-switchover 221:1
*A:Dut-B# show service id 1 endpoint
=====
Service 1 endpoints
=====
Endpoint name : mcep-t1
Description : (Not Specified)
Revert time : 0
Act Hold Delay : 0
Ignore Standby Signaling : false
Suppress Standby Signaling : false
Block On Mesh Fail : true
Multi-Chassis Endpoint : 1
MC Endpoint Peer Addr : 3.1.1.3
Psv Mode Active : No
Tx Active : 221:1(forced)
Tx Active Up Time : 0d 00:00:17
Revert Time Count Down : N/A
Tx Active Change Count : 6
Last Tx Active Change : 02/14/2009 00:17:32

Members

Spoke-sdp: 221:1 Prec:1 Oper Status: Up
Spoke-sdp: 231:1 Prec:2 Oper Status: Up
=====
*A:Dut-B#
```

## eval-pw-template

<b>Syntax</b>	<b>eval-pw-template</b>
<b>Context</b>	tools>perform>service
<b>Description</b>	<b>Platforms Supported:</b> 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12. This command re-evaluates the pseudowire template policy.
<b>Parameters</b>	<i>policy-id</i> — Specifies the pseudowire template policy.

## eval-expired-fec

<b>Syntax</b>	<b>eval-expired-fec</b> spoke-sdp-fec-id <b>eval-expired-fec</b> all
<b>Context</b>	tools>perform>service>pw-routing
<b>Description</b>	<b>Platforms Supported:</b> 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12. This command will reset the retry counter and retry timer for the specified spoke-SDP and attempt to establish the spoke-SDP again.

## spoke-sdp-fec-release

<b>Syntax</b>	<b>spoke-sdp-fec-release</b> <i>global-id[:prefix[:ac-id]]</i>
<b>Context</b>	tools>perform>service
<b>Description</b>	<b>Platforms Supported:</b> 7210 SAS-M, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-X, 7210 SAS-R6, and 7210 SAS-R12. This command clears the MS-PW bindings associated with a particular SAII or TAIID on an S-PE.
<b>Parameters</b>	<i>global-id</i> — The Global ID of either the SAII or TAIID of the MS-PW to be released. <b>Values</b> [1..4294967295] <i>prefix</i> — The prefix of either the SAII or TAIID of the MS-PW to be released. <b>Values</b> ipv4-formatted address: a.b.c.d   [1..4294967295] <i>ac-id</i> — The AC-ID of either the SAII or TAIID of the MS-PW to be released. <b>Values</b> [1..4294967295]

# Common CLI Command Descriptions

---

## In This Chapter

This chapter provides CLI syntax and command descriptions for SAP and port commands.

Topics in this chapter include:

- [SAP Syntax on page 502](#)
- [Port Syntax on page 502](#)

## Common Service Commands

### sap

**Syntax** [no] sap *sap-id*

**Description** This command specifies the physical port identifier portion of the SAP definition.

**Parameters** *sap-id* — Specifies the physical port identifier portion of the SAP definition.

The *sap-id* can be configured in one of the following formats:

Type	Syntax	Example
port-id	<i>slot/mda/port[.channel]</i>	1/1/5
null	[ <i>port-id</i>   <i>lag-id</i> ]	<i>port-id</i> : 1/1/3 <i>lag-id</i> : lag-3
dot1q	[ <i>port-id</i>   <i>lag-id</i> ]:qtag1	<i>port-id</i> :qtag1: 1/1/3:100 <i>lag-id</i> :qtag1:lag-3:102
qinq	[ <i>port-id</i>   <i>lag-id</i> ]:qtag1.qtag2	<i>port-id</i> :qtag1.qtag2: 1/1/3:100.10 <i>lag-id</i> :qtag1.qtag2: lag-10:

### port

**Syntax** port *port-id*

**Description** This command specifies a port identifier.

**Parameters** *port-id* — The *port-id* can be configured in one of the following formats.

Values	port-id	slot/mda/port[.channel]
	lag-id	<i>lag-id</i>
	lag	keyword
	id	1— 200

# STANDARDS AND PROTOCOL SUPPORT

---



**Note:** The information presented is subject to change without notice.

Nokia assumes no responsibility for inaccuracies contained herein.

M(A,N) means 7210 SAS-M in both Access-uplink mode and Network mode; Similarly M(N) means 7210 SAS-M in network mode only

T(A,N) means 7210 SAS-M in both Access-uplink mode and Network mode; Similarly T(N) means 7210 SAS-T in network mode only

K5 means 7210 SAS-K 2F2T1C

K12 means 7210 SAS-K 2F4T6C

Sx/S-1/10GE means all variants of 7210 SAS-Sx 1/10GE and 7210 SAS-S 1/10GE platforms

Sx-1/10GE means only the variants of 7210 SAS-Sx 1/10G

R6 means 7210 SAS-R6

R12 means 7210 SAS-R12

D means 7210 SAS-D and 7210 SAS-D ETR; if a line item applies to 7210 SAS-D ETR, then it is indicated as D-ETR

E means 7210 SAS-E

X means 7210 SAS-X

## BGP

draft-ietf-idr-add-paths-04, Advertisement of Multiple Paths in BGP is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

draft-ietf-sidr-origin-validation-signaling-04, BGP Prefix Origin Validation State Extended Community is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

- RFC 1772, Application of the Border Gateway Protocol in the Internet is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 1997, BGP Communities Attribute is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 2385, Protection of BGP Sessions via the TCP MD5 Signature Option is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 2439, BGP Route Flap Damping is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 2545, Use of BGP-4 Multiprotocol Extensions for IPv6 Inter-Domain Routing is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 2858, Multiprotocol Extensions for BGP-4 is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 2918, Route Refresh Capability for BGP-4 is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx-10/100GE, R6, and R12
- RFC 3107, Carrying Label Information in BGP-4 is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 3392, Capabilities Advertisement with BGP-4 is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 4271, A Border Gateway Protocol 4 (BGP-4) is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 4360, BGP Extended Communities Attribute is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 4364, BGP/MPLS IP Virtual Private Networks (VPNs) is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 4456, BGP Route Reflection: An Alternative to Full Mesh Internal BGP (IBGP) is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 4659, BGP/MPLS IP Virtual Private Network (VPN) Extension for IPv6 VPN is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 4724, Graceful Restart Mechanism for BGP (Helper Mode) is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 4760, Multiprotocol Extensions for BGP-4 is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 4798, Connecting IPv6 Islands over IPv4 MPLS Using IPv6 Provider Edge Routers (6PE) is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 4893, BGP Support for Four-octet AS Number Space is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 5004, Avoid BGP Best Path Transitions from One External to Another is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 5291, Outbound Route Filtering Capability for BGP-4 is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12



RFC 5668, 4-Octet AS Specific BGP Extended Community is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 6811, Prefix Origin Validation is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

## Circuit Emulation

RFC 4553, Structure-Agnostic Time Division Multiplexing (TDM) over Packet (SAToP) is supported on M(N)

RFC 5086, Structure-Aware Time Division Multiplexed (TDM) Circuit Emulation Service over Packet Switched Network (CESoPSN) is supported on M(N)

RFC 5287, Control Protocol Extensions for the Setup of Time-Division Multiplexing (TDM) Pseudowires in MPLS Networks is supported on M(N)

## Ethernet

IEEE 802.1AB, Station and Media Access Control Connectivity Discovery is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

IEEE 802.1ad, Provider Bridges is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, and R12

IEEE 802.1ag, Connectivity Fault Management is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

IEEE 802.1ah, Provider Backbone Bridges is supported on M(N), X, and T(N)

IEEE 802.1ax, Link Aggregation is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

IEEE 802.1D, MAC Bridges is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

IEEE 802.1p, Traffic Class Expediting is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

IEEE 802.1Q, Virtual LANs is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

IEEE 802.1s, Multiple Spanning Trees is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

IEEE 802.1w, Rapid Reconfiguration of Spanning Tree is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

IEEE 802.1X, Port Based Network Access Control is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

IEEE 802.3ab, 1000BASE-T is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

- IEEE 802.3ac, VLAN Tag is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- IEEE 802.3ad, Link Aggregation is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- IEEE 802.3ae, 10 Gb/s Ethernet is supported on M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- IEEE 802.3ah, Ethernet in the First Mile is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- IEEE 802.3ba, 40 Gb/s and 100 Gb/s Ethernet is supported on R6 and R12
- IEEE 802.3i, Ethernet is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- IEEE 802.3u, Fast Ethernet is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- IEEE 802.3z, Gigabit Ethernet is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- ITU-T G.8032, Ethernet Ring Protection Switching is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- ITU-T Y.1731, OAM functions and mechanisms for Ethernet based networks is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

### Fast Reroute

- draft-ietf-rtgwg-lfa-manageability-08, Operational management of Loop Free Alternates is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 5286, Basic Specification for IP Fast Reroute: Loop-Free Alternates is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

### IP — General

- draft-grant-tacacs-02, The TACACS+ Protocol is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 768, User Datagram Protocol is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 793, Transmission Control Protocol is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 854, TELNET Protocol Specifications is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 951, Bootstrap Protocol (BOOTP) is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

- RFC 1034, Domain Names - Concepts and Facilities is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 1035, Domain Names - Implementation and Specification is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 1350, The TFTP Protocol (revision 2) is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 1534, Interoperation between DHCP and BOOTP is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 1542, Clarifications and Extensions for the Bootstrap Protocol is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 2131, Dynamic Host Configuration Protocol is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 2347, TFTP Option Extension is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 2348, TFTP Blocksize Option is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 2349, TFTP Timeout Interval and Transfer Size Options is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 2428, FTP Extensions for IPv6 and NATs is supported on D, E, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 2865, Remote Authentication Dial In User Service (RADIUS) is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 2866, RADIUS Accounting is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 3046, DHCP Relay Agent Information Option (Option 82) is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 3596, DNS Extensions to Support IP version 6 is supported on D, E, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 3768, Virtual Router Redundancy Protocol (VRRP) is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 4250, The Secure Shell (SSH) Protocol Assigned Numbers is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 4251, The Secure Shell (SSH) Protocol Architecture is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 4254, The Secure Shell (SSH) Connection Protocol is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 4632, Classless Inter-domain Routing (CIDR): The Internet Address Assignment and Aggregation Plan is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

- RFC 5880, Bidirectional Forwarding Detection (BFD) is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 5881, Bidirectional Forwarding Detection (BFD) IPv4 and IPv6 (Single Hop) is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 5883, Bidirectional Forwarding Detection (BFD) for Multihop Paths is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 6528, Defending against Sequence Number Attacks is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

### **IP — Multicast**

- RFC 1112, Host Extensions for IP Multicasting is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 2236, Internet Group Management Protocol, Version 2 is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 3306, Unicast-Prefix-based IPv6 Multicast Addresses is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 3376, Internet Group Management Protocol, Version 3 is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 3446, Anycast Rendezvous Point (RP) mechanism using Protocol Independent Multicast (PIM) and Multicast Source Discovery Protocol (MSDP) is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 4601, Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised) is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 4604, Using Internet Group Management Protocol Version 3 (IGMPv3) and Multicast Listener Discovery Protocol Version 2 (MLDv2) for Source-Specific Multicast is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 4607, Source-Specific Multicast for IP is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 4608, Source-Specific Protocol Independent Multicast in 232/8 is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 4610, Anycast-RP Using Protocol Independent Multicast (PIM) is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 5059, Bootstrap Router (BSR) Mechanism for Protocol Independent Multicast (PIM) is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 5384, The Protocol Independent Multicast (PIM) Join Attribute Format is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 6513, Multicast in MPLS/BGP IP VPNs is supported on T(N), Mxp, R6, and R12
- RFC 6514, BGP Encodings and Procedures for Multicast in MPLS/IP VPNs is supported on T(N), Mxp, R6, and R12

RFC 6515, IPv4 and IPv6 Infrastructure Addresses in BGP Updates for Multicast VPNs is supported on T(N), Mxp, R6, and R12

RFC 6625, Wildcards in Multicast VPN Auto-Discover Routes is supported on T(N), Mxp, R6, and R12

RFC 6826, Multipoint LDP In-Band Signaling for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Path is supported on T(N), Mxp, R6, and R12

RFC 7385, IANA Registry for P-Multicast Service Interface (PMSI) Tunnel Type Code Points is supported on T(N), Mxp, R6, and R12

## **IP — Version 4**

RFC 791, Internet Protocol is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 792, Internet Control Message Protocol is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 826, An Ethernet Address Resolution Protocol is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 1519, Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 1812, Requirements for IPv4 Routers is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 1981, Path MTU Discovery for IP version 6 is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 2401, Security Architecture for Internet Protocol is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 2460, Internet Protocol, Version 6 (IPv6) Specification is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

## **IP — Version 6**

RFC 2464, Transmission of IPv6 Packets over Ethernet Networks is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 3021, Using 31-Bit Prefixes on IPv4 Point-to-Point Links is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 3122, Extensions to IPv6 Neighbor Discovery for Inverse Discovery Specification is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 3587, IPv6 Global Unicast Address Format is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

- RFC 4007, IPv6 Scoped Address Architecture is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 4193, Unique Local IPv6 Unicast Addresses is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 4291, Internet Protocol Version 6 (IPv6) Addressing Architecture is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 4443, Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 4861, Neighbor Discovery for IP version 6 (IPv6) is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 4862, IPv6 Stateless Address Autoconfiguration (Router Only) is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 5095, Deprecation of Type 0 Routing Headers in IPv6 is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 5952, A Recommendation for IPv6 Address Text Representation is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 6106, IPv6 Router Advertisement Options for DNS Configuration is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 6164, Using 127-Bit IPv6 Prefixes on Inter-Router Links is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

### **IPsec**

- RFC 2401, Security Architecture for the Internet Protocol is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 2406, IP Encapsulating Security Payload (ESP) is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

### **IS-IS**

- draft-ietf-isis-mi-02, IS-IS Multi-Instance is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- draft-kaplan-isis-ext-eth-02, Extended Ethernet Frame Size Support is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- ISO/IEC 10589:2002, Second Edition, Nov. 2002, Intermediate system to Intermediate system intra-domain routing information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode Network Service (ISO 8473) is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 1195, Use of OSI IS-IS for Routing in TCP/IP and Dual Environments is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

- RFC 3359, Reserved Type, Length and Value (TLV) Codepoints in Intermediate System to Intermediate System is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 3719, Recommendations for Interoperable Networks using Intermediate System to Intermediate System (IS-IS) is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 3787, Recommendations for Interoperable IP Networks using Intermediate System to Intermediate System (IS-IS) is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 4971, Intermediate System to Intermediate System (IS-IS) Extensions for Advertising Router Information is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 5120, M-ISIS: Multi Topology (MT) Routing in IS-IS is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 5130, A Policy Control Mechanism in IS-IS Using Administrative Tags is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 5301, Dynamic Hostname Exchange Mechanism for IS-IS is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 5302, Domain-wide Prefix Distribution with Two-Level IS-IS is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 5303, Three-Way Handshake for IS-IS Point-to-Point Adjacencies is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 5304, IS-IS Cryptographic Authentication is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 5305, IS-IS Extensions for Traffic Engineering TE is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 5306, Restart Signaling for IS-IS (Helper Mode) is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 5308, Routing IPv6 with IS-IS is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 5309, Point-to-Point Operation over LAN in Link State Routing Protocols is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 5310, IS-IS Generic Cryptographic Authentication is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 6232, Purge Originator Identification TLV for IS-IS is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 6233, IS-IS Registry Extension for Purges is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

## Management

- draft-ietf-snmpv3-update-mib-05, Management Information Base (MIB) for the Simple Network Management Protocol (SNMP) is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- draft-ietf-idr-bgp4-mib-05, Definitions of Managed Objects for the Fourth Version of Border Gateway Protocol (BGP-4) is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- draft-ietf-isis-wg-mib-06, Management Information Base for Intermediate System to Intermediate System (IS-IS) is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- draft-ietf-mpls-ldp-mib-07, Definitions of Managed Objects for the Multiprotocol Label Switching, Label Distribution Protocol (LDP) is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- draft-ietf-mpls-lsr-mib-06, Multiprotocol Label Switching (MPLS) Label Switching Router (LSR) Management Information Base Using SMIV2 is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- draft-ietf-mpls-te-mib-04, Multiprotocol Label Switching (MPLS) Traffic Engineering Management Information Base is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- draft-ietf-ospf-mib-update-08, OSPF Version 2 Management Information Base is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- ianaaddressfamilynumbers-mib, IANA-ADDRESS-FAMILY-NUMBERS-MIB is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- ianaiftype-mib, IANAifType-MIB is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- ianaiprouteprotocol-mib, IANA-RTPROTO-MIB is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- IEEE8021-CFM-MIB, IEEE P802.1ag(TM) CFM MIB is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- IEEE8021-PAE-MIB, IEEE 802.1X MIB is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- IEEE8023-LAG-MIB, IEEE 802.3ad MIB is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- LLDP-MIB, IEEE P802.1AB(TM) LLDP MIB is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 1157, A Simple Network Management Protocol (SNMP) is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 1215, A Convention for Defining Traps for use with the SNMP is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12



- RFC 1724, RIP Version 2 MIB Extension is supported on Mxp
- RFC 2021, Remote Network Monitoring Management Information Base Version 2 using SMIV2 is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 2115, Management Information Base for Frame Relay DTEs Using SMIV2 is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 2138, Remote Authentication Dial In User Service (RADIUS) is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 2206, RSVP Management Information Base using SMIV2 is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 2213, Integrated Services Management Information Base using SMIV2 is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 2494, Definitions of Managed Objects for the DS0 and DS0 Bundle Interface Type is supported on M(N)
- RFC 2571, An Architecture for Describing SNMP Management Frameworks is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 2572, Message Processing and Dispatching for the Simple Network Management Protocol (SNMP) is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 2573, SNMP Applications is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 2574, User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3) is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 2575, View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP) is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 2578, Structure of Management Information Version 2 (SMIV2) is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 2579, Textual Conventions for SMIV2 is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 2787, Definitions of Managed Objects for the Virtual Router Redundancy Protocol is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 2819, Remote Network Monitoring Management Information Base is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 2856, Textual Conventions for Additional High Capacity Data Types is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 2863, The Interfaces Group MIB is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

- RFC 2864, The Inverted Stack Table Extension to the Interfaces Group MIB is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 2933, Internet Group Management Protocol MIB is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 3014, Notification Log MIB is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 3164, The BSD syslog Protocol is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 3165, Definitions of Managed Objects for the Delegation of Management Scripts is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 3231, Definitions of Managed Objects for Scheduling Management Operations is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 3273, Remote Network Monitoring Management Information Base for High Capacity Networks is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 3416, Version 2 of the Protocol Operations for the Simple Network Management Protocol (SNMP) is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 3417, Transport Mappings for the Simple Network Management Protocol (SNMP) (SNMP over UDP over IPv4) is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 3419, Textual Conventions for Transport Addresses is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 3584, Coexistence between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 3593, Textual Conventions for MIB Modules Using Performance History Based on 15 Minute Intervals is supported on K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 3635, Definitions of Managed Objects for the Ethernet-like Interface Types is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 3826, The Advanced Encryption Standard (AES) Cipher Algorithm in the SNMP User-based Security Model is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 3877, Alarm Management Information Base (MIB) is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 3895, Definitions of Managed Objects for the DS1, E1, DS2, and E2 Interface Types is supported on M(N)
- RFC 4001, Textual Conventions for Internet Network Addresses is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

- RFC 4022, Management Information Base for the Transmission Control Protocol (TCP) is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 4113, Management Information Base for the User Datagram Protocol (UDP) is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 4220, Traffic Engineering Link Management Information Base is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 4292, IP Forwarding Table MIB is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 4293, Management Information Base for the Internet Protocol (IP) is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 6241, Network Configuration Protocol (NETCONF) is supported on K5, K12, R6, and R12
- RFC 6242, Using the NETCONF Protocol over Secure Shell (SSH) is supported on K5, K12, R6, and R12

## **MPLS — General**

- RFC 3031, Multiprotocol Label Switching Architecture is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 3032, MPLS Label Stack Encoding is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 3443, Time To Live (TTL) Processing in Multi-Protocol Label Switching (MPLS) Networks is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 4182, Removing a Restriction on the use of MPLS Explicit NULL is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 5332, MPLS Multicast Encapsulations is supported on T(N), Mxp, R6, and R12

## **MPLS — GMPLS**

- draft-ietf-ccamp-rsvp-te-srlg-collect-04, RSVP-TE Extensions for Collecting SRLG Information is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

## **MPLS — LDP**

- draft-pdutta-mpls-ldp-adj-capability-00, LDP Adjacency Capabilities is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- draft-pdutta-mpls-ldp-v2-00, LDP Version 2 is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- draft-pdutta-mpls-tldp-hello-reduce-04, Targeted LDP Hello Reduction is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

- RFC 3037, LDP Applicability is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 3478, Graceful Restart Mechanism for Label Distribution Protocol (Helper Mode) is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 5036, LDP Specification is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 5283, LDP Extension for Inter-Area Label Switched Paths (LSPs) is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 5443, LDP IGP Synchronization is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 5561, LDP Capabilities is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 6388, Label Distribution Protocol Extensions for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

### **MPLS — MPLS-TP**

- RFC 5586, MPLS Generic Associated Channel is supported on T(N), R6, and R12
- RFC 5921, A Framework for MPLS in Transport Networks is supported on T(N), R6, and R12
- RFC 5960, MPLS Transport Profile Data Plane Architecture is supported on T(N), R6, and R12
- RFC 6370, MPLS Transport Profile (MPLS-TP) Identifiers is supported on T(N), R6, and R12
- RFC 6378, MPLS Transport Profile (MPLS-TP) Linear Protection is supported on T(N), R6, and R12
- RFC 6426, MPLS On-Demand Connectivity and Route Tracing is supported on T(N), R6, and R12
- RFC 6428, Proactive Connectivity Verification, Continuity Check and Remote Defect indication for MPLS Transport Profile is supported on T(N), R6, and R12
- RFC 6478, Pseudowire Status for Static Pseudowires is supported on T(N), R6, and R12
- RFC 7213, MPLS Transport Profile (MPLS-TP) Next-Hop Ethernet Addressing is supported on T(N), R6, and R12

### **MPLS — OAM**

- RFC 6424, Mechanism for Performing Label Switched Path Ping (LSP Ping) over MPLS Tunnels is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 6425, Detecting Data Plane Failures in Point-to-Multipoint Multiprotocol Label Switching (MPLS) - Extensions to LSP Ping is supported on T(N), Mxp, R6, and R12

## MPLS — RSVP-TE

- RFC 2702, Requirements for Traffic Engineering over MPLS is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 2747, RSVP Cryptographic Authentication is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 2961, RSVP Refresh Overhead Reduction Extensions is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 3097, RSVP Cryptographic Authentication -- Updated Message Type Value is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 3209, RSVP-TE: Extensions to RSVP for LSP Tunnels is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 3477, Signalling Unnumbered Links in Resource ReSerVation Protocol - Traffic Engineering (RSVP-TE) is supported on M(N), T(N), X, Mxp, R6, and R12
- RFC 4090, Fast Reroute Extensions to RSVP-TE for LSP Tunnels is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 4561, Definition of a Record Route Object (RRO) Node-Id Sub-Object is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 4875, Extensions to Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label Switched Paths (LSPs) is supported on T(N), Mxp, R6, and R12
- RFC 4950, ICMP Extensions for Multiprotocol Label Switching is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 5712, MPLS Traffic Engineering Soft Preemption is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 5817, Graceful Shutdown in MPLS and Generalized MPLS Traffic Engineering Networks is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

## OSPF

- draft-ietf-ospf-prefix-link-attr-06, OSPFv2 Prefix/Link Attribute Advertisement is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 1765, OSPF Database Overflow is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 2328, OSPF Version 2 is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 3101, The OSPF Not-So-Stubby Area (NSSA) Option is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 3509, Alternative Implementations of OSPF Area Border Routers is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

- RFC 3623, Graceful OSPF Restart Graceful OSPF Restart (Helper Mode) is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 3630, Traffic Engineering (TE) Extensions to OSPF Version 2 is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 4222, Prioritized Treatment of Specific OSPF Version 2 Packets and Congestion Avoidance is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 4552, Authentication/Confidentiality for OSPFv3 is supported on M(N), T(N), X, Mxp, R6, and R12
- RFC 4576, Using a Link State Advertisement (LSA) Options Bit to Prevent Looping in BGP/MPLS IP Virtual Private Networks (VPNs) is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 4577, OSPF as the Provider/Customer Edge Protocol for BGP/MPLS IP Virtual Private Networks (VPNs) is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 4970, Extensions to OSPF for Advertising Optional Router Capabilities is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 5185, OSPF Multi-Area Adjacency is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 5187, OSPFv3 Graceful Restart (Helper Mode) is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 5243, OSPF Database Exchange Summary List Optimization is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 5250, The OSPF Opaque LSA Option is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 5309, Point-to-Point Operation over LAN in Link State Routing Protocols is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 5340, OSPF for IPv6 is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 5838, Support of Address Families in OSPFv3 is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 6987, OSPF Stub Router Advertisement is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

### **Pseudowire**

- draft-ietf-l2vpn-vpws-iw-oam-04, OAM Procedures for VPWS Interworking is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 3916, Requirements for Pseudo- Wire Emulation Edge-to-Edge (PWE3) is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

- RFC 3985, Pseudo Wire Emulation Edge-to-Edge (PWE3) is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 4385, Pseudo Wire Emulation Edge-to-Edge (PWE3) Control Word for Use over an MPLS PSN is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 4446, IANA Allocations for Pseudowire Edge to Edge Emulation (PWE3) is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 4447, Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP) is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 4448, Encapsulation Methods for Transport of Ethernet over MPLS Networks is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 5659, An Architecture for Multi-Segment Pseudowire Emulation Edge-to-Edge is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 6073, Segmented Pseudowire is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 6310, Pseudowire (PW) Operations, Administration, and Maintenance (OAM) Message Mapping is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 6391, Flow-Aware Transport of Pseudowires over an MPLS Packet Switched Network is supported on Mxp, R6, and R12
- RFC 6718, Pseudowire Redundancy is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 6870, Pseudowire Preferential Forwarding Status bit is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 7023, MPLS and Ethernet Operations, Administration, and Maintenance (OAM) Interworking is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 7267, Dynamic Placement of Multi-Segment Pseudowires is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

## Quality of Service

- RFC 2430, A Provider Architecture for Differentiated Services and Traffic Engineering (PASTE) is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 2474, Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 2598, An Expedited Forwarding PHB is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12
- RFC 3140, Per Hop Behavior Identification Codes is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 3260, New Terminology and Clarifications for Diffserv is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

### **RIP**

RFC 1058, Routing Information Protocol is supported on Mxp

RFC 2082, RIP-2 MD5 Authentication is supported on Mxp

RFC 2453, RIP Version 2 is supported on Mxp

### **Timing**

GR-1244-CORE, Clocks for the Synchronized Network: Common Generic Criteria, Issue 3, May 2005 is supported on D-ETR, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

GR-253-CORE, SONET Transport Systems: Common Generic Criteria. Issue 3, September 2000 is supported on D-ETR, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

IEEE 1588-2008, IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems is supported on D-ETR, K5, K12, M(A,N), T(A,N), X, Mxp, Sx-1/10GE, R6, and R12

ITU-T G.781, Synchronization layer functions, issued 09/2008 is supported on D-ETR, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

ITU-T G.813, Timing characteristics of SDH equipment slave clocks (SEC), issued 03/2003 is supported on D-ETR, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

ITU-T G.8261, Timing and synchronization aspects in packet networks, issued 04/2008 is supported on D-ETR, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

ITU-T G.8262, Timing characteristics of synchronous Ethernet equipment slave clock (EEC), issued 08/2007 is supported on D-ETR, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

ITU-T G.8264, Distribution of timing information through packet networks, issued 10/2008 is supported on D-ETR, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

ITU-T G.8265.1, Precision time protocol telecom profile for frequency synchronization, issued 10/2010 is supported on D-ETR, K5, K12, M(A,N), T(A,N), X, Mxp, Sx-1/10GE, R6, and R12

ITU-T G.8275.1, Precision time protocol telecom profile for phase/time synchronization with full timing support from the network, issued 07/2014 is supported on X, Mxp, R6, and R12

RFC 5905, Network Time Protocol Version 4: Protocol and Algorithms Specification is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12



## VPLS

RFC 4761, Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4762, Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 5501, Requirements for Multicast Support in Virtual Private LAN Services is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 6074, Provisioning, Auto-Discovery, and Signaling in Layer 2 Virtual Private Networks (L2VPNs) is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12



# Index

## C

continuity check 157  
CPE ping 103

## D

DDMAP 86

## E

Ethernet CFM 136

## L

linktrace 155  
loopback 154  
LSP diagnostics 82

## M

MAC ping 101  
MAC populate 104  
MAC purge 104  
MAC trace 102  
Mirror  
    overview 18  
        implementation 19  
        source and destination 20  
    configuring  
        basic 34  
        classification rules 35  
            IP filter 36  
            MAC filter 36  
            port 35  
            SAP 35  
    command reference 51  
    local mirror service 38  
    management tasks 44  
    overview 32  
    remote mirror service 40

## O

OAM 81  
    overview 81  
    configuring  
        command reference 217

## P

ping  
    VCCV 106

## S

SAA test parameters 177  
SDP diagnostics 98  
SDP ping 98  
service assurance agent 176, 179  
service diagnostics 100

## T


Tools 443

## V

VCCV ping 106  
VLL diagnostics 106  
VPLS MAC diagnostics 101



# Customer Document and Product Support



## Customer documentation

[Customer Documentation Welcome Page](#)



## Technical Support

[Product Support Portal](#)



## Documentation feedback

[Customer Documentation Feedback](#)

